

IMPROVING ROUTE DIRECTIONS ON MOBILE DEVICES

Sabine Geldof

Centre for Language Technology

Macquarie University

Sydney, Australia

sabine@ics.mq.edu.au

Robert Dale

Centre for Language Technology

Macquarie University

Sydney, Australia

rdale@ics.mq.edu.au

Abstract The provision of information on mobile devices introduces interesting challenges. The most obvious of these is that ways have to be found of optimising the use of the limited available space; however, we also have to take account of the fact that, unlike many desktop-based tasks, activities carried out on mobile devices often require the user to attend to the external environment. In such cases, it is important that the device be able to provide relatively transparent assistance to the user's performance of a task in the real world.

Our focus is on the delivery of route descriptions via mobile devices: our contention is that, in this context, meaningful segmentation of information is a key element in meeting both of the above challenges. This paper describes our approach to developing a mode of interaction which supports the cognitive involvement of the user in performing the task of following a route description; we describe the technological underpinnings of the work and report on a pilot evaluation in a real task setting.

Keywords: mobiles, instructions, navigation, segmentation, usability, cognitive load, language technology

1. Introduction

A great advantage of mobile computing devices is that they allow the provision of what we might think of as ‘run-time support’: assistance in the physical context where a real world task is to be performed. In such contexts, interaction with the device becomes cognitively subordinate to performing the task at hand, giving rise to the notion of the invisible computer (Norman, 1998). For applications which exhibit this property, interface design needs to be focused on supporting the user in performing her task: an interface that complies with the way users cognitively engage with the performance of the task at hand will enhance the process of switching between the physical world and the world as represented through the device.

One such application area is broadly termed ‘wayfinding’ or ‘navigational assistance’: provided with a route description, it is the task of the route taker to apply this description to the real world in order to arrive at some target destination. An increasing number of real-world navigation assistance systems are now available for pre-trip planning via the Web (see, for example, www.mapquest.com and www.whereis.com.au) as well as for in-car GPS-based navigation; the distinction between these two types of system is blurring as the web-based systems increasingly provide functionality for downloading routes to hand-held devices. Our contention is that the type of navigation support these systems provide does not address the user’s cognitive involvement in the task, as is required when delivered via a mobile device. To explore this hypothesis, we compared the output of current automated navigations systems to human-generated route directions.

In general, the turn-by-turn instructions produced in existing navigation systems are generated by a straightforward mapping from a graph-based representation of a route, where edges are travelable paths and nodes are turning points. This approach generally results in sequences of instructions like the following:

Take ramp to I-95 North (just ahead).¹

However, human-generated route directions deviate significantly from this pattern in two ways. Firstly, people select particular aspects of reality based on their saliency, and leave out other details that they assume will be inferred by the route taker. The reasoning processes and underlying knowledge required here are beyond the capabilities of current systems, and although extensive research in the automatic generation of route description (for an overview see (Moulin and Kettani, 1999)) has formulated hypotheses on how these capabilities could be developed, the cost of constructing the required underlying databases puts this out of the question for most practical purposes. Secondly, people tend to provide meta-information and structural information about the route, as in the following example:

Starting from the Macquarie campus, first get on to Lane Cove Road.
The simplest way to do this is to go out of the main campus gate and
keep straight ahead on Waterloo Road ... ²

Here, the route giver provides a high level description of a component of a route, and then provides some elaboration on how the higher level instruction might be carried out.

Focusing on salient properties of the environment and reasoning about the route at a higher level of abstraction seem to be key aspects of the way people are cognitively involved in the route guidance process, with the consequence that their understanding of instructions is a complex interaction between the instructions themselves, their understanding of the global navigation task, and their perception of the environment.

