What's In a Message? Interpreting Geo-referenced Data for the Visually-impaired

Kavita E. Thomas

Department of Computing Science
University of Aberdeen
Aberdeen, Scotland
tkavita@abdn.ac.uk

Somayajulu Sripada

Department of Computing Science
University of Aberdeen
Aberdeen, Scotland
yaji.sripada@abdn.ac.uk

Abstract

In this paper we describe content determination issues involved in the Atlas.txt project, which aims to automatically describe geo-referenced information such as census data as text for the visually-impaired (VI). Texts communicating geo-referenced census information contain census data abstractions and their corresponding geographic references. Because visually impaired users find interpreting geographic references hard, we hypothesized that an introduction message about the underlying geography should help the users to interpret the geographic references easily. We performed user studies to design and evaluate the introduction message. An initial evaluation study with several sighted users and one partially sighted user showed that an introduction message is certainly preferred by most participants. Many of them used an introduction message themselves when they described maps textually. But the study also showed that the introduction message made no difference when the participants were asked to draw maps using the information in the textual descriptions.

1 Introduction

Atlas.txt is a data-to-text natural language generation (NLG) system which aims to communicate geo-referenced information to VI people. Geo-referenced data are data which has a geographic component, and is distributed over a region on Earth. Such data is often visualised as thematic maps, as shown in Figure 1, where colour shading indicates

density of distribution, so that dark or intense regions on such maps indicate where the variable, for example unemployment, is most frequently located. Geo-referenced data is frequently stored in long data tables which hinder VI users, who must rely on screenreaders to read out the values in the data table, from quickly and easily forming an overview of how the data is distributed. For example, although the table corresponding to Figure 1 only has 31 rows (corresponding to the 31 council areas of Scotland), VI users still need to listen to each of these values and remember them in order to form a mental picture of the overall trends in population distribution which sighted users can infer from the map in under a minute. Furthermore, tables do not communicate where regions are located. This means that although they might be able to pick out maxima and minima, inferring trends is a much harder task, and mentally visualising how the data is distributed is impossible for unfamiliar areas given named locations in tables. Atlas.txt aims to automatically interpret this tabular geo-referenced data and communicate its central information as texts which can then be read out to VI users via a screenreading application. We start by describing some of the issues involved in communicating geo-referenced information in the census domain. Then we address content determination issues in the system and present some user needs requirements for communicating geo-referenced information to VI users, principle among which is the need for texts to enable mentally visualising the region and how the data is distributed. This motivates our decision to include introduction messages which describe the region before communicating how the

Figure 1: A typical map showing geo-referenced data

data is distributed, and present an experiment into this choice, focusing on evaluating whether introduction messages are beneficial.

2 Related Work

Systems which generate descriptions of numerical data are not uncommon, e.g., the FOG system (Goldberg et al, 1994), TREND (Boyd, 1997) and MULTI-METEO (Coch, 1998) all generate textual summaries of numerical weather data. (Ferres et al, 2006) generates textual descriptions of information in graphs and enables user querying. SumTime¹ summarises time-series data and RoadSafe² generates travel advisories and looks at spatial time-series data. While there is no prior work on generating textual descriptions of geo-referenced data, there have been studies on describing spatial data in the context of route directions ((Geldof, 2003), (Marciniak and Strube, 2004)), scene descriptions (Novak, 1986), geometric descriptions (Mitkov, 1990) and spatial descriptions (Ebert et al, 1996).

Data-to-text generation projects like Atlas.txt differ from the traditional three-stage pipeline generation architecture of most NLG systems (Reiter and Dale, 2000) because they need to analyse the data in order to determine data abstractions before then going on to the traditional three stages of document planning, micro-planning and realisation, as was put forward by (Sripada, Reiter et al, 2001). Content determination has been described as involving a collective classification problem (Barzilay and Lapata,

2005) and in the context of data-to-text systems. Since our system involves data-to-text NLG, content determination has more in common with the two-stage process of (Sripada, Reiter et al, 2001) than the classification process of (Barzilay and Lapata, 2005). However a notable difference between our domain and the meteorology domain described in (Sripada, Reiter et al, 2001), is that for us, arriving at an overview before communicating the data to the end-user involves not just data analysis but also scene description to aid in our visualisation goal.

