# Shared Experiences in Personalized Route Planning *

**Lorraine McGinty** and **Barry Smyth**
Smart Media Institute, University College Dublin
Belfield, Dublin 4, Ireland.

**Abstract**

In this paper we discuss personalized planning (as opposed to personalized information retrieval) where instead of recommending atomic information assets to users, the goal is to construct composite plans that reflect the complex problem solving preferences of users within a particular domain. Specifically, we describe our dual approach to personalized route planning and bring together, for the first time, results to demonstrate the potential of this approach.

## Introduction

The mass availability of the Internet and the range of information services that it supports has been largely responsible for the information overload problems that we are now routinely faced with in our everyday lives. On the one hand, users have never had such ready access to such large amounts of archival and live information. But on the other hand, it is becoming increasingly difficult for users to locate relevant information quickly and easily using modern search-engine or portal technology (Bradley *et al.* 2000; Smyth & Cotter 2000).

In response, there has been a revival of interest in areas such as user profiling, adaptive user interfaces and information filtering in order to deliver the next-generation of personalized information services. The advantage offered is the ability to better respond to the needs of individuals by not only considering their explicit information queries, but also their learned personal preferences, both long-term and short-term (Bradley *et al.* 2000; Cheverst *et al.* 2000; Smyth & Cotter 2000; Strapparava *et al.* 2000).

Today there are many examples of what might be termed *first-generation personalized information systems* that provide similar functionality to traditional search engines except that they filter search results not only on the basis of the relevance of the result to the target user query, but also according to how relevant the result is to a learned user profile (Bradley *et al.* 2000; Konstan *et al.* 1997; Shardanand & Maes 1995; Smyth & Cotter 2000). Such systems are performing an *information retrieval task* and, for the record, are often known as recommender systems.

*This research was funded in part by grant N00014-00-1-0021 from the US Office of Naval Research.


In our work we are interested in a different type of personalization problem, namely *personalized problem solving*. Accordingly the task at hand is a planning one, rather than an information retrieval one, in which a complex plan must be assembled in response to a target problem specification in a way that reflects the preferences of the target user. In the next section we discuss this issue in more detail and introduce our planning task, namely route planning. Section 3 outlines our dual planning algorithm which uses two alternate techniques to solve routing problems depending on the level of coverage the target users case-base offers over the target problem. A more in-depth explanation of these techniques (i.e. pure case-based route planning and collaborative case-based route planning) can be found in previous papers (McGinty & Smyth 2000b; 2000c; 2001a; 2001b).

In the past we have been asked to comment on how these techniques compare in terms of efficiency and solution quality. Thus, the central contribution of this paper is to provide an extended comparative analysis of non-collaborative and collaborative route planning (see Section 4). To do this we re-analyse previous evaluation results in order to further demonstrate of the effectiveness of our route planner. In this analysis we focus on the cost implications of integrating the collaborative component with respect to efficiency and solution quality.

## Personalization: IR vs Planning

The GroupLens project is perhaps the seminal work on personalization and is an excellent example of a personalized information retrieval application (Konstan *et al.* 1997). The target domain is Usenet News articles and GroupLens uses collaborative filtering techniques in order to proactively recommend individual articles to users based on learned user profiles. Each profile represents a set of articles previously rated by the user in question, and the collaborative filtering technique makes its recommendations by identifying users with similar profiles to act as recommendation partners; articles that these users have rated highly and that have not been seen by the target user are recommended.

The CASPER project operates in the online recruitment domain and, unlike GroupLens, uses case-based reasoning techniques to identify job adverts that are relevant to a target user (Bradley *et al.* 2000). It does this by comparing these job cases to the job cases that the user has previously ranked, recommending those that are similar to preferred jobs and dissimilar to jobs the user has not liked.

Recently a number of researchers have started to develop multi-strategy personalization techniques that combine collaborative filtering and case-based methods. For example, the PTV system does precisely this in the domain of TV programme recommendation (Smyth & Cotter 2000). PTV constructs personalized TV guides for users based on their learned viewing preferences by combining recommendation lists that are independently generated using a collaborative filtering and case-based reasoning strategy.