Our aim is to investigate how the task of navigation can be optimally supported given the specific requirements and opportunities encountered in mobile interface design. The first requirement holds for the development of mobile interfaces in general, the obvious problem of limited real estate: screen sizes are small, and so the presentation of information must be optimised. Limitations on the amount of content that can be provided also apply where voice delivery is available, since verbose messages are typically not appropriate. Secondly, our analysis of cognitive involvement in the navigation task needs to be translated into more precise requirements. Höppner formulated requirements for route descriptions in general (Höppner, 1995): a route description needs to be both *recognizable* and *rememberable*. Both are particularly relevant in our context and relate to our above analysis. Current mobile devices do not permit the user to simultaneously attend to both the device and the real world; she needs to be able to relate the description of the world to the physical world (hence recognizability), and to be able to switch back and forth between the device and the real world (rememberability).

Our view is that both the cognitive requirements and the limited real estate can be addressed to some extent by introducing *segmentation* and *structuring* into the flat sequences of instructions provided by existing systems: a hierarchy of instructions that can be explored during route execution then allows the user to follow the instructions at a level that is cognitively appropriate at a given point in time.

The remainder of this paper is structured as follows. In Section 2 we provide an outline of the approach we are exploring, with examples, in the context of the Coral project; in Section 3, we present the underlying technology that Coral uses to achieve this solution; and in Section 4 we describe the results of a pilot evaluation of an interface that uses these ideas. Finally, in Section 5, we draw some conclusions.

2. Structured Information Delivery

Given a flat sequence of instructions of the kind delivered by a typical navigational assistance system, our approach is to segment this sequence of instructions in a meaningful way, and to generate a summary for each resulting segment. This approach has two distinct advantages: first, the result is a structure that lends itself to interactive presentation within a constrained display space; and second, the use of a hierarchy reduces the cognitive load on the user and enhances the rememberability of the route description.

The following example is taken from a version of our system that guides people around our department: the indented italicized material corresponds to the detailed contents of segments, whereas the Roman face lines present the summaries of the segments:

- From Rolf's office go to the lift and turn left.

From Rolf's office go left. Walk down the corridor past one room on your left and two rooms on your right. On the wall straight ahead of you you will see the Language Technology Noticeboard. When you reach it, turn right. Walk down the corridor past four rooms on your left and one room on your right. On the wall on your right you will see the lift. When you reach it, turn left.

- Jim's office is straight ahead of you.

Walk down the corridor past two rooms on your left and one room on your right. Jim's office is straight ahead of you.

The summary lines provide the route taker with a set of higher-level entities that she can choose to explore in further detail as she executes the navigation task. The segmentation structure allows her to easily memorize the overall structure of the route and to recognize progress in the execution of the task.

Using this approach to segmenting a route, we have developed an interface to a real-world route description system that allows the user to browse a route on a mobile device as she performs the navigation task. Figure 1 shows how the above example is rendered on a Palm handheld computer.

When the route is first presented, both the map and the segment summaries are displayed; clicking on a \oplus button leads to the display of the detailed instructions within the corresponding segment. To increase screen space for the text, the user may click the 'Hide Map' button.




Figure 1. Example of a segmented route description presented on a Palm handheld computer; the second screen shows the expanded text for the second segment.

3. Techniques

Our solution is based on two elements. First, the route to be described needs to be segmented and summarized in a meaningful way. In an ideal world this might correspond to the top-down structure developed in a hierarchical planner; however, existing systems do not make use of or provide such structures, and so we have explored the use of bottom-up heuristics for the identification of appropriate segmentations. Currently, our range of segmentation strategies makes use of both significant landmarks as segmentation points, and specific features of the constituent turns and paths (such as status in the road hierarchy and path length).

