

UBIQUITOUS GIS, DATA ACQUISITION AND SPEECH RECOGNITION

Andrew Hunter^{1,*} and Dr. C. Vincent Tao²

¹ Department of Geomatics Engineering, University of Calgary,

²Department of Earth and Atmospheric Science, York University
Commission IV, WG IV/7 - Data Integration and Digital Mapping

KEY WORDS: Mobile GIS, Wireless Computing, Wearable Computers, Speech Recognition, Data Acquisition

ABSTRACT:

The research asked whether Mobile GIS that incorporated speech recognition was a viable tool for locating defects in pavement, curbs and footpaths. A Xybernaut MA IV wearable computer and Dragon NaturallySpeaking's speech recognition software were tested, and the Geography Markup Language, GML 2.0, was used to implement an application schema for street condition surveys. Both technical and overall accuracy exceeded 95% over six speech recognition tests in environments that were quiet or constantly loud. However, for three tests while walking along a busy road during which the noise level varied, the accuracy of the speech recognition plummeted to 58%. A "standing" test for capturing the position of defects (N=30) gave an error of 0.41m at the 95% confidence interval. Finally, a web-based questionnaire was answered by 80 GIS project managers, who indicated that they are unhappy with the quality of their data, although they do not require the data in real-time.

1. INTRODUCTION

This paper investigates the use of a data acquisition tool that has been designed to minimise both positional and attribute errors. A key objective of this work has been to provide the field operator with a tool that allows them to take ownership of the whole data acquisition process, thus improving workflow by removing duplicated effort and the need to communicate difficulties to office based colleagues. The fundamental research question is whether a Mobile GIS, that includes speech recognition and wireless connectivity for real time access to spatial data, is a viable tool for data acquisition?

The mobile GIS tool has been designed around a wearable computer, which utilizes a multi-modal interface. Speech recognition (SR) and Text to Speech (TTS) technologies have been combined with the traditional keyboard, mouse, and visual display unit interface in an effort to create a hands-free computing environment. Wireless communication has been incorporated to provide increased mobility by allowing the GIS to be free from location constraints such as physical network connections (see Figure 1).

Traditionally, due to the high cost of field acquisition, the capture of geographic information (GI) for inclusion in a Geographic Information System (GIS) has been undertaken using techniques such as digitization or scanning of paper maps. However, the scale of the map, the data collection and manipulation techniques utilized, the symbolization employed by the cartographer, and the medium upon which the map is produced all determine the accuracy of any measurement, and therefore the accuracy of any information derived from the measurements. Data accuracy is further compounded by the need to produce clear maps. This often requires that the cartographer rectify, enlarge or move elements on a map in order to clarify a situation. This practice leads to local differences in scale, rotation and translation (de Knecht *et al*, 2001). The result is that the positional quality of data digitized from maps can be significantly more than map accuracy requirements promoted by organizations such as the US Geological Survey.

Personal experience also indicates that GI users are inclined to believe that the accuracy of their spatial data is better than it really is. For example, while working for Truebridge Callender Beach Ltd. capturing water facilities for Wellington City, New Zealand, it was found that the digitized positions of features

with respect to positions obtained from the City Mapping Database had a positional accuracy of 2.38m at a 99% Confidence Interval (CI) for Johnsonville, a suburb of Wellington City (Truebridge Callender Beach Ltd., 1999). The Council required an accuracy of 0.15m (the accuracy of the City Mapping Database) at the 99% CI, and refused to acknowledge that the plans provided for conversion could not achieve their accuracy requirements.

This paper commences with a synopsis of the mobile computing paradigm followed by an examination of SR. Development of a Geography Markup Language (GML) Schema is discussed as a vehicle for the storage and transport of GI between field and office, followed by a review of current GI requirements assembled from a Data Accuracy Requirements Survey undertaken for this research. The paper concludes with initial test results of the mobile GIS under development and a



summary of future research activities.

2. MOBILE COMPUTING AND GIS

Mobile computing is the use of computing devices at some distance from the normal, fixed workplace that enables the mobile worker to: create, access, process, store, and communicate information without being constrained to a single location (Zimmerman, 1999). It can be viewed as a combination of three important and related properties: computation, communication and mobility. Computation includes the computing devices at either end of the network; communication systems include the different wireless and wired networks that link the computing devices; and mobility is an aspect of user behaviour (Liu *et al*, 1995).