These systems, and many more like them, all share one important feature. They are all focused on personalized information retrieval tasks involving the selection of *atomic* information items, be they news articles, job adverts or TV programmes. In other words, relevant items are selected from an existing data repository and presented unchanged to the user. There is no motive nor opportunity to adapt the items themselves as part of the personalization process. This begs the question of whether there are other personalization tasks where this particular view does not hold?

To answer this question it is appropriate to take a brief detour into the world of route planning. Traditional first-principles approaches to route planning have focused on ways of generating optimal plans by using so-called shortest-path algorithms to minimise some well-defined fitness function governing plan quality. Recent advances have focused on ways of improving the efficiency of such techniques by using case-based methods to generate new route plans by reusing parts of existing plans in order to bypass the need for first-principles search (Branting & Aha 1995; Haigh, *et al.* 1997; Liu 1996; Smyth & Cunningham 1996). Independently, a variety of cognitive science studies have highlighted that a key problem with these route planning strategies is the lack of recognition given to the planning preferences of individual users (Bovy & Stern 1990; Winsum 1989a). These studies have shown that such *one-size-fits-all* strategies are inappropriate for real users because real users generally have very different preference models. Indeed some studies indicate that very often users do not even have conscious access to their own implicit preferences (Rogers & Langley 1998; Rottengatter 1993). The result is the need for techniques that are capable of generating routes that conform to the preference models of individual users.

This has been the goal of our research and we have proposed a case-based reasoning strategy which is innovative on a number of accounts. First of all, unlike other approaches (Rogers & Langley 1998), it does not rely on an explicit model of user preferences, preferring instead to build new route plans by reusing and combining part of previous plans that have been preferred by a user and stored in her profile case-base. The assumption here is that these routes have been preferred by the user because they conform to her implicit preferences, whatever they might be, and that by reusing parts of these preferred routes it will be possible to generate similarly high quality routes in the future. The second innovation is the use of a distributed architecture to allow individual users to benefit from the experience of other similar users in order to generate personalized plans in unfamiliar territories.

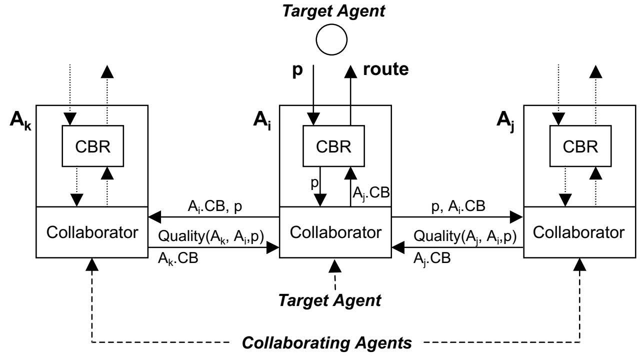


Figure 1: The top-level C-CBR agent architecture.

In the following section we describe this so-called collaborative, case-based route planning approach (C-CBR). However the point here is that route planning is an example of a personalization task that extends beyond the recommendation of individual atomic items. Instead, a new personalized plan is constructed by revising, reusing and recombining parts of previous plans stored in the users profile case-base. Thus the recommendation results are no longer atomic items, but instead are composite structures which must be reassembled as part of the planning process, and chosen and tuned with respect to the target user's preferences as part of the personalization process.

## C-CBR for Personalised Route Planning

Before outlining the C-CBR approach it is important to highlight two important types of route planning scenarios: *Type 1*, solving an unfamiliar problem in a familiar territory; or *Type 2*, solving an unfamiliar problem in an unfamiliar territory. While a pure case-based approach can deal with Type 1 problems (McGinty & Smyth 2000b; 2000c), it fails to adequately handle Type 2 problems, since by definition the user will not have access to relevant cases from the target territory. Instead to deal with Type 2 challenges we implemented a collaborative case-based reasoning technique, where we leverage the experience of other agents with the relevant planning expertise (McGinty & Smyth 2001a; 2001b).