3 Content Determination in Atlas.txt System

The Atlas.txt system aims to be able to take geo-referenced data in tabular form, perform data analysis to determine significant features in the data, and then communicate these significant features via text. Data analysis for Atlas.txt involves clustering to find regions in the geography with similar values of census variable. Data analysis also involves trending to find spatial trends in the data. Information obtained from data analysis forms the input to the content determination module.

In order to select which information to communicate (i.e., content selection), we ran a corpus study on online census texts to try and determine what messages they communicate (Thomas and Sripada, 2008). Of course, these documents are not written for the VI user group. On the other hand, without evidence to the contrary, we cannot assume that VI users require census information (i.e., census statistics) that is different from what sighted users require. But VI users do need additional content that helps them to interpret any geographic references associated with census information. An experiment conducted in Ottawa on a group of 5 blind and 5 sighted participants indicated that both blind and sighted participants (without access to a map) need additional information to ground proper name reference to unfamiliar places. Because we did not find examples of these kind in our corpus texts, we hypothesised that all texts produced by Atlas.txt system should start with what we called an 'introduction' messages which help VI users to interpret geographic information better.

This means, we need knowledge of two kinds of

¹More information can be found at <http://www.csd.abdn.ac.uk/research/sumtime>.

²See <http://www.csd.abdn.ac.uk/research/roadsafe>.

messages for building Atlas.txt, ones that communicate census information and the others that help VI users understand the geography better. Designing the first type of census messages requires a good understanding of user goals. We have noticed a broad range of goals census users have, and we consider a number of use cases illustrating particular user goals below:

1. You want to buy a house somewhere safe, i.e., with a low crime rate, and within commuting distance of City X.
2. You want to start a business selling expensive furniture and you want to pick a location near potential clientele who are profiled as wealthy based on income level data.

This illustrates the huge range of purposes people might have when accessing census data, and many of these examples were found via online documents on census uses.³ We determined which content to communicate from the data in this domain by considering both evidence from use cases like the ones above and also online census documents. In many online census documents there is no corresponding data available in tabular form. This could be because the information in different modalities is intended to be complimentary rather than equivalent. The lack of corresponding data in most cases makes it hard to learn what sorts of results constitute interesting data which should be reported. From a qualitative study of corpus documents, we notice that most such texts at the least describe where the variable they focus on (e.g., unemployment) has its high values, low values, and spatial trends of the variable's distribution if there are any. Some scenarios require additional information, e.g., the vicinity of a major road, the coast, a city, etc. We aim for a scenario-generic solution for now as a first step given that the system is not interactive. However we hope that communicating the basic statistics in a meaningful way (i.e., via useful spatial expressions) and also introducing the region and its salient features will improve accessibility to this information. The goal of such a system

is of course to eventually become interactive so that it can eventually address users' specific goals and answer more specific questions about how the data is distributed. As explained above Atlas.txt includes the following census messages:

- Maxima message: location/s of high values from Data Analysis
- Minima message: location/s of low values from Data Analysis
- Trend message: location/s of trends from Data Analysis

All these census messages make explicit references to geographic locations which are hard to interpret for VI users. We assume that the introduction messages introduced earlier enable hearers to visualise the underlying geography which helps them to make sense of spatial descriptions in census messages. The introduction message also gives a general summary by communicating the overall average value-range of the data (e.g., "Overall, unemployment frequencies are low"), inspired by Gestalt principles of grouping in perception (Metzger, 2006) and the visual information seeking mantra (Schneiderman, 1996) which argues that one should start with an overview, zoom and filter, and only provide details upon demand (the latter of which is only applicable to interactive systems). The following components of introduction messages are currently considered:

