Personalizing Routes

Kayur Patel[1], *Mike Y. Chen*[2], *Ian Smith*[2], *& James A. Landay*[1,2]
{kayur, landay}@cs.washington.edu
{mike.y.chen, ian.e.smith}@intel.com

[1]DUB Group
Computer Science and Engineering
University of Washington
Seattle, WA 98105

[2]Intel Research Seattle
1100 NE 45th Street
Suite 600
Seattle, WA 98105

ABSTRACT
Navigation services (e.g., in-car navigation systems and online mapping sites) compute routes between two locations to help users navigate. However, these routes may direct users along an unfamiliar path when a familiar path exists, or, conversely, may include redundant information that the user already knows. These overly complicated directions increase the cognitive load of the user, which may lead to a dangerous driving environment. We have developed a system, called MyRoute, that reduces route complexity by creating user specific routes based on a priori knowledge of familiar routes and landmarks. MyRoute works by *compressing* well known steps into a single contextualized step and *rerouting* users along familiar routes.

ACM Classification H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General Terms Design, Human Factors

Keywords: maps, directions, driving, navigation, personalization

INTRODUCTION
Routing services are becoming pervasive, popular, and indispensable. Every day, millions of people search for driving directions online; four out of five internet users visit online routing services, making searching for directions the second most popular search activity and the third most popular online activity [6]. As the number of in-car navigation systems increases [7] and as local search and mapping services move from personal computers to mobile phones, users will have the ability to ask for directions at any time from any place.

Routes generated by these services are often overly complex, and the increased cognitive load can lead to dangerous driving situations [12]. Directions printed from online sites divert users' attention, causing them to lose focus in an eyes-busy driving environment. Interruptions and distractions caused by in-car navigation systems are even worse than than those of printed directions [7]. These problems could be mitigated if directions had fewer, more relevant steps.

Unlike computers, humans use a priori knowledge to *compress* known steps and *reroute* through known landmarks [11]; they sacrifice speed and completeness for conciseness and familiarity [8, 13]. By mimicking human strategies for giving directions, we attempt to create routes that are more concise and familiar.

We present MyRoute, a routing service that generates personalized routes by exploiting each user's familiar landmarks and routes. By reducing the step count and directing users along familiar routes, MyRoute reduces the number of distractions to provide users with a safer and less stressful driving experience. MyRoute operates below the presentation layer, allowing it to deliver better routes regardless of the presentation method. Since MyRoute is built upon existing routing infrastructure, it can be easily integrated into current routing methods.

MyRoute optimizes for familiarity through *route compression* and *rerouting*. Route compression reduces familiar segments into one contextualized step. For example, a route may contain directions to navigate from the user's home to a local highway. If the user frequently travels to that highway it is better to remove the steps leading from home to the highway and replace them with one step that says "go to local highway." Table 1 illustrates how MyRoute compresses an example route.

Traditional route finding algorithms optimize for distance or time, however there may be a slightly longer route that allows the user to stay on familiar roads. Rerouting considers routes *from* familiar landmarks to the destination to reduce the amount of time users spend on unknown roads. Routes *to* familiar landmarks can also be compressed (e.g., "go to work"). Figure 2 presents a LineDrive [1] example of MyRoute rerouting a user.

Our paper proceeds as follows: First, we present our target user and usage scenario. Then, we describe our compression and rerouting algorithms and our implementation. We go on

UIST'06, October 15–18, 2006, Montreux, Switzerland..

Table 1: The directions below were generated during our informal user study. The first six steps overlapped with a familiar route; MyRoute compressed the familiar six steps into one contextualized step.

Step	Directions	Personalized Directions
1	Start on 10th Ave E (South)	
2	Turn RIGHT (West) onto E Roanoke St	
3	Turn LEFT (South) onto Boylston Ave E	Drive towards Pike Place Market
4	Take Ramp onto I-5	until you are on Union St
5	At exit 165B, turn RIGHT onto Ramp	
6	Road name changes to Union St	
7	Turn LEFT (South) onto 2nd Ave	Turn LEFT (South) onto 2nd Ave
8	Turn RIGHT (West) onto Madison St	Turn RIGHT (West) onto Madison St
9	Turn RIGHT (North) onto Alaskan Way	Turn RIGHT (North) onto Alaskan Way

to present the procedure and results of an informal user study. Finally, we discuss conclusions and our planned future extensions to MyRoute.

TARGET USER AND USAGE SCENARIO

Since rerouting and route compression require a priori knowledge of a user's familiar landmarks, MyRoute provides the most benefit when users navigate within a familiar area. MyRoute would not work in instances where users are completely traveling in unfamiliar areas (e.g., on vacation in an unfamiliar city), however these instances are uncommon. In the most common situation, users drive within their local area and, therefore, will have a number of familiar landmarks.

MyRoute returns the best possible route when either the departure or destination point of the route query is known to the user (e.g., driving to or from home). Route compression requires knowledge of at least one of the endpoints. Rerouting can work if both endpoints are unknown, however reroutes with known endpoints are much more concise than those without. In most situations, users search for routes from known locations to unknown locations and vice versa, and MyRoute provides the most benefit for this common case.

METHOD

MyRoute operates in two steps. First, it collects known landmarks and route data into a personalized profile. Second, MyRoute uses the personalized profile to perform route compression and rerouting. MyRoute picks the best route using a cost function and presents that route to the user.

Implementation

We used ASP.NET to store user profiles and build the front-end, C# to build the back-end, and the Microsoft MapPoint Web Service for our mapping and routing functions. MapPoint uses the NavTech database, which is also used by other major online routing services, such as Google Maps and MapQuest.