The C-CBR architecture, shown in Figure 1, consists of a collection of homogeneous CBR agents, $A_1, ..., A_n$. For our personalized route planning requirement each agent corresponds to an individual user and the CBR component consists of a personalized case-based route planner with a case-base of their preferred past route planning experiences (see Figure 1). Upon receiving a target route planning problem, $p$ (consisting of a start and goal location), the agent $A_i$ first determines if $p$ falls within its area of expertise. In the context of route planning this capability check determines whether any of the junctions that make up the routes in the agent's case-base are near to the start and goal locations. If they are then the user agent has planning experience in the

```
RoutePlan(start, end, CB, threshold)
  If Dist(start,end) < threshold then
    route    A*(start,end)
  Else if
    case     RetrieveCase(start, end, CB)
    If case then
       section    AdaptCase(start, end, case)
       route    RoutePlan(start, Start(section),CB,threshold)
                + section + RoutePlan(End(section),end,CB,threshold)
    Else route    A*(start,end)
  End if
  Return(route)

RetrieveCase(start, end, CB)
For each case C CB
   C.X'   junction in C with min Dist(start,X')
   C.Y'   junction in C with min Dist(Y',end)
End For
C    case with min Dist(start,C.X')+Dist(C.Y',end)
Return(C)

AdaptCase(start, end, C)
C.X'   junction in C with min Dist(start,X')
C.Y'   junction in C with min Dist(Y',end)
section    road segments in C from C.X' to C.Y'
Return(section)
```

Figure 2: The CBR route planning algorithm.

problem territory and the problem should be covered by the agent's case-base and planning proceeds using the agent's local CBR Planner (see the algorithm given in Figure 2).

Otherwise, if $p$ is not covered by $A_i$ then the collaborative component seeks to locate cases from a set of similar users that do have the required planning experience, with $A_i$ solving $p$ using these *borrowed* cases. Specifically, $A_i$ broadcasts $p$ to each remote agent ($A_j$ and $A_k$ in Figure 1) and selects that agent with the highest quality score.

### The Case-Based Component

The core case-based route planning algorithm is shown in Figure 2. Very briefly, each new route is generated recursively by retrieving and adapting multiple cases to fit the current problem. Each recursive call attempts to solve part of the overall target problem by reusing a case segment such that the remaining uncovered distance is maximally reduced. If at any stage a suitable case cannot be found, or if the distance between the current start and goal locations is below a set threshold, the standard distance-based A* planning algorithm is used to complete the route. In addition, a geometric case indexing scheme is employed in order to deliver a cost-efficient and accurate retrieval system. The details of this so-called *fast-indexing* strategy are provided in (McGinty & Smyth 2000b; 2000c) along with further details on the case-based planning component.

### The Collaborator Component

To deal with Type 2 problems (problems from unfamiliar territories) a target user agent $A_i$ must be provided with access to relevant cases from another user with the necessary experience. The key to success is the identification of agents that not only have experience in the target problem territory, but that also have similar route planning preferences to the target user. Each agent, $A_j$, is equipped with a *collaborator component*, whose job it is to assess whether it is capable of solving the current target problem, $p$, and whether its preference model is similar to that of $A_i$. The agent that reports with the highest quality score is chosen as the collaborator and its case-base is merged with $A_i$ prior to final problem solving by $A_i$.

The quality of an agent $A_j$ depends on the coverage that $A_j$ provides of $p$ and the similarity of $A_i$ and $A_j$. The quality metric used is shown in Equation 1, which allows the relative weight of user similarity and problem coverage to be adjusted.

$$Qual(A_j, A_i, p) = (1-w)*Cov(A_j, p) + w*Sim(A_j, A_i) \quad (1)$$

Once again, the precise details of how problem coverage and agent similarity are computed are discussed elsewhere (McGinty & Smyth 2000b; 2000c), and for reasons of space cannot be reproduced fully here. Very briefly though, problem coverage is determined by measuring the overlap between the set of junctions covered by agent $A_j$'s case-base and those junctions in the region of the target problem $p$. The higher the overlap, the better the likelihood that agent $A_j$ will contain cases that cover $p$.

Measuring agent similarity is more complex. The objective is to measure the similarity of agent $A_i$'s and agent $A_j$'s underlying preference models by looking for similarities in their case-bases. This involves comparing the solutions that these agents have for any problems that they share, under the assumption that if these agents solve the same shared problems in similar ways then they must do so because of similarities in their underlying preference models.