At the heart of mobile computing lies a need to deliver intelligence to the field to improve productivity and provide a competitive edge in the marketplace. In order for a mobile GIS to be successful it must emulate existing field practices and eliminate repetitive time-consuming tasks.

The mobile workforce often has little computing experience (Wilson, 2000), as such a mobile application should help a mobile user to automate their entire workflow and improve their efficiency. However, a mobile GIS application should also support primary functions such as mapping and navigation (zoom and pan), data collection, query, update and transmission, remote data and component access (wireless), location determination (GPS), coordinate transformations. But the most important requirement is that the application is able to work in the same environment as that of a field user.

2.1 Wearable Computers

Conventional mobile devices include Personal Digital Assistants (PDA's). While cost and convenience make devices such as a PDA attractive, limited disk space, memory and battery capacity can impose considerable restrictions on mobile applications. An alternative is the wearable computer, whose most distinguishing feature is that it can be used while moving around (Rhodes, 1997). They emphasize hands-free computing, and concentrate on speech input and heads-up displays. Wearable computers are designed to be usable at any time with the minimum amount of cost or distraction from the wearer's primary task, i.e., work or recreation.

3. SPEECH RECOGNITION

Speech recognition has long been recognized as an appealing computer input mode (Licklider, 1960; Murray, Jones, and Frankish, 1996; Tyfa, Howes, 2000). It is generally accepted that speech recognition can play a useful role in hands-busy, eyes-busy environments such as field acquisition of GI.

Speech recognition transforms human voice input into commands and characters using one of two modes; command-and-control (discrete SR) and dictation. Typically, discrete SR is used to navigate an application. For example, a user speaks the command "Launch [Application]" to launch the application. The user is required to speak discrete words or short phrases that imitate an application menu. Dictation is typically used for word processing based activities and requires a considerably larger vocabulary.

To perform SR, an application requires an audio source object (a microphone) to accept speech input, a SR engine that resolves speech input into text or commands, and a list of speech commands, a grammar, supplied to the SR engine.

3.1 Vocabularies and Templates

The words that a discrete SR engine resolves are its vocabulary. A vocabulary consists of one or more templates for each word in the vocabulary. In SR, a template is a pattern that can be used to identify or be matched to a speaker's pronunciation of a word. When a SR engine tries to resolve a word, it compares the audio input stream to its templates until it finds a match, or determines no match is available. If a command is recognized the engine then notifies the associated application.

Templates typically fall into two categories, Speaker-independent and Speaker-dependent. Speaker independent templates combine multiple pronunciations of a word and are intended to require no user training. Speaker dependant templates contain single words or phrases and require the user to train the speech engine so that their particular pronunciation can be recognized.

Many speaker-independent applications are field-driven. That is the application will "anticipate" certain vocabulary only when a particular form or field has been activated. In a street condition survey application, for example, saying the word "defect" could activate the Defect vocabulary, for which there may only be a limited number of acceptable inputs; "distortion", "ravelling" and "cracking", for instance. By limiting the number of acceptable values for this field, accuracy and computing efficiency are maximized. By utilizing unique commands, these inputs are not acceptable in any other field, minimizing the risk that they will be inserted in the wrong place by accident.

3.2 Grammar

A grammar is a subset of the available vocabulary that contains only the words used by the application. A grammar is normally created to avoid comparing a spoken command to all the words in a speech engines vocabulary. A grammar should contain only those words that the application requires. A single application may utilize more than one grammar. As a user moves from one application context to another, the active grammar changes to the commands required for the new context. Managing grammars efficiently makes speech-enabled applications easier to use and speech recognition more accurate.

For speech recognition to work effectively, there are a few guidelines that should be adhered to when developing a grammar. Firstly, choose grammars that are consistent with other applications. Use intuitive words for commands that identify the tasks that the application performs, or the objects that are being captured. Speech menus should be short and should use words that are easy to remember. Avoid using commands that sound alike, and use vocabulary included with speech applications when possible.

3.3 Speech Algorithms: The Basis of Speech Technology

There are many words that sound alike but have different spellings and meanings. This presents challenges for speech recognition software vendors and speech-enabled application developers.

