

Generating Personal Travel Guides from Discourse Plans

Shijian Lu, François Paradis, Cécile Paris,
Stephen Wan, Ross Wilkinson, MingFang Wu
CSIRO, Division of Mathematical and Information Science, 723 Swanston St, Carlton 3053, Australia

Abstract

This paper describes a system that delivers travel guides tailored to individual needs. It does so by integrating a discourse planner with a system for querying the web and generating synthesised web pages using document prescriptions. We show by way of example how a user model can lead to a personal travel guide, and show what this might look like in different media. We briefly describe the investigations we are undertaking to determine the utility of such approaches.

1. Tailored Travel Guides

There is an increasing emphasis on the delivery of information that is tailored to individual needs. This paper describes an approach to delivery that assumes a user profile, a discourse model, and a set of web data sources, and delivers information from those sources in a coherent form to a variety of media including paper, Palm Pilots, and the web. The approach requires developing a system that integrates a discourse planner with a web query and document synthesis technology. Both components need to be informed by the user model, and a model of travel guides.

There are currently several commercial travel information web guides that allow some degree of tailoring. These sites operate by having a user provide answers to a set of questions, and the information is delivered by applying a filter to the information resources held at the site, and then deliver a standard structured web page or pages.

2. Discourse Planning and Document Synthesis

In producing a tailored document, information should be coherent and relevant to the user. To achieve this, the resulting document is planned automatically by resolving goals representing mental states of the user, and coherence relations representing how sections of text meaningfully relate to each other. This planner, modelled on a one described by Moore and Paris (1994), produces an intermediate tree structure called a Discourse, which uses nodes and arcs to represent the content of the document and the coherence relations between content parts, respectively. The particular theory of discourse structure used to represent coherency is Mann and Thompson's (1988) Rhetorical Structure Theory, which lends itself easily to computational planning. By using such a planner, only relevant content is selected and assembled for the user.

Complementary to discourse planning is the problem of pulling out bits of information from various sources and combining them to form a valid document. Several models have been proposed to query data sources and transform them for inclusion in a virtual document (Paradis, Vercoestre & Hills 1998). Not surprisingly, the Web community is very active in this domain. The following are elements of a comprehensive solution: script languages, document object models, style sheets and transformation languages, etc.

Perhaps the best examples of tailored document generation at the moment can be seen on the Web. It has been argued that users already perform some level of customisation merely by navigating a site (Nielsen 1998). Good Web sites are designed to facilitate that navigation, and will do it better by taking into account the user community they are designed for. For example, the difference between Internet and Extranet sites is not only about restricting information, but also about providing different information to different people (the public versus the staff). Likewise, most Web portals are designed to suit a particular user community in a particular context. More sophisticated sites try to dynamically adapt their content to the user. An interactive dialogue is established; the system can learn about the user need by presenting forms to be filled out or multiple choice of links, it can then present information according to what it knows about the user (possibly taking into account what it has "said" before), ask for more precision, etc. We claim that to support a meaningful dialogue, a system must keep a current "view" or profile of the user, and have an appropriate model of the discourse.

3. The Personal Travel Guide Approach

The personal travel guide system starts with two things, a user model appropriate to travel, and a general discourse model appropriate to travel guides.

Let us assume that Frank is visiting Melbourne in Australia, from the 15-17th of October 1999. We will assume that Frank wishes to stay in hotels around the \$100-150 range, and that he enjoys opera and cycling, and has an interest in exploring an important local historical figure, Major Mitchell. Frank needs information delivered on the web, for initial exploration, and also in Palm Pilot form, which he will take with him on his trip. Frank has no interest in a paper travel guide. This will be the user profile.

Next we need to develop a discourse model. This is derived by an analysis of a corpus of travel guides, including travel books, travel leaflets, and on-line guides. This leads to a model where after a general introduction, there are generally, depending again on the user model, needs to provide information about accommodation, restaurants, and activities.

The discourse model is now used to develop a discourse. This is done in three stages. First, the general discourse plan is created by evaluating various elements of the user model. The second stage creates the presentation by querying databases of information held as web pages or structured databases. We then have the content of the tailored guide. The third step is to generate a particular form of the content that is dependent upon the medium of delivery. Thus a Palm Pilot presentation has lots of navigation, with little information on any particular page, whereas the Web form will use less navigation, and more layout in generating an interesting set of Web pages.