## Evaluation

In previous work we have separately reported evaluation results governing efficiency and solution quality for Type 1 and Type 2 problems[1]. In this section we revisit these results to provide a meaningful comparative analysis. Specifically, what additional computational burden does the collaborative component place on the existing case-based route planner? And, how is the quality of the final routes affected when we *borrow* cases from imperfectly similar users to the target?

### Set-Up

We compare the performance of two route planning algorithms: (1) Distance-Based A* - a standard A* planner with a distance-based cost function; (2) C-CBR - a collaborative, case-based planner with fast-indexing. Specifically this evaluation focuses on how the alternative techniques our C-CBR planner uses to deal with Type 1 and Type 2 problems compare in relation to eachother, and also, relative to the traditional A* approach.

In order to evaluate our technique we simulate user profiles in a similar manner to the *dummy profile* strategy described in (Rogers & Langley 1998). We define a user cost function by assigning random weights to the road segments of our digital map, and the cost of an individual road segment is computed according to the cost function shown in Equation 2.

$$Cost(seg) = length(seg) * weight(seg) \quad (2)$$

[1] Our dual C-CBR planner solves Type 1 problems using only the case based technique and extends this technique by integrating a collaborative component to deal with Type 2 problems.

One can view these weights as being inversely proportional to the *desirability* of the road segment for a given user, where desirability is based on some complex and hidden user preference model. Road segments with a high weighting have a low desirability and present with a higher cost than similar length road segments with a low weight (high desirability). In this way we can generate arbitrarily large user case-bases for given user by using their specific cost function in association with an A* planner to solve a set of selected route planning problems. This will guarantee to generate a case-base of routes that minimize the user's cost function. These are the ideal routes for the user, the routes that the user would prefer in a real-life route planning scenario. We also localise each case-base to a specific map territory by selecting the route problems from this territory according to a specific probability function.

## Method

A set of 260 test users is generated, each with a different preference model (cost function). These preference models are generated such that they display varying degrees of similarity. For each user we generate a case-base of 200 cases in a target territory and 60 test problems inside (Type 1) and outside (Type 2) of this territory.

It is important to note here that user profiles are constructed such that the user's personal cost function is not correlated with the distance-based A* cost function. In this way, we are assuming a minimal relationship between the A* distance-based cost heuristic and the user's preference model. The optimal route for each target problem and user is calculated such that the cost function for that user is minimised. These optimal routes are the routes that would be ideally chosen by the user and serve as a benchmark against which to evaluate the routes produced by our collaborative case-based planner and the distance-based A* method, which of course do not have access to the actual cost functions.

All target problems are solved by the two test algorithms, using different sized profile case-bases for each user, ranging from 50 to 200 cases. The mean problem solving efficiency and solution quality values are measured for the resulting target solutions, to give Type 1 and Type 2 efficiency and quality results for each algorithm. The above is repeated across all 260 users and the efficiency and quality results are averaged for each of the case-base sizes examined.

## Planning Efficiency

For the Type 1 and Type 2 problems we measure the mean planning efficiency of each of the planning algorithms in terms of the mean problem solving time for the target problems across each case-base size. Figure 3 reports these results in terms of the mean speed-up found for the C-CBR approach relative to the A* approach; for instance a speed-up of 2 indicates that the C-CBR approach generates solution plans in half the time needed by the A* approach. Figure 3 shows the speed-up results for Type 1 and Type 2 problems separately. Clearly the results indicate that the C-CBR technique enjoys significant efficiency benefits, compared to A*, on both Type 1 and Type 2 problems. For example, for Type 1 problems, average speed-ups of between 10 and 20 are experienced depending on the case-base size; with the higher speed-ups associated with smaller case-bases.