There are two basic approaches to deciphering spoken commands. One approach uses matching algorithms that compare bit patterns of the spoken commands to standard bit patterns stored in the SR software's library of patterns. The commands and related bit patterns are matched for associated actions. The library of bit patterns is created by averaging the patterns of a large multicultural sampling of pronunciations of a finite vocabulary.

The second approach relies on speech patterns that are created by users “training” the speech software. The SR software creates patterns based on the user’s pronunciation of the phonemes*. This is the approach employed for this research. Spoken words are then matched against patterns in the libraries and those created during the user’s training of the software.

3.3.1 How Speech Recognition Works

Speech recognition consists of three broad processes, being the transformation of digital audio from a sound card into a better acoustic representation; determination of the phonemes that have been spoken; and matching the phonemes with a grammar that has been supplied to the SR engine so as to establish the spoken words.

Firstly the Pulse Code Modulation (PCM) digital audio produced by the computer’s sound card is transformed into the frequency domain using a band limited Fast-Fourier Transform (FFT). This allows the frequency components of a sound to be identified. A sound card typically samples an audio stream 16,000 times per second. In order to reduce processing time the FFT samples the PCM audio stream every 1/100th of a second. From this sampling process a graph of the amplitudes of each frequency component is created to describe the sound heard for that 1/100th of a second.

The SR engine contains a database of graphs that identify different types of sounds the human voice can make. The graph produced by the FFT is then “identified” by matching it to the closest entry in the database.

Background noise and speaker variation problems are minimized via the use of statistical modelling to determine which phoneme was spoken. Statistical analysis is possible because a phoneme often lasts for 50 to 100 frequency components, and it is likely that one or more sounds dominate during that time that can be used to predict the phoneme spoken. Determination of a phoneme by a speech engine works by hypothesizing a number of different states at once. Each state contains a phoneme with a history of previous phonemes. The state with the highest score (statistically most likely state) is used as the final recognized phoneme.

Once all phonemes have been identified they are compared against a dictionary of pronunciations to determine the word that was spoken. However, if a word is not pronounced as described in the dictionary no match will be found, or the recognizer may select an incorrect word.

4. A DATA MODEL FOR THE TRANSFER OF GEOGRAPHIC INFORMATION

A data model is used to describe and represent some aspect of the real world in a computer. The Open GIS Consortium’s (OGC) Interoperability program was been used in the development of a data model for testing the mobile GIS under development. This research has used Street Condition surveys as the vehicle for testing the application. The purpose of a Street Condition survey is to locate defects within the street pavement, along curbs and on sidewalks.

4.1 Geography Markup Language

The Geography Markup Language is an eXtensible Markup Language (XML) encoding for the transport and storage of geographic information, including both the spatial and non-

spatial properties of geographic features. GML defines an XML Schema syntax that provide an open, vendor-neutral framework for the definition of GI and objects. It supports the description of GI application schemas for specialized domains and enables the creation and maintenance of linked geographic application schemas and data sets, such as a national road network that passes through a city. The municipal organization may develop a schema that describes the type of road features within its domain and import a schema that describes the types of national road features that pass through the city, rather than duplicating road features that have already been described by the national organisation.

GML also supports the storage and transport of application schemas and data sets (Open GIS Consortium, 2001b). GML allows organisations to either store geographic application schemas and information in GML, as the schema can be directly mapped to a database, or to convert from some other storage format on demand and use GML only for schema and data transport. For this research, data was transported using GML and is stored on the server in a database conforming to the GML schema.

GML supports interoperability through the provision of basic geometry tags, a common data model (features/properties), and a mechanism for creating and sharing application schemas. In general terms a schema defines the characteristics of a class of objects: how the data should be marked up. By publishing a schema an organization can enable or enhance the interoperability of their applications.

GML 2.0 defines three base schemas for encoding spatial information. The Geometry schema (geometry.xsd) includes the detailed geometry components. The Feature schema (feature.xsd) defines the general feature-property model (as feature types) and includes common feature properties such as *fid* (a feature identifier), *name* and *description*. The XLink schema (xlink.xsd) provides the attributes that can link different data sets or schemas. The schema documents alone do not provide a schema suitable for constraining data instances; rather, they provide base types and structures, which may be used by an application schema. An application schema declares the actual feature types and property types of interest for a particular domain using components of GML.