- Boundary geographic information, which indicates geographic features which bound the region as a whole on any side
- Gestalt-driven overview of data (average value-range)
- Location of where population is densest

In order to communicate the boundary geographic information, we ran a survey eliciting spatial descriptions of Aberdeenshire which indicated that people's spatial descriptions when presented with a census map seem to fall into the following three categories of geographic information:

- Human geography: cities, roads, towns, urban/rural distinctions, census ward boundaries, etc.

³Business Uses of Census Data: www.incontext.indiana.edu/2001/july01/spotlight.html, Uses of Census Data: www.laria.gov.uk/content/articles/2003_conference_pdf/boag.pdf

- Physical geography: lakes, mountains, sea, desert, etc.
- Orientation geography: cardinal directions

In other words, we saw many descriptions which used terms like “urban”, “rural”, which pointed out major roads as indicators of where population might be denser, and also many of the descriptions which referred to areas as “coastal”, “inland”, “[in the] hills”, etc. Cardinal directions (e.g., “North”, “Southwest”, etc.) were also frequently used to locate phenomena and are absolute coordinates in the survey frame of reference involved in Atlas.txt, which describes maps from a bird’s eye view, (Taylor and Tversky, 1996).

Determining salient geographic information which describes where maxima, minima and trends are involves overlaying the clusters (and ends of trends) with geographic layers in the human and physical geographies discussed above. Then if we find that any of the geographic elements contained in any of the geographic overlays aligns spatially with the cluster, we add this to a list of multiple possible spatial expressions which can be used to describe the location of the cluster. Appropriately describing clusters via selecting informative spatial expressions is a central aspect of communicating the maximum, minimum and trend messages. Overlaying different layers like this is a standard “intersect overlay” operation in Geographic Information Systems (GIS) and data containing altitude information, urban/road information and physical geography exist online which can be used as the alignment layers in a GIS. We want to select features which intersect our clusters even minimally, as they help to describe where features are located. If no intersecting regions exist, we choose the nearest. Note that the layers considered all fall within the different geographic types, i.e., human and physical geographies. Orientation information does not require overlaying maps to identify which parts are North, etc. Orientation information can be calculated for the region in question by overlaying it with a square 8 by 8 grid for the various cardinal directions. The location of the maxima and minima clusters can then be described as lying in a particular cardinal direction by simply checking which square/s of the grid they

fall in. Determining cardinal direction has been implemented in the Roadsafe project.

In future we will consider how to handle situations in which multiple possible features intersect (or are near to) our clusters, which will involve us ranking some of these features higher than others. We will also want to consider the effects of other factors than proximity on selection of spatial descriptions, e.g., discourse context and causality could affect the choice of spatial description. The results from our survey eliciting spatial descriptions of Aberdeen showed more spatial descriptions from the human geography (58%) than from the physical geography (42%). The survey involved a between-subjects design where the 24 participants were blocked into 3 groups, one group getting coarse maps with only 3 or 4 large wards covering all of Aberdeenshire while the other two groups had medium (around 28 wards) and fine grained maps (around 50 wards) covering the shire. Coarse-grained maps had higher percentages of scores for physical geography descriptions than for human geography descriptions (62% vs. 37%) while both medium and fine-grained maps had larger percentages for human geography descriptions (62% and 57% respectively). There were also a higher number of human descriptions used per turn on average, 4.2 vs. 2.6 for physical descriptions.

Most census variables describe information which pertains to human populations, (e.g., unemployment statistics, or employment in hunting, agriculture and forestry in our survey), indicating that human geographic features might be causally linked to most census data. As was indicated in our initial survey, human geographic features like the locations of cities and roads seem to be most related to these population variables and might form the most meaningful description for the hearer, since they may help the hearer to infer useful correlations between the presence of (e.g.) a maxima and a geographic feature. However there might also be correlations between physical geographic features and these variables, for example, logging in forests, fishing industry in coastal areas, etc.