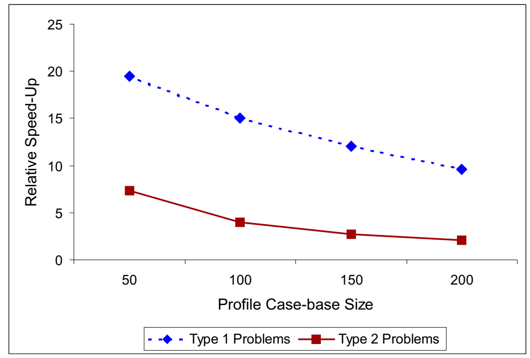


Figure 3: Speed-up for Type 1 and Type 2 problems for C-CBR relative to A* algorithms.

The speed-up found for Type 2 problems (between 2 and 7 according to case-base size) is significantly less than the Type 1 speed-ups due to the extra cost associated with the collaborative planning component and the network transfer times needed for the exchange of information between agents. Nevertheless, the speed-ups still represent a significant efficiency gain when compared to the A* approach. In addition, it must be pointed out that these experiments assume a limited network bandwidth of only 56k bps. As such the results represent a lower-bound on speed-up since we could legitimately expect greater network bandwidth in reality.

## Solution Quality

Of course the key motivation for developing our C-CBR personalized route planner is not an efficiency one, but rather a solution quality one. Ultimately we are interested in generating solution plans that better reflect the route planning preferences of individual users. If this can be achieved more efficiently than with traditional methods then all the better, but we can only claim that our technique is successful if there is also a demonstrable improvement in solution quality.

We measure the quality of a route plan generated by one of our test algorithms by computing its percentage segment overlap between the plan and the optimal plan generated with reference to the underlying preference model. Thus an overlap value of 70% for a plan generated by C-CBR means that the plan shares 70% of its route segments with the optimal plan for the user in question.

Figure 4 reports the quality results for Type 1 and Type 2 problems in terms of the relative quality improvement for C-CBR compared to A*. Thus a relative quality value of 2 means that C-CBR generates plans with twice the percentage overlap of those plans produced by A*. Once again we find that there are significant benefits to be gained from the C-CBR planning approach. For both Type 1 and Type 2 problems we find relative quality values of between 1.5 and 1.9; that is, relative quality improvements of between 50%

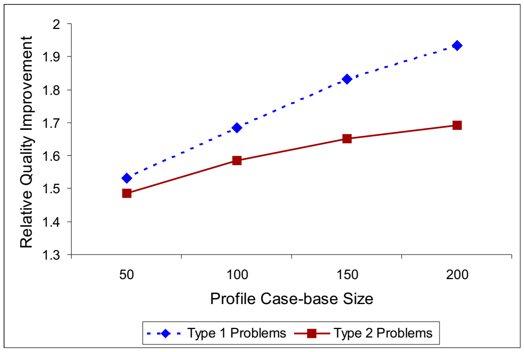


Figure 4: Solution quality improvement for Type 1 and Type 2 problems for C-CBR relative to A* algorithms.

and 90% for C-CBR compared to A*. As expected the quality improvement for Type 1 problems is consistently greater than that found for Type 2 problems. The main reason for this of course is that the Type 2 problems are solved via indirect reference to the user model of another user, which although similar to the target user, is unlikely to have the same preferences as the target. The loss of quality, therefore, is likely to be a direct result of discrepancies between the underlying preference model of the target user and the preference model of the collaborating user. The closer the similarities between these two users, the better the quality results are likely to be.

## Conclusions

The goal of personalization is the automatic adaptation of an information service in response to the learned implicit and explicit needs of an individual user. In recent years a range of techniques from a number of research communities (artificial intelligence, user modeling, information retrieval, and machine learning, to name but a few) have been brought together in the pursuit of effective, real-time personalization solutions. Their core focus has been on the personalization of information retrieval services.

In this paper we have argued for the importance of personalization in problem solving and planning tasks. In particular, we argue that our dual C-CBR route planning approach is capable of generating personalized route plans without the need to draw on explicit user preference models. Evaluation results have been presented to show that the techniques our planner uses to deal with Type 1 & Type 2 route planning problems perform well under a variety of conditions.

## References

Bovy, P., and Stern, E. 1990. *Route Choice: Wayfinding in Transport Networks, Studies in Industrial Organisation*. Kluwer Academic Publisher.