4.2 Street Defect Schema Implementation

Figure 2 is a UML diagram for the Road Defect data model. As shown, allowable defect members must be *Road* or *Footpath* instances. The <<restriction>> stereotype applied to the generalization relationship indicates that subtypes defined in the schema are derived by restriction from their super type. For instance, a *RoadDefect* can only contain class members that are members of the class *RoadDefectMember*, in this case, *Road* and *Footpath*.

* Phonemes are a set of similar speech sounds, which are perceived to be a single distinctive sound in a language.

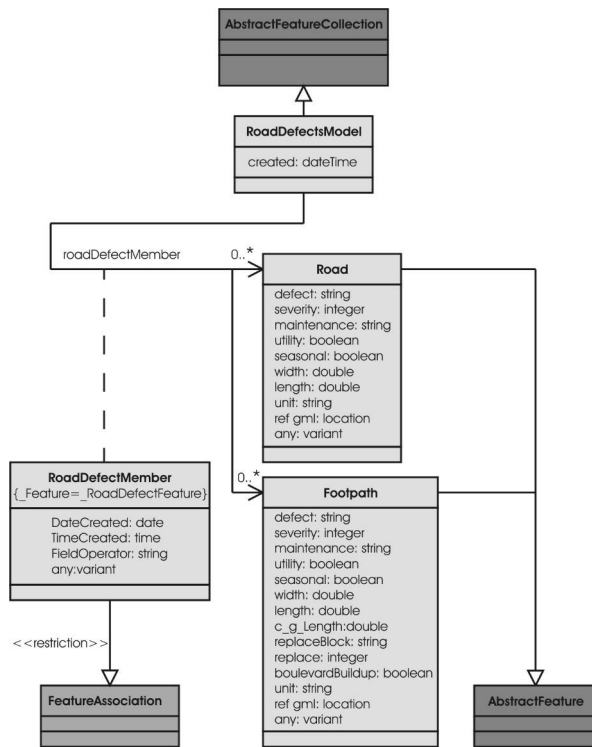


Figure 2: UML Data Model of Defects Schema

Each member class then defines its attributes by extending the base AbstractFeatureType class in the GML feature schema. Road Defects Instance

An instance of the data model for this research (see Figure 3) has a single feature collection of type RoadDefectsModel and contains one feature using a containment relationship called DefectMember. The feature collection has a string property called <Created> which holds the date and time that the GML file was created and a geometric property called <boundedBy> with a <box> value. The box geometry (the 'bounding box' of the feature collection) is expressed in terms of the spatial reference system (SRS) identified by the value of the srsName attribute. The srsName attribute of the geometry types are those described by the European Petroleum Survey Group (EPSG) as proposed by the OGC (Open GIS Consortium, 2001c).

The feature member is an instance of RoadType representing a "Distortion" <Defect> of <Severity> "2", etc. The member has a <defectID> of "1395" and contains other properties describing the date and time the feature was captured and by whom. It has a geometric property called <location> with a <point> value, which are derived from the base GML schemas.

5. DATA ACCURACY REQUIREMENTS SURVEY

A Data Accuracy Requirements Survey was created and sent out to a number of GIS and governmental bulletin boards such as those managed by GITA, AURISA, URISA, the NZ ESRI User Group and the GISList maintained by the GeoCommunity at the beginning of June, 2001. The purpose of Survey was to determine: the spatial accuracy needs of end users; present data capture methods; "time to use" requirements and the level of validation of captured data. The survey was intended for Utility, Local, Provincial and Federal/National Government GIS project managers.