Context will certainly also play a large role, e.g., if the prior message uses a description from the physical geography (e.g., “coastal”), this might result in “inland” being preferred over “urban”, par-

ticularly if the messages have opposite content (e.g., one reports maxima while the other reports minima).

We also saw that participants tended to mix geographies more often than not, involving both physical and human geography descriptions in their texts, as 60% of the number of turns which had descriptions from both geographies. Describing cluster locations with two geographies leads to sentences like “Unemployment is highest in *rural* and *coastal* areas,” which leads to a range of inferences about how coastal and rural areas might be similar where unemployment is concerned. Mixed geography descriptions are more informative, since they convey additional information, so our view is that extra descriptors, particularly if they come from other geographies, will only clarify and help listeners to ground the locations of clusters.

So the content determination module takes the document plan containing the four messages to be communicated and gets the locations of maxima, minima and trends (which come from data analysis). These are then overlaid with data of the area containing physical or human geographies for the region. Proximity is our current selection factor, so aligning features are currently selected if they intersect (though as discussed earlier, if no features intersect, the nearest ones are chosen) and are saved as possible spatial expressions to be used in the message. In future work we will also need to account for causality and context in this selection of appropriate spatial descriptions.

4 Evaluation

Unlike many of the other data-to-text systems described, the Atlas.txt system’s goal is much more general. Rather than generating specific texts for specific audiences with particular expertise and tasks in mind, this project aims to communicate census data to a VI audience of varying expertise who, unlike for most other NLG systems, have a range of tasks in mind. As we argued earlier, the messages we choose to communicate enable the vast majority of users to at least get an overview of the salient features in the data and we leave more specific informational requirements for an interactive version of Atlas.txt.

One issue that arises for a broad-focus non-task-

specific domain like this is that evaluation methods are not as well-specified as for other NLG systems for content determination. (Dale and Mellish, 1998) present four questions one can ask when evaluating content determination: (1) truth (is the output true), (2) quantity—does the output say enough, (3) quantity—does the output say too much, and (4) appropriateness (is the output inappropriate to the context). But what if we want to ask the more high level question of utility? That is, do we really need the introduction message (assuming it is not misleading)? The two quantity considerations—have we said too much or too little—are probably the best ways to gauge whether the introduction message is informative enough (and not too verbose). But again, this is a bit too fine-grained for us, as we want to first ask whether the introduction message is necessary or not in the first place before attempting to refine its contents.

The general-purpose task aspect of this domain also implies that we cannot compare our texts with an expert’s texts based on the underlying data, as it is difficult to find simple summarising expert-written census texts online. The vast majority of census texts are far longer than ours and tend to discuss the implications of what the data presents which is beyond our scope. Furthermore, we cannot compare Atlas.txt output against expert-written text because there is no expert-produced parallel data where texts summarise the information contained in tables or maps. Also, experts (statisticians) are not used to summarising tables as simple, general-purpose texts like this and they vary in the sorts of texts they produce. Our experiment in Ottawa indicated that different statisticians employed at Statistics Canada vary both in terms of the analyses they make of both tabular and textual census data (under experimental conditions) and in their post-edits of expert-written census texts (Thomas and Sripada, 2008). This means that we need to rethink how we should evaluate Atlas.txt’s texts.

4.1 Study

Given that instructions rich in imagery help subjects to form a spatial mental model (Gyselinck et al, 2006), our end goal is to discover if introduction messages actually aid in visualisation. Given the lack of conclusive evidence to the contrary, we

operate under the premise that VI people have spatial comprehension abilities (Kitchin et al, 1997) and can benefit from texts which enable visualisation. We run the study with sighted people because they can draw maps which they visualise from our texts.

