Case-based User Profiling in a Personal Travel Assistant

Pawel Waszkiewicz[1], Padraig Cunningham[1], and Ciara Byrne[2]

[1] Department of Computer Science, Trinity College, Dublin, Ireland
[2] Broadcom Research Ireland Ltd., Dublin, Ireland

Abstract. This paper presents an architecture for a Personal Travel Assistant (PTA) which can elaborate on a users travel request and evaluate travel offers. Information on user behaviour is stored in a case base in the form of a Case Retrieval Net. This lazy approach has the advantages that it is incremental, extendible and allows flexible reuse of the information.

1 The FIPA PTA Scenario

FIPA is a standards organisation concerned with specifying generic agent technologies. FIPA has proposed a set of prototype applications (FIPA (1997)) in which proposed standards can be tested. One of these applications is a travel scenario where agents assist a user both during pre-trip planning and during travel. The travel scenario is being implemented as part of the FACTS (FIPA Agent Communication and Services) (FACTS (1998)) project which attempts to validate the work of FIPA and other standards bodies by building demonstrator systems.

In the travel scenario envisaged, there will be three different types of agent: Personal Travel Agents (PTA) which represent the interests of an individual user or organisation, Travel Broker Agents (TBA) which broker travel services, and Travel Service Agents (TSA) which represent a particular travel service or have expertise in a particular area, e.g. Danish hotels. A user requests a trip by giving the PTA the trip details and specifying some preferences, such as preferred airline. The PTA then negotiates with the TBA over the details. The TBA composes a trip from several different elements (e.g. flight, hotel, etc.) supplied by the TSAs. This paper focuses on the design of the PTA and in particular on the pre-trip planning phase. We have identified two tasks which the PTA can learn to perform for the user during this phase: adding extra user preferences to a travel request (Task 1) and evaluating a travel offer (Task 2).

2 User Modeling in the PTA

The PTA stores a user profile as a collection of cases. Case Base Reasoning (CBR) was chosen as the learning method because it is an incremental learning system, cases can be acquired easily and the decisions of the agent can be easily explained to the user. CBR is a form of lazy learning. Lazy learning systems often use local approaches (problems are solved by reference to nearby examples in the problem space) which can lead to highly adaptive behaviour not usually found in eager algorithms (Aha (1997)). Case information is stored in a Case Retrieval Net (Lenz et al. (1998)). Two types of entity appear in a Case Retrieval Net(CRN): cases and information entities (IE). An information entity is a basic knowledge item such as a single feature-value pair. Each case is connected to a number of IEs which describe its features. Similarity links exist between

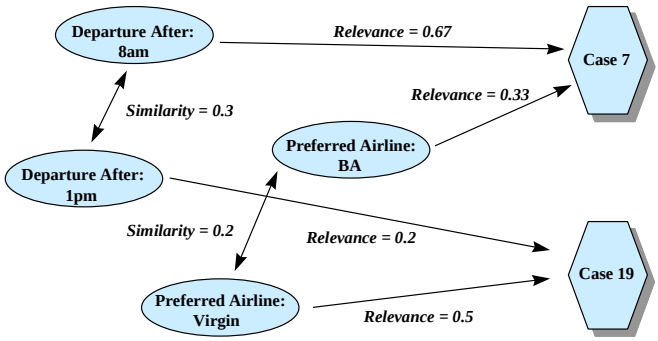

Figure 1. Partial Case Retrieval Net

IEs which have the same feature name and indicates the similarity of a particular pair of values, e.g. how similar a 1-star hotel is to a 3-star hotel. Relevance links between a case and an IE which specifies the importance of that feature in the case (See Figure 1). A spreading activation process is used in order to retrieve cases. IEs which describe the features of a new trip are activated. Activation is spread through similarity links to neighbouring IEs. Finally, cases activations are calculated from the activation of the IEs linked to the case and the relevance of each active IE. The cases with the highest activations are retrieved.