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- File: Defect.xml -->
<RoadDefectsModel
xmlns="http://www.ucalgary.ca/~ahunter/gml"
xmlns:gml="http://www.opengis.net/gml"
xmlns:xlink="http://www.w3.org/1999/xlink"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation=
"http://www.ucalgary.ca/~ahunter/gml/defects.xsd">
  <gml:boundedBy>
    <gml:Box
srsName="http://www.opengis.net/gml/srs/epsg.xml#26911"
>
      <gml:coord>
        <gml:X>750921.789</gml:X>
        <gml:Y>5322419.521</gml:Y>
      </gml:coord>
      <gml:coord>
        <gml:X>750921.789</gml:X>
        <gml:Y>5322419.521</gml:Y>
      </gml:coord>
    </gml:Box>
  </gml:boundedBy>
  <DefectMember>
    <Road>
      <Defect>Distortion</Defect>
      <Severity>2</Severity>
      <Maintenance>Hot box</Maintenance>
      <Utility>True</Utility>
      <Seasonal>False</Seasonal>
      <Width Units="meters">4.3</Width>
      <Length Units="meters">5.8</Length>
      <gml:location>
        <gml:Point
srsName="http://www.opengis.net/gml/srs/epsg.xml#26911"
>
          <gml:coord>
            <gml:X>750921.789</gml:X>
            <gml:Y>5322419.521</gml:Y>
          </gml:coord>
        </gml:Point>
      </gml:location>
    </Road>
    <DefectID>1395</DefectID>
    <DateCreated>2001-12-13</DateCreated>
    <TimeCreated>08:42:53</TimeCreated>
    <FieldOperator>Andrew Hunter</FieldOperator>
  </DefectMember>
  <Created>2000-12-13T08:45:22</Created>
</RoadDefectsModel>
```

Figure 3: Defect Schema Instance

By mid-October 2001, 80 responses had been received. The respondents were grouped into functional groups: AM/FM, Electrical and Gas (12 respondents), Business GIS (5 respondents), Environment (10 respondents), Local Government (36 respondents), Water Resources (4 respondents), and others (13 respondents).

	Accuracy > 1m	Accuracy < 1m	Significance
Actual Accuracy	34	46	
Practical Accuracy [†]	45	35	$p = 0.11$
Desired Accuracy	63	17	$p < 0.01$

Table 1: Comparison of Accuracy Requirements in Relationship to Cost of Acquisition

There is no significant difference between the accuracy of data currently being used and end-users' accuracy requirements. They make do with what they have. However, comparing the respondents who require data at an accuracy of less than one metre (with cost as a constraint), to the respondents who would

[†] Practical accuracy was defined as the accuracy specification considered ideal for an application when cost of data acquisition is a consideration.

like data with an accuracy of less than one metre if cost was not an issue it is evident that respondents desire more accurate data (refer to Table 1). In other words, the cost of data acquisition significantly affects the accuracy of the data.

Improved data accuracy was most desirable for Local Government organisations, however all domains would like more accurate data if they could afford it. The significance values indicated in the table have been obtained by performing a Chi-squared test between actual and practical accuracy, and actual and desired accuracy.

Respondents were provided a number of options ranging from 1 day or less through to 6 months or more, and were asked to indicate actual time-to-use and preferred time-to-use for their organization. Seventeen respondents are currently using data within one day; the majority (61) are using their data within one month. In terms of preferred time-to-use, 32 respondents would like their data within one day and the majority (68) would like to use their data within 2 weeks. However, while the survey indicates that virtually all users would like data to be obtained rapidly, the most desirable timeframe is somewhere between one week and two weeks (see Table 2).

	Actual Time-to-use	Desired Time-to-use	Significance (p)
$t \leq 1$ day	17	32	0.017
1 day $< t \leq 3$ days	10	9	0.037
3 days $< t \leq 1$ wk	12	23	<0.001
1 wk $< t \leq 2$ wks	5	4	<0.001
2 wks $< t \leq 1$ mth	17	8	0.016
1 mth $< t \leq 2$ mths	9	1	0.042
2 mths $< t \leq 6$ mths	5	1	0.246
$t \geq 6$ months	5	2	

Table 2: Comparison of Time to Use (t) Requirements

If the respondents are grouped into Local Government and “Others” we see that “Others” have a greater preference ($p = 0.005$) for 1 week or less and Local Government has a greater preference ($p = 0.002$) for two weeks or less.

Half of the respondents use conflation, of which 53% consider the results acceptable. Only 60% of the respondents do any form of validation, 53% digitize maps in order to acquire data, and 54% are not happy with the quality of their data. To conclude, GI users would like more accurate data, although they do not require the data in real-time. The majority of GI users (79%) continue to use scanning and digitizing for the capture of data, both of which are methodologies that add to the existing spatial error in source data.

6. APPLICATION TESTING

6.1 Speech Recognition Testing