We ran a study which presented a group of 20 randomly recruited sighted participants with a pair of texts where one member of the pair included the introduction message and the other member did not include the introduction message. The introduction message contains the two geographic elements we discussed earlier: locations of boundary information, the location of where population is densest and the trend for population density. The messages were identical in content (though names and directions were changed so that participants would not assume that the same country was being described), and both texts have a maxima, minima and trend message (each), as can be seen in Examples 1 and 2 below:

1. The land of Bosk is bounded to the East by the Bale Ocean, and its largest city is Wristwath, on the East coast. Population is densest around Wristwath and gradually decreases as one moves Westward, though there are clusters of medium-high population in cities like Slangin in Central Bosk. Crime rates are highest in Wristwath and Southern suburbs of Wristwath. Additionally, Slangin also has high crime rates. Crime is lowest in Northwestern and East-Central Bosk. Crime tends to decrease as one goes East or West from Central Bosk with the exception of the Wristwath area.
2. Poverty rates are highest in Galen and Western suburbs of Galen. Additionally, Boroos also has high poverty rates. Poverty rates are lowest in Northwestern and East-Central Dagor. Poverty rates tends to decrease as one goes East or West from Central Dagor with the exception of the Galen area.

For each text, participants were asked to draw a map of the area they visualise based on the text, indicating where the maxima and minima are and where all other information mentioned in the text is located. Then they were asked two task-based questions per text which aim to check whether they can use the

information appropriately. The first question asks whether they can identify the minima/maxima, e.g., “You want to relocate to Bosk. Considering only crime rates, where might you want to move?” The second question gives them a piece of additional information which sighted people would be able to pick off from the map, e.g., that there is a coastline, harbours, roads, etc., e.g., “You work in Wristwath and need to commute in daily. Given this and the above info, where might you want to live in Bosk?” After each of the two pairs of texts and questions, participants are asked to indicate which text was easier to visualise and to state why. Participants were also presented with two thematic maps and asked to write descriptions of them so that someone who is blind can visualise them.

We hypothesized that Ex.2 above will cause difficulties for map drawing. There is no introduction message helping readers to orient themselves to the region, and one does not know a priori whether Galen is the name of a city or the region itself, or where it is in the region if it is indeed a city. This might seem purposefully difficult, but this is what one typically finds in census texts online. For example, if you want to know where in Canada Saskatoon is, you need to look it up in a map. This makes it hard (even for sighted people) when you read many unfamiliar place names and are trying to get an overall picture of where places are.

4.2 Results and Discussion

The maps drawn for the text stimuli are scored on a number of features, but the primary goal is to ascertain whether readers visualised the locations of the maximum, minimum and trend correctly. Each text is based on an underlying map which is used as the gold standard against which participants’ maps are compared. This gold standard map includes a range of variation in location of maximum, minimum and trend based on the ambiguity present in the text. For example, “East-Central” in Ex. 2 is hard to define. Participants were given a point each for correctly identifying the maximum, minimum and trend in each message, or a zero if they failed to do so correctly or if their answer was uninterpretable. These three scores were also totalled giving a total score per text for the maps drawn. Here we did not see a huge difference in scores for the different mes-

Table 1: Average & std. deviation scores for map drawing

Text	Maximum	Minimum	Trend	Total
Text 1	1; 0	0.80; 0.41	0.80; 0.41	0.87; 0.23
Text 2	0.85; 0.37	0.70; 0.47	0.65; 0.49	0.73; 0.32

sages, as can be seen in the table; Text 1 has an introduction message (and is seen in Ex. 1) and Text 2 (Ex. 2) lacks the introduction message. The total column corresponds to the average score for the three messages (maximum, minimum and trend). A t-test on the total scores for Texts 1 and 2 gives a non-significant result with a p -value of 0.110. On the other hand, almost all the participants preferred Text 1 to Text 2; a signed ranks test on the preference judgements (participants ticked which text they preferred, 1 or 2) was significant with a p -value < 0.001 . When asked to state their reasons for preferring Text 1, participants identified the geographic information from the introduction message as the only reason. Also when participants wrote texts for map stimuli they tended to include introduction messages in them; a binomial test on whether introduction messages were included in either of their two texts was significant with a p -value < 0.041 . Interestingly enough, shape was a feature used in many of these introduction messages when the map had a particularly non-square shape. In the part of the experiment where participants answered task-based questions based on texts with and without introduction messages they performed well on all questions following both types of texts.