Bradley, K.; Rafter, R.; and Smyth, B. 2000. Case-based User Profiling for Content Personalization. In P. Brusilovsky, O. S., and Strapparava, C., eds., *Proc. of the International Conference on Adaptive Hypermedia and Adaptive Web-based Systems*, 62–72. Springer-Verlag.

Branting, L., and Aha, D. 1995. Stratified Case-Based Reasoning: Reusing hierarchical problem solving episodes. In *Proc. of the Fourteenth International Joint Conference on AI*, 384–390. Morgan Kaufmann.

Cheverst, K.; Davis, N.; Mitchell, K.; and Smith, P. 2000. Providing Tailored(Context-Aware) Information to City Visitors. In P. Brusilovsky, O. S., and Strapparava, C., eds., *Proc. of the International Conference on Adaptive Hypermedia and Adaptive Web-based Systems*, 73–85. Springer-Verlag.

Haigh, K.; Shewchuk, J.; and Veloso, M. 1997. Exploiting Domain Geometry in Analogical Route Planning. *Journal of Experimental and Theoretical Artificial Intelligence* 9:509–541.

Konstan, J.; Miller, B.; Maltz, D.; Herlocker, J.; Gorgan, L.; and Riedl, J. 1997. GroupLens: Applying collaborative filtering to Usenet news. *Communications of the ACM* 40(3):77–87.

Liu, B. 1996. Intelligent Route Finding: Combining Knowledge, Cases and An Efficient Search Algorithm. In *Proc. of the 12th European Conference on AI*, 380–384.

McGinty, L., and Smyth, B. 2000b. Personalised Route Planning: A Case-Based Approach. In Blanzieri, E., and Portinale, L., eds., *Proc. of the Fifth European Workshop of Case-Based Reasoning*, 431–442. Springer-Verlag.

McGinty, L., and Smyth, B. 2000c. Turas: A Personalised Route Planning System. In *Proc. of the Sixth Pacific Rim International Conference on AI, (PRICAI'00)*. Springer-Verlag. Australia.

McGinty, L., and Smyth, B. 2001a. Collaborative Case-Based Reasoning: Applications in Personalised Route Planning. In Aha, D., and Watson, I., eds., *Proc. of the International Conference on Case-Based Reasoning (ICCBR01)*, 362–376. Springer-Verlag.

McGinty, L., and Smyth, B. 2001b. Collaborative CBR for Real-World Route Planning. In Arabnia, H., ed., *Proc. of the 2001 International Conference on AI (IC-AI'2001)*, 254–260. CSREA Press. Las Vegas, Nevada.

Rogers, S., and Langley, P. 1998. Personalized Driving Route Recommendations. In *Proc. of the AAAI Workshop on Recommender Systems*, 96–100. Madison,WI.

Rottengatter, J. 1993. Road User Attitudes and Behaviour. In Grayson, G., ed., *Behavioural Research in Road Safety III*. Transport Research Laboratory, United Kingdom.

Shardanand, U., and Maes, P. 1995. Social Information Filtering: Algorithms for Automating "Word of Mouth". In *Proc. of the Denver ACM CHI 1995*, 210–217.

Smyth, B., and Cotter, P. 2000. A Personalized TV Listings Service for the Digital TV Age. *Journal of Knowledge-Based Systems* 13(2-3):53–59.

Smyth, B., and Cunningham, P. 1996. The Utility Problem Analysed: A Case-Based Reasoning Perspective. In Smith, I., and Faltings, B., eds., *Proc. of the AAAI Spring Symposium on Agents with Adjustable Autonomy*, 392–399. Springer-Verlag.

Strapparava, C.; Magnini, B.; and Stefani, A. 2000. Sense-Based User Modelling for Web Sites. In P. Brusilovsky, O. S., and Strapparava, C., eds., *Proc. of the International Conference on Adaptive Hypermedia and Adaptive Web-based Systems*, 388–391. Springer-Verlag.

Winsum, W. V. 1989a. A MAUT study on car drivers' preferences of routes. In W. Van Winsum, H. Alm, J. S., and Rottengatter, J., eds., *Laboratory and field studies on route representation and drivers' cognitive models of routes*, 11–33. Traffic Research Center, University of Groningen.