Speech recognition was undertaken by recording spoken commands with a Dictaphone and then comparing them with a log of commands recognised by the speech engine. Agreement is binary, in that the computer recognised command either matches the recorded command or not. In order to obtain a reasonable picture of speech recognition performance, testing was carried out in three different environments with respect to background noise. The first test was performed in a quiet environment (office) where background noise was minimised. The second test was in an environment where the background noise was relatively loud but constant (tested while

driving the car along a motorway). The background noise in the final environment was variable in that there were moments of very low background noise and very high background noise, while walking along a busy street around the University. Quiet periods were observed when there was no traffic; noisy periods occurred when traffic lights turned green or when public transit buses passed. In the quiet environment, the system was also tested on two different computers to determine if computing power affected SR performance.

Table 3 summarises the results in each of the three environments described above. The results have been characterized into Technical and Overall accuracy. Technical accuracy looks only at the recognition errors made by the speech engine, which where either the misrecognition of a command, or the inability to recognise a command. Overall accuracy includes user errors, which consist of poor enunciation of commands or giving an unacceptable command.

Environment	Numbers of Commands	Technical Accuracy	Overall Accuracy
Quite	507	99.3%	98.2%
Constantly Loud	463	96.8%	95.5%
Variable	234	58.8%	57.3%

Table 3: Speech Recognition Results

The results for the first two categories are acceptable (correct better than 95% of the time); however the third category is not, given that this is the environment in which most data acquisition will be performed. Having said that, when analysing the results more broadly it was often found that the speech engine did recognise a command correctly when in this environment, but it took an intolerable amount of time to do so. With Internet based applications a time of 5 to 8 seconds is typically used as the maximum criteria that a user must wait for a response. This is the same criteria used for testing in this research. While waiting for the recognition of a command, additional commands were often issued, hence the very poor result. What is actually happening is that the speech engine has to process all the sounds that it hears; if traffic is busy the computer hears this and tries to make sense of it. Speech recognition is extremely processor-intensive, so in times of high background noise it was found that it could take several minutes before the speech engine actually caught up and recognised a valid command.

The reason that speech recognition worked in a constantly noisy environment is that the speech engine can sample the background noise at the beginning of an exercise and then attempt to remove this from everything that it hears. As is evidenced by these results the major source of error is the microphone. The wearable computer comes with its own specialized directional microphone, but it is clearly not adequate for working in variable noise environments.

There is also an apparent reduction in ability to remember pieces of information, numbers for example, that are otherwise possible in a normal environment minus the computer. This does point to the question of user safety. Does talking to a computer interfere with your ability to recognise what is happening around you? Studies such as those performed by Strayer *et al*, 2001, have provided evidence which shows that those engaged in phone conversations miss twice as many simulated traffic signals as when they were not talking on the phone, and that they took longer to react to the signals that they

did detect. Is talking on the phone any different than talking to a computer in terms of user safety?

6.2 Positional Accuracy

Various modes of transportation will be used while testing positional accuracy in order to match the application to different data acquisition requirements. Four methods are to be analysed: standing on a feature to be captured, and capturing features while walking, riding a bike or driving a car over them. To date, only the initial phase has been completed. The “standing” test involves walking up to a local survey control monument, standing on it, observing a differential GPS position, and then moving away. This process was repeated sufficiently (30 times) to ensure a statistically reliable result. The mean error for the standing test is 0.25m with a standard deviation of 0.08m. This equates to a positional error of ± 0.41 m at a Confidence Interval of 95%.

7. CONCLUDING REMARKS

One of the primary tasks of this research has been to look at an alternative method of data acquisition, the objectives being to improve spatial accuracy, improve attribute accuracy, minimised acquisition timeframes, and removal of intermediate processes typically required to get GI from the field into an application.