It is encouraging that participants preferred texts with introduction messages and also used them in their own texts. But it is disappointing that the absence of introduction messages seems to make no difference to mental visualisation of Text 2. In fact, many of the maps for Text 2 were incomplete; participants simply marked the required information and their nearby geographic entities. The problem with evaluating maps is that they partly measure the geometrical relationships in the map which could bias such evaluations against participants with poor drawing skills. Poor drawing skills could also obscure whether they have visualised the correct information.

Many participants said that drawing the map and accounting for all of the information contained in

the introduction message was very difficult, so there might have additionally been some cognitive overload on drawing maps for Text 1 which biased the results in favour of Text 2. Generally all the participants complained that the whole experiment took longer than they had expected. This indicates that the texts were challenging enough and not trivial. (Murray, 1979) claims that there are variations among people's ability to draw mentally visualised geographic information, indicating that evaluating elicited maps might have a large amount of individual variation. Ideally we need to find alternative ways (other than eliciting map drawing) of assessing the effectiveness of introduction messages.

Other factors could also have affected these results. For example, we did not randomise the order of stimuli. Furthermore, our task-based questions proved too easy for participants because they can be answered by simply searching for the named region with the highest value etc. Another issue which we will want to explore in future work is whether using spatial descriptions (e.g., "a medium-sized city" instead of "Slangin" from Text 1, etc.) in both texts would have helped people to make more sense of the texts, particularly Text 2.

We also ran the text stimuli of the experiment with a VI computing undergraduate who clearly indicated that the texts without the introduction message were much harder to visualise, as he did not know where the cities (referred to by proper name alone) were located. He also indicated that some information required to partially visualise the region could be inferred indirectly from information presented in these texts lacking introduction messages. For example, in Ex. 2, he could deduce that poverty rates are high in Central Dagor, since it decreases as one goes East or West from Central Dagor. However there was not enough information for him to place Boroos in Central Dagor.

5 Conclusions and Future Directions

In this paper we describe how content determination in Atlas.txt is achieved and report on evaluation of one message in our content determination schema: the introduction message. Our end goal is to enable VI users to mentally visualise the region being described in order to locate where areas with high

and low values of the census variable described are and also to visualise spatial trends for this variable in the region. In order to evaluate whether introduction messages in our texts help users to mentally visualise the region, we elicited drawings of maps visualised by sighted participants based on our texts with and without the introduction message. This was based on the assumption that VI people also have spatial cognition abilities, and because sighted participants can draw the maps they visualise. We did not see a significant difference between texts with and without introduction messages from the elicited drawings of maps which were scored on the correct locations of maximum, minimum and trend information. However we saw that participants overwhelmingly preferred texts with introduction messages to those without, and also tended to include introduction messages in texts describing maps when asked to describe the maps to blind readers.

The question then arises as to how one should ideally evaluate message inclusion when communicating spatial data. Although drawing maps seemed a natural way to evaluate visualisation, we saw that it was fraught with difficulties. We leave the question of finding an appropriate evaluation methodology for message choice in spatial data-to-text systems such as ours for future work.

Evaluating the utility of the individual components of the introduction message has also been left as future work. Furthermore, we also plan on investigating the factors involved in selecting description geographies, e.g., causality and context, in order to produce informative spatial descriptions in our messages.