Then, we need techniques that support flexible interaction with the segmented route in conjunction with task execution. We have developed a markup language called RPML (Route Plan Markup Language) that allows for delivery via combinations of different modalities (textual, graphical, and eventually also voice): our segmentation mechanism produces RPML structures as output, and these are then rendered accordingly on different device types. The rendering device can then provide an interaction mode appropriate for the consumption of the route directions via that device. So, for example, for web delivery we use XLST to deliver the entire route at once via a web page (see Figure 2), whereas the Palm

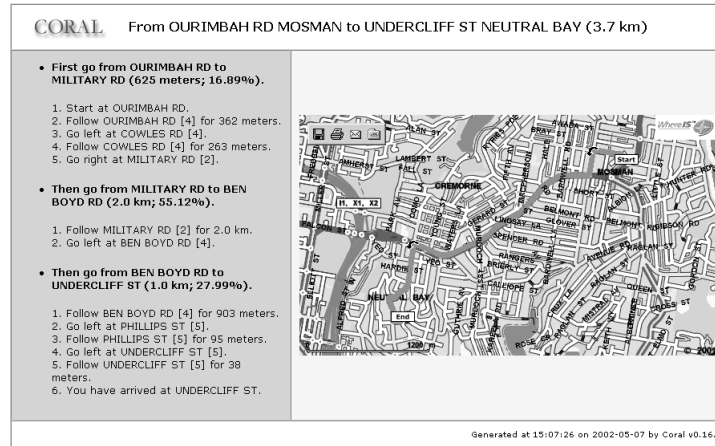


Figure 2. Example of a segmented route presentation via the Web.

renderer allows interactive step-by-step exploration of the description as the user performs the navigation task (see Figure 1).

The following subsections further explain each of these elements.

3.1 Segmentation

As noted above, existing route planning systems provide flat sequences of instructions, consisting of alternating paths and turns, rather than hierarchical structures. The process of segmentation therefore consists in grouping these path and turn instructions into higher level entities that we call *segments*. The notion of segmentation we are working with here bears some relation to the notion of a discourse segment as discussed by (Grosz and Sidner, 1986): elements that are more related are seen as aggregating together to form segments within a larger structure, and in theory this analysis applies recursively to produce a hierarchy.³

The concept of hierarchy in wayfinding is not, of course, new. The process of human spatial knowledge acquisition is often assumed to result in a hierarchical structure, referred to as the cognitive map by (Kuipers, 1978); and Pailhous's observation of wayfinding behaviour by experts (i.e. taxi drivers in Paris) confirmed the hypothesis of the existence of a hierarchical strategy, where first a route between regions is constructed at a higher level before being refined into concrete path components (Pailhous, 1970).

As so far described, segmentation can be viewed as a way of coherently organising and structuring information. However, it can also be seen as addressing a key question in the provision of information in dialogic contexts: how do we convey information in installments so that the course of information exchange approximates the way humans interact? The segmentation of information in human dialogue responds to the need for decreasing the cognitive effort required from the interlocutor (Clark and Schaefer, 1989).

Of course, only a subset of the mathematically possible segmentations of a stream of information is meaningful, and so a key task is to determine which segmentations should be used. We have explored two alternative strategies: one determines optimal break points in the sequence of paths that make up the route, and the other aggregates several paths into a higher level structure on the basis of properties of the constituent elements. These strategies have been applied to the output of existing route description systems.

3.1.1 Landmark-based segmentation. Our first strategy relies on the experimentally verified idea that landmarks at decision points constitute useful cognitive entities that improve the effectiveness of route descriptions (Lovelace et al., 1999; Denis et al., 1999; Burnett, 2000). Although what constitutes a landmark remains vague and ill-defined, attempts have been made to distinguish different categories of landmarks. (Sorrows and Hirtle, 1999), for example, identify visual landmarks (objects such as churches and towers which are clearly distinguishable from their environment by virtue of salient visual features), cognitive landmarks (for example, the desk of a receptionist, which may be significant because it has a particular function for a user), and structural landmarks (entities such as Trafalgar Square in London, which assists in structuring a spatial environment). The saliency of these objects can be exploited to structure route descriptions.

A landmark at a decision point delimits a part of the route to be followed, so the navigator will be aware whether she has reached that point in the route and will thus know how far she has progressed in the navigation task at hand.