Figure 1 – A User Model



Figure 2 – Delivery modes

presentation.

4. The Tiddler System

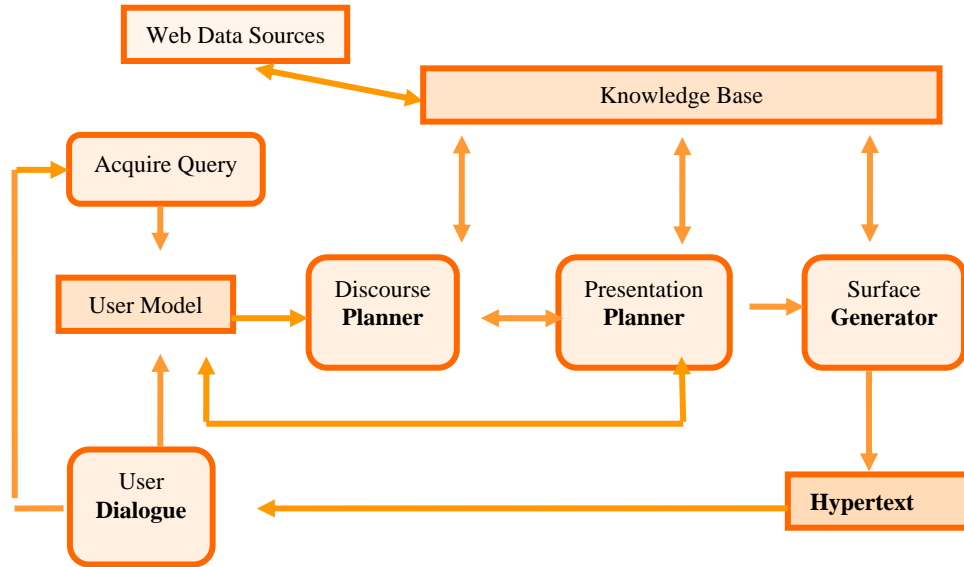


Figure 3 – System Architecture

The architecture of the tailored information delivery system is shown in Figure 3. The core elements of the system are a discourse planner, a presentation planner, and a surface generator. These correspond to the three stages of generation described above. A user dialogue system is also needed to capture the user model and to deliver the information. The discourse planner takes the user model and the prototypical discourse for a travel guide to create a discourse plan, which uses a tree structure. This tree structure is manipulated by a series of RST transformations. The presentation planner then populates this structure with a series of queries against web sources, which are either managed locally, and has a well understood structure, or as general as a list of sources returned from a general web search engine. The Norfolk system is used to create a set of virtual web pages by evaluating document prescriptions lying in the nodes of the tree structure. The answers to these queries are synthesised into a set of HTML answer fragments. The surface generator then takes these fragments and generates a hypertext that is appropriate to the desired medium for delivery. The surface generator again uses RST transformations to take the tree of document fragments and determine a final presentation form.

5. Remarks

We have described a system that delivers travel guides tailored to individual needs. By basing this system around a series of transformations of a basic discourse plan, quite different applications, quite different user models, and quite different delivery platforms can be supported by just changing the data and the transforms. There is no need to change the software. While engineering benefits accrue from this approach, the key issue is the level of benefit of a discourse approach to tailored information delivery. This issue is currently being investigated.

6. Bibliography

- Mann, W.C., & Thompson, S.A. (1988). Rhetorical structure theory: Towards a functional theory of text organization. In *TEXT*, 8(3), (pp 243- 281).
- Moore, J.D. & Paris, C.L. (1993). Planning Text for Advisory Dialogues: Capturing Intentional and Rhetorical Information. In *Computational Linguistics*, Cambridge, MA. Vol 19(4), (pp.651- 694).
- Nielson, J. (1998). Personalization is Over-Rated, in *The Alertbox: Current Issues in Web Usability* (Web bi-weekly), October. URL: <http://www.useit.com/alertbox/981004.html>

Paradis, F., Vercoustre, A.M., & Hills, B. (1998), A Virtual Document Interpreter for Reuse of Information, in *Lecture Notes in Computer Science 1375, Proceedings of Electronic Publishing*, Saint-Malo, France (pp. 487-498).