3 PTA Tasks

The PTA acquires profile information by observing the users actions when planning a trip. This information is converted into cases which are stored in a CRN. A minimal amount of direct user feedback will also be used. Figure 2 shows the flow of data through the system when the user requests a trip. The travel preferences supplied by the user are used to query the case base, and one or more cases similar to the new case are retrieved. Travel preferences defined in these cases are extracted. Preferences from several different cases may be combined. The travel preferences specified by the user and those extracted from the case base are merged. The PTA then constructs a trip request message from this information, gets user approval and sends it to the TBA. A new case is created from the travel request.

Task 2 involves evaluating an offer sent by the TBA whose features do not exactly match the preferences specified by the user. Some basic details in the offer and the request sent by the PTA, e.g. destinations, are compared to make sure that they match. If there is a mismatch between these details, then the offer is rejected immediately. Otherwise, the rest of the users preferences are compared with the offer details and the differences between them are identified. This information is passed to the retrieval module which retrieves cases in which there were similar discrepancies. If the offer is categorised as being of interest (depending on whether the offer was accepted in the retrieved cases), then it is passed to the user for evaluation. Otherwise it is rejected.

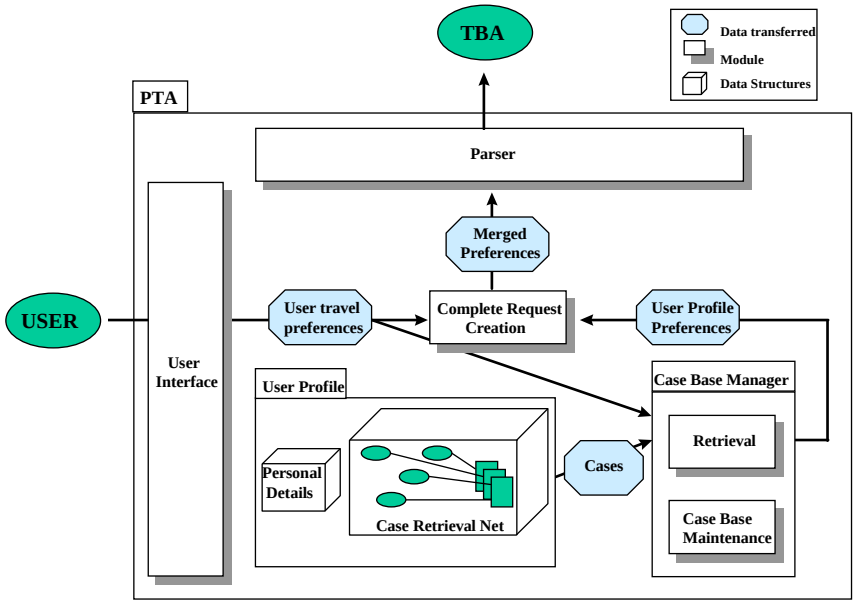

Figure 2. System Diagram

4 Conclusions and Future Work

The PTA is currently being implemented. We feel that the case-based approach to user profiling is suitable in this domain, in particular because of the use of the CRN. The CRN offers several advantages as a method of storing a user profile: it is easy to extend, flexible and content-sensitive. A new case feature can be added to the CRN without making any major changes to its structure. The CRN can be adapted to address more than one problem (i.e. adding preferences and deciding whether to reject an offer) without having to modify the basic case representation or retrieval method. Finally, the CRN allows very flexible specification of similarity functions. The weights on relevance links can be adjusted to implement context-dependent similarity measures where, for example, each case evaluates its input differently. Potential disadvantages include the need to automatically calculate relevance weights and construct similarity metrics. These issues will be investigated as the work proceeds.

References

Aha, D. W. (1997). Editorial on lazy learning. *Artificial Intelligence Review* 11:7–10.

ACTS Project AC317. (1998). *FIPA Agent Communication Technology and Services: Project Summary*.

Foundation for Intelligent Physical Agents. (1997). *Draft Specification, Part 4, Personal Travel Assistance.*

Lenz, M., Auriol, E., and Manago, M. (1998). Diagnosis and decision support. In M. Bartsch-Sporl, H. D. B., and Wess, S., eds., *Case-Based Reasoning Technology: From Foundations to Applications*, volume 1400 of *Lecture Notes in Computer Science*. Springer-Verlag.