We have applied this idea to an earlier version of Coral which provided indoor route descriptions for our department (Williams, 1998). The knowledge representation used in that system includes landmarks as domain objects, and these are included in the intermediate representation from which the textual route description is generated. The route plan representation consists of a sequence of alternating path and turn specifications as shown in the following example, which underlies part of the route presented in Figure 1:

```

⟨start(r333), via(⟨⟩), end(p333)⟩,
turn(lhs),
...
⟨start(p2), via((c1, pass(lhs:[4,room]), pass(rhs:[1,room]), final(rhs:lift1))),
end(lift1)⟩
turn(lhs),
...
⟨start(p362), via(⟨⟩), end(r362)⟩

```

Our segmentation strategy makes use of a separate knowledge source that indicates which domain objects are plausible landmarks; in the present case, the lift is one such object. Since this appears at a decision point (just before the final left-hand turn in the fragment above), it is selected as a segment border and included in the summary for this segment. Consequently, the route is decomposed into one segment leading to the lift and a second segment from this landmark to the destination.

The intuition behind this approach to segmentation is quite straightforward: if the user is familiar with the environment, she will recognize the landmark that terminates the segment and realise that she does not need the detailed instructions for that segment. It is also easy for the user to keep this landmark in mind as an intermediary target and to remember that, once she has reached it, she should revert back to the instructions.

There are, however, limitations to this strategy, since it depends on the presence of landmarks at appropriate locations along the route. Applied blindly, it can lead to segments of significantly varying lengths, which can be confusing. Overall, then, whereas a landmark-based segmentation might be feasible for route descriptions on a small dataset (such as an indoor area), where it is relatively easy to determine which objects of the domain constitute landmarks, it becomes more difficult to apply to a larger scale

3.1.2 Path-based segmentation. Another approach to segmentation is to investigate characteristics of the constituent paths of the route to determine whether they belong to a meaningful higher-level entity. Other work (see, for example, (Höök, 1991)) has explored the hypothesis that recurring higher-level patterns can be found in route descriptions. A frequently occurring pattern consists of three segments corresponding to the beginning, middle and end of a route; typically these involve, respectively, getting onto a main thoroughfare or higher-level road, travelling along that road, and then leaving that road to reach the destination via a number of lower-level roads. We refer to this route pattern as ‘BME’. For example:

- How do I get from Macquarie University to the Queen Victoria Building, in the City?
- [Well, first you get onto Epping Road B], [then you continue ahead via the freeway, following signs to the City M]. [Exit at Druit Street, then the QVB is not far from there E].

Given a flat sequence of paths and turns, we need to determine how these constituents are allocated to segments within such a structure. Our analysis of a small corpus of human-generated routes led us to formulate the hypothesis that three features of paths and turns play a role in this segmentation:

Road status hierarchy: Routes often involve travelling on roads of different status with the road network, from freeways down through main roads to side roads. Our analysis demonstrated that a series of consecutive paths of the same or similar road status is likely to be perceived as constituting a higher-level entity.

Path length: For some routes, segmentation on the basis of road status alone can result in a large number of segments. In such cases, the total length of a segment can help to decide which one of the segments is the stable middle segment.

Turn typology: A turn that is very salient (for example, a T-junction) or that requires careful navigation (for example, a right turn in a drive-on-the-left road context) is a likely segment border candidate.

These principles are very prominent in the prototypical BME route, as demonstrated in the example above: the middle segment consists of a long stretch of one or more steps on higher level roads, and the absence of explicit or difficult turns along this middle segment reinforces the perception of a stable section in the route. However, when examining a larger number of routes, it becomes clear that many variants on this pattern exist, and that these three features interact in a complex manner.

To allow for a systematic exploration of the space, we implemented a segmentation module that takes as input a route (obtained from a route planning system available on the web) augmented with road status information (derived from a widely used street directory). We used 23 routes of different length and in various suburbs of Sydney in our initial exploration.⁴ Our main criterion for segmentation quality was approximation to the prototypical BME pattern. We experimented with various combinations of road status-based and length-based heuristics for segmentation; our conclusions from this study were that road status

is a good indicator for segmentation (in 43% of the cases); in most other cases (another 34%), segmentation can be improved by augmenting this with heuristics that combine segments on the basis of path length. We have not yet investigated the use of turn type to improve segmentation in the remaining cases, but we plan further verification of these preliminary results on a larger set of routes.