Speech recognition plays a major role in this application in terms of system architecture and user interaction. Results indicate that the current system does not respond as instructed, nor in a timeframe that is acceptable when working in a typical environment. However, if the background noise issue can be resolved it is apparent that this mode of interaction with a computer shows immense promise. When in a quiet environment, talking to the computer is fast, and accuracy is acceptable. Speech recognition’s ability to interact with the user via text to speech capabilities (play back to the user what it has heard) provides an opportunity to verify captured data on the fly, rather than having somebody else verify it. If head mounted displays are to be utilised, speech recognition is a necessity. Without it, interaction with the computer via mouse and miniature keyboard is difficult.

Positional accuracy tests completed to date would indicate that the system developed can meet or better the accuracy requirements of most GI users’ who desire accurate (sub-metre) spatial data. Further testing will provide an indication of system accuracy while on the move.

Given the difficulties encountered with speech recognition and in light of the survey undertaken as part of this research, we ask whether it really is necessary to deliver data in real-time to a local government? The study would indicate that, as yet, there is not a great need for instant geographic information.

Further research is required to investigate alternative methods of capturing spoken commands issued by the user. Alternatives include throat and skull microphones, which record sound directly from vocal cord vibrations, rather than through the air. Significant cognitive based testing is also required to ensure that the system is simple and intuitive to use. Lastly, with regards to the wireless/Internet component further investigation with respect to reliability, data transfer capabilities and latency are warranted.

8. BIBLIOGRAPHY

de Knecht, J. P., Schavemaker, J. G. M., Reinders, M. J. T., Vossepoel, A. M., (2001) Utility Map Reconstruction,

International Journal of Geographical Information Science, Vol. 15(1), pp. 7-26

Licklider, J.C.R., (1960) Man-Computer Symbiosis, *IRE Transactions of Human Factors in Electronics*, Vol. 1, pp. 4-11

Liu, G., Marlevi, A., Maguire, G. Q. Jr., (1995) A Mobile Virtual Distributed System Architecture for Supporting Wireless Mobile Computing and Communications *Proceedings of the First Annual International Conference on Mobile Computing and Networking (MOBICOM '95)*, Berkeley, California, USA, pp. 111-118

MobileInfo.com, Mobile Computing Hardware & Software Components, *MobileInfo Hardware & Software Components*, http://www.mobileinfo.com/hswsw_components.htm, (accessed May 15, 2001)

Murray, A. C., Jones, D. M., Frankish, C. R., (1996) Dialog Design in Speech-mediated Data-entry: The Role of Syntactic Constraints and Feedback, *International Journal of Human-Computer Studies*, Vol. 45, pp. 263-286

United States National Map Accuracy Standards - Version 6/1947 <http://rockyweb.cr.usgs.gov/nmpstds/nmas647.html>, (accessed 6/2/02)

Open GIS Consortium (2001a) OGC Document Number: 01-029, Geography Markup Language (GML) 2.0, *OpenGIS® Implementation Specification*, 20 February 2001, <http://www.opengis.net/gml/01-029/GML2.html>, (accessed 8 July, 2001)

Open GIS Consortium (2001b) *The Open GIS™ Abstract Specification Topic 1: Feature Geometry (ISO 19107 Spatial Schema) Version 5* <http://www.opengis.org/techno/abstract/01-101.pdf>, (accessed 8 July, 2001)

Open GIS Consortium (2001c) *The Open GIS™ Abstract Specification Topic 2: Spatial Reference Systems Version 4* <http://www.opengis.org/techno/abstract/99-102r1.pdf>, (accessed 15 July, 2001)

Rhodes, B. J., (1997) The Wearable Remembrance Agent: A System for Augmented Memory, *Personal Technologies*, Issue 1, pp. 218-224

Strayer, D., Drews, F., Albert, R. and Johnston, W. (2001) *Does Cell Phone Conversation Impair Driving Performance?* National Safety Council | Cellphone Article, <http://www.nsc.org/library/shelf/inincell.htm> (accessed 15 July, 2001)

Trubridge Callender Beach Ltd. (1999) *Data Characteristics: Data Conversion of Water Network Record* – Johnsonville Portion Report to Wellington City Council, New Zealand

Tyfa, D. A., Howes, M., (2000) Speech Recognition for Command Entry in Multimodal Interaction, *International Journal of Human-Computer Studies*, Vol. 52, pp. 637-667

Wilson, J. D., (2000). Mobile Technology takes GIS to the Field, *GEOWorld*, Vol. 13(6), pp. 32-36

Zimmerman, J. B., (April 1999) Mobile Computing: Characteristics, Business Benefits, and the Mobile Framework, *Mobile Computing: Characteristics, Business Benefits, and the Mobile Framework*, <http://faculty.ed.umuc.edu/~meinkej/inss690/zimmerman/INSS690CC-MobileComputing.htm> (accessed May 24, 2001)