References

- R. Barzilay and M. Lapata. 2005. Collective Content Selection for Concept-to-text Generation. *Human Language Technology Conference*.
- S. Boyd. 1997. Detecting and Describing Patterns in Time-varying Data Using Wavelets. *Advances in Intelligent Data Analysis: Reasoning About Data, Lecture Notes in Computer Science*, 1280.
- J. Coch. 1998. Multimeteo: Multilingual Production of Weather Forecasts. *ELRA Newsletter*, 3(2).
- R. Dale and C. Mellish. 1998. Towards the Evaluation of Natural Language Generation. *Proceedings of the First International Conference on Evaluation of Natural Language Processing Systems*.
- C. Ebert, D. Glatz, M. Jansche, R. Meyer-Klabunde and R. Porzel. 1996. From Conceptualization to Formulation in Generating Spatial Descriptions. *Proceedings of the 5th European Conference on Cognitive Modelling*.
- L. Ferres, A. Parush, S. Roberts and G. Lindgaard. 2006. Helping People with Visual Impairments Gain Access to Graphical Information Through Natural Language: The iGraph System. *Proceedings of ICCHP 2006, Lecture Notes in Computer Science*, 4061.
- S. Geldof. 2003. Corpus Analysis for NLG. *Proceedings of the 9th EWNLG*.
- E. Goldberg, N. Driedger and R.L. Kittredge. 1994. Using Natural-Language Processing to Produce Weather Forecasts. *IEEE Expert*, 9(2).
- V. Gyselinck, R. De Beni, F. Pazzaglia, C. Meneghetti and A. Mondoloni. 2006. Spatial Knowledge Acquisition from Maps and from Navigation in Real and Virtual Environments. *Memory and Cognition*, 27.
- R. Kitchin, M. Blades and R. Golledge. 1997. Understanding Spatial Concepts at the Geographic Scale Without the Use of Vision. *Human Geography*, 21(2).
- T. Marciniak and M. Strube. 2004. Classification-based Generation Using TAG. *Proceedings of Natural Language Generation: 3rd International Conference, Lecture Notes in Artificial Intelligence*, 3123.
- W. Metzger. 2006. *Laws of Seeing*. MIT Press.
- R. Mitkov. 1990. A Text-Generation System for Explaining Concepts in Geometry. *Proceedings of the 13th Conference on Computational Linguistics*.
- D. Murray and C. Spencer. 1979. Individual Differences in the Drawing of Cognitive Maps: the Effects of Geographical Mobility, Strength of Mental Imagery and Basic Graphic Ability. *Transactions of the Institute of British Geographers, New Series*, Vol.4, No.3.
- H.J. Novak. 1986. Generating a Coherent Text Describing a Traffic Scene. *Proceedings of the 11th Conference on Computational Linguistics*.
- E. Reiter and R. Dale. 2000. *Building Natural Language Generation Systems*. Cambridge University Press.
- B. Schneiderman. 1996. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualisations. *Proceedings of Visual Language*.
- S. Sripada, E. Reiter, J. Hunter and J. Yu. 2001. A Two-Stage Model for Content Determination. *Proceedings of the 8th ACL-EWNLG*.
- S. Sripada, E. Reiter, J. Hunter and J. Yu. 2003. Exploiting a Parallel TEXT-DATA Corpus. *Proceedings of Corpus Linguistics*.
- H. Taylor and B. Tversky. 1996. Perspective in Spatial Descriptions. *Journal of Memory and Language*, vol. 35.
- K. Thomas and S. Sripada. 2008. Atlas.txt: First Steps Towards Bridging the Gap Between Geo-Referenced Data and Text. *Proceedings of the Cross-Disciplinary Conference on Web Accessibility*.
- J. Yu, E. Reiter, J. Hunter and C. Mellish. 2005. Choosing the Content of Textual Summaries of Large Time-series Data Sets. *Natural Language Engineering*, vol. 11.