3.2 A Route Planning Markup Language

Our goal is to produce one route description that can be rendered via a variety of devices; in the first instance we have been exploring rendering via both standard desktop web browsers and via handheld computers (specifically, the Palm), and we are also extending this to voice delivery via VoiceXML.

To support this variety of outputs, we have defined an intermediate, device-independent representation called RPML (for Route Planning Markup Language). Two principle features of this representation are that (a) it allows for the annotation of a route description with segmentation information that can be used for navigation by the rendering device; and (b) it allows for multi-modal content, such as links to graphical representations of the described route and to voice output. Using this representation, we use XSLT to produce web pages for pre-trip planning like those found at <http://www.ics.mq.edu.au/coral/Routes/Sydney/>, and the same input is used by a specially written renderer on the Palm that formats the output for interactive display to support incremental exploration of the route description while travelling. Figure 3 shows a fragment of RPML; this demonstrates how individual instructions can be provided both as canned output (*First go from BAY RD to PACIFIC HWY*) and as more abstracted specifications (as in the contents of the top level <summary> element) which the renderer can decide how to realise.

4. Evaluation

At the outset of this paper, we made the point that a key feature of applications such as the one we describe here is that the application should allow the attention of the user to be focused on performing the task, not on interacting with the device (Norman, 1998). Consequently, the evaluation of a mobile system requires a different approach to that required for desktop systems where the user only has to attend to the device itself; in particular, it is important to evaluate a mobile system in a real-world task setting.

```
<route-plan context="Sydney">
  <summary>
    <from>BAY RD ARCADIA</from>
    <to>UNIVERSITY AV MACQUARIE PARK</to>
    <distance>35.0 km</distance>
  </summary>
  <map url="http://www.ics.mq.edu.au/~coral/Routes/Sydney/map302.gif"/>
  <segment sid="1">
    <summary>
      <string>First go from BAY RD to PACIFIC HWY.</string>
    </summary>
    <detail>
      <utterance uid="1">
        <string>Start at BAY RD.</string>
      </utterance>
      ...
    </detail>
  </segment>
</route-plan>
```

Figure 3. A fragment of RPML.

In their overview of evaluation methods used for in-car navigation systems, (Ross and Burnett, 2001) distinguish four dimensions along which evaluation methods can be specified: context of use (task-based vs. desk-based), technique (task based vs. checklist), measures (subjective vs. objective), and evaluators (expert vs. user). Their comparison of advantages of different methods as used in cited evaluation studies of navigation systems leads to recommending task-based evaluation on the road by experts, possibly followed by user evaluations to broaden the scope of the evaluation. More generally, the choice of an HCI evaluation method also depends on the stage in the development cycle at which it is to be used: in an early design stage, group discussions, questionnaires and checklists are appropriate; prototype systems might be evaluated by experts; and user trials (possibly on a larger scale) require a fully working system.

In the light of these considerations we decided to perform a small scale expert evaluation in a task-based context.

4.1 The Goals of the Experiment

Our major aims were to obtain feedback on the use of segmented route descriptions and their incremental presentation on a mobile device. Although some objective information was collected, such as the number of navigation errors and the duration of travel, the form of feedback

we were looking for was essentially subjective: significant quantitative evaluation would require a larger scale experiment. In this pilot evaluation, our goal was to get a feel for how the proposed delivery mode and segmented content were perceived in conjunction with navigation task execution.

4.2 Experimental Set-up

There are a great many parameters that might influence the results of a task-based evaluation of a mobile navigation system: user acquaintance with mobile devices, familiarity with the environment to be navigated, general spatial skills, direction of travel, time and context of travel, type of route, quality of signage, to mention but a few. In order to minimize the influence of these factors, we set up a well-defined experimental scenario, where three teams each consisting of a navigator and a driver would use route descriptions on a Palm handheld computer to drive from the University to a particular location and back again. For safety reasons, only the navigator would use the handheld computer, relaying verbal instructions to the driver as appropriate. Different routes, each sourced from an online navigation system and each of approximately 30km in length, were used for the outward and return trips. In order to compare segmented with non-segmented route descriptions, each team followed a segmented route description in one direction and a non-segmented in the other; Figure 4 provides examples of the two types of descriptions.

Two observers also participated in the experiment: each accompanied a team on one of their trips. The observers' task was to observe the navigators without interfering, taking notes on any points of interest. The navigator was asked to log any event or observation related to the navigation task (such as hesitations, navigation errors, lack of information, suggestions for improvement and so on). A post-travel questionnaire included a question asking which presentation mode the navigator preferred and why. The questionnaires, briefing and forms were carefully formulated so as to avoid a suggested preference for one or another presentation mode: for example, we used the terms 'list-based' versus 'tree-based' rather than making reference to the notion of segmentation.

4.3 Experiment Outcomes

4.3.1 Observations. In general, no major problems occurred during task execution; only one navigation error occurred, and no hesitations led to another information source being consulted. All three teams had difficulties with initial orientation, and one also with recog-

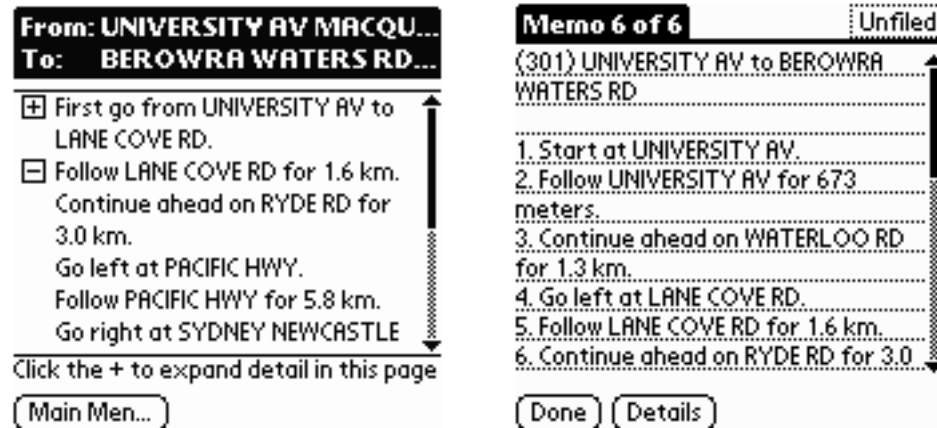


Figure 4. Segmented and non-segmented route descriptions.

nizing the destination, but this was due to a lack of information in the source route descriptions. The observers noted a striking difference in navigation style: whereas one navigator relied heavily on distance information provided in the description (as in *Follow Lane Cove Road for 1.6km*), another navigator did not communicate this information to his driver at all. The latter navigator explored the route description at summary level, but did not systematically provide this information to the driver. Where summaries were read out, they did not provide enough information for the driver to find the correct route (except on the way back to the familiar starting point).

4.3.2 Feedback from navigators. Two of the three navigators preferred the segmented over the non-segmented presentation, indicating that it seemed more appropriate for a limited bandwidth device and that it assisted in keeping track of progress in the directions. The third navigator observed that the use of numbering in the non-segmented presentation mode was helpful and that a typographic distinction between summary and detailed information (for example, the use of a different font) might help to improve the usability of the segmented presentation mode. It should be noted that the preference for one or another presentation mode in this experiment co-occurred with the earlier mentioned difference in navigation style. Our provisional conclusion therefore is

that navigators who rely on factual information (such as exact distance measures) are less inclined to use the structural information provided through segmentation than navigators who use more general information, such as general orientation, topology, and road layout. Whether navigators of the former type could be encouraged to use structural information through a more explicit typographic distinction between the two types of information remains to be investigated. Other feedback and suggestions for improvement relate to the type of information which might be useful, such as the names of cross-streets preceding turns and confirmatory landmarks along the road.

5. Conclusion

Starting from limitations of existing navigation assistance systems and characteristics of human-provided route directions, we have identified key elements that might lead to more effective navigation support systems: emphasizing salient properties of the routes and segmenting routes. The latter aspect is particularly useful in improving navigation support provided via mobile devices as it also allows optimal usage of the limited (screen) space while supporting the navigator's switching back and forth between the presented information and the task at hand. We have investigated and implemented the use of segmentation and summarisation of route descriptions provided via mobile devices. A pilot evaluation of our prototype in a real-world setting is suggestive of the utility of this approach, although a certain type of navigator might prefer a non-segmented description. Larger scale user experiments aimed at collecting quantitative data would be required to refine these conclusions.

Acknowledgments

The authors wish to thank David Hood for his work on the Palm renderer and the participants of the pilot evaluation experiment.

Notes

1. Based on a PocketCopilot example; see <http://www.travroute.com/>.
2. Example occurring in our corpus of human generated route descriptions.
3. In practice, we have so far only found need for one level of hierarchy in our structures.
4. These routes can be inspected online, providing you have an XML-enabled browser such as Internet Explorer 6. See <http://www.ics.mq.edu.au/coral/Routes/Sydney/Segm/rte002.rpml> as an example.

References

- Burnett, G. (2000). 'Turn right at the traffic lights' The requirements for landmarks in vehicle navigation systems. *The Journal of Navigation*, 53(3):499–510.
- Clark, H. and Schaefer, E. (1989). Contributing to discourse. *Cognitive Science*, 13:259–294.
- Denis, M., Pazzaglia, F. C., and Bertolo, L. (1999). Spatial discourse and navigation: an analysis of route directions in the city of venice. *Applied Cognitive Psychology*, 13:145–174.
- Grosz, B. and Sidner, C. (1986). Attention, intentions and the structure of discourse. *Computational Linguistics*, 12(3):175–204.
- Höök, K. (1991). An approach to a route guidance interface. Licentiate Thesis, Dept. of Computer and Systems Sciences, Stockholm University, ISSN 1101-8526.
- Höppner, W. (1995). Modularity in nlg. In *Proceedings of the 4th Intl. Colloquium on Cognitive Science*, pages 117–121, San Sebastian.
- Kuipers, B., J. (1978). Modelling spatial knowledge. *Cognitive Science*, 2:129–153.
- Lovelace, K., Hegarty, M., and Montello, D. (1999). Elements of good route directions in familiar and unfamiliar environments. In Freksa, C. and Mark, D., editors, *Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science. Proceedings of the International Conference COSIT'99*, pages 65–82, Stade, Germany. Berlin: Springer Verlag.
- Moulin, B. and Kettani, D. (1999). Route generation and description using the notions of object's influence area and spatial conceptual map. *Spatial Cognition and Computation*, 1:227–259.
- Norman, D. (1998). *The invisible computer*. MIT Press, Cambridge, MA.
- Pailhous, J. (1970). *La representation de l'Espace Urbain - l'exemple du chauffeur de Taxi*. Presses Universitaires de France.
- Ross, T. and Burnett, G. (2001). Evaluating the human-machine interface to vehicle navigation systems as an example of ubiquitous computing. *Int. Journal of Human-Computer Studies*, 55:661–674.
- Sorrows, M. E. and Hirtle, S. C. (1999). The nature of landmarks for real and electronic spaces. In Freksa, C. and Mark, D. M., editors, *Spatial Information Theory*, number 1661 in Lecture Notes in Computer Science. Springer.
- Williams, S. (1998). Generating pitch accents in a concept-to-speech system using a knowledge base. In *Proceedings of the 5th International Conference on Spoken Language Processing*, Sydney, Australia.