User Profiles

Personal profiles provide our algorithm with a basis for reducing the complexity of a route. MyRoute stores these profiles as graphs, where landmarks are vertices and routes are edges. Our prototype requires users to manually input landmarks and specify how they are connected. However, these profiles can be gathered automatically using geolocation systems, such as GPS, found in navigation systems and mobile phones [10], greatly reducing the need for manual entry. In addition, a user's learning rate for new landmarks as well the utility of added landmarks can be personalized through known mixed initiative techniques such as preference elicitation [5].

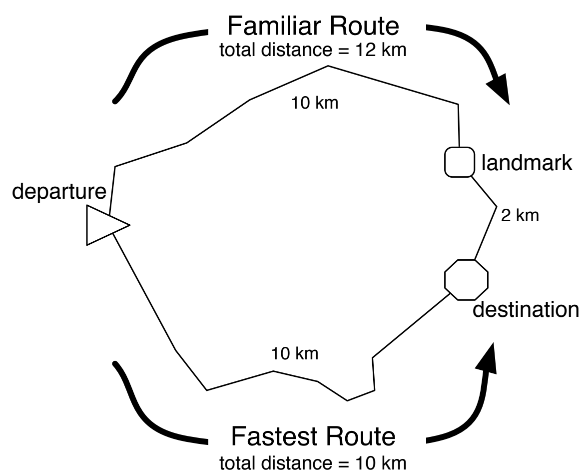

Figure 1: The example above contains a trade off between familiarity and distance. Driving directly from the departure point to the destination is shorter but involves more unfamiliar turns, while driving from the departure point to a landmark and then to the destination is longer but more familiar.

Route Compression

To find overlapping routes, we create a step tree for each landmark, l. Each node in a step tree is a step in a route; each distinct leaf path is a route originating at l. The step tree is generated as follows, given two destination landmark l_1 and l_2 from the departure landmark l_d, we compute two routes r_1 and r_2. Each route consists of steps $r_1 = \{s_{11}, s_{12}, ..., s_{1a}\}$ and $r_2 = \{s_{21}, s_{22}, ..., s_{2b}\}$. If $s_{11} = s_{21}$ then they would be stored in the same node n in the step tree. The node also keeps a list of destination landmarks $n.L$ that have routes containing step $n.s$. Each node has a set of children which describe the next step in some route. For instance if s_{12} and

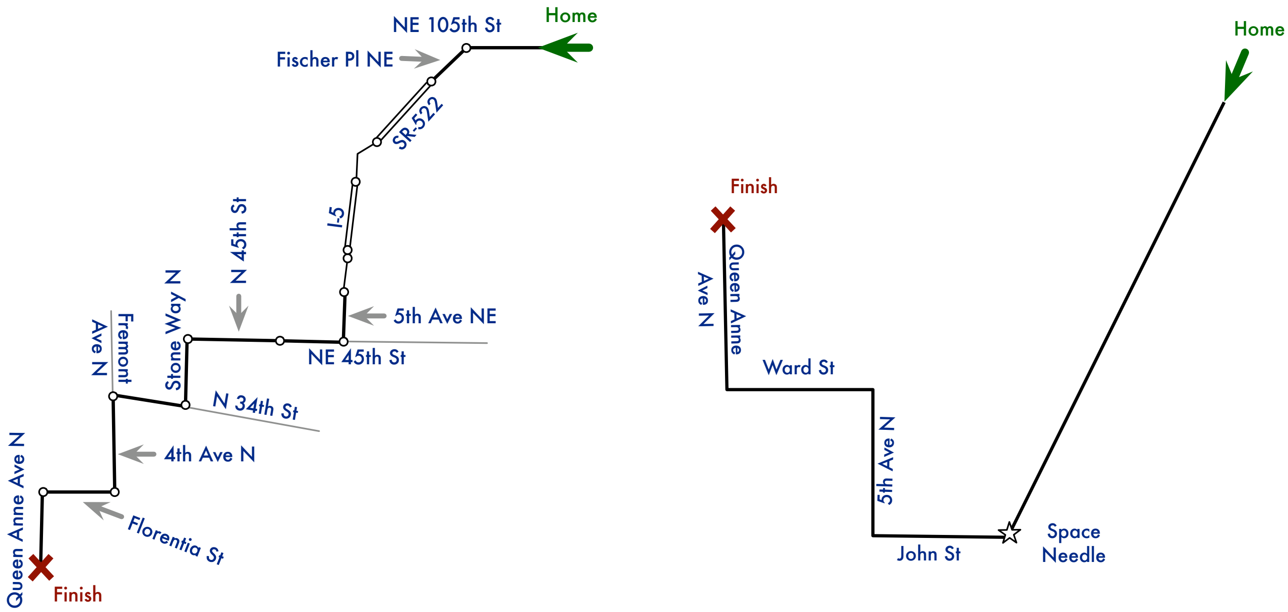

(a) LineDrive (14 steps, 18:34)

(b) Rerouted LineDrive (6 steps, 19:54)

Figure 2: The directions above were generated during the informal user study. The rerouted path that passes by the Space Needle (b) is one minute, 20 seconds longer than the direct route (a). However, the direct route has 14 steps, while the familiar route has six. The user preferred the longer route with fewer steps. LineDrive maps were generated using MapPoint and edited using OmniGraffle.

s_{22} diverge then n would have children $n.C = \{n_1, n_2\}$, where $n_1.s = s_{12}$ and $n_2.s = s_{22}$. A simple traversal of the step tree can find the overlapping region of a new route.

Rerouting

Traveling on a familiar route may be longer in distance or time than traveling the direct route. Drivers prefer both familiar routes and shorter routes, Figure 1 illustrates this tradeoff. We resolve this conflict through a cost function $C(r)$ based on distance (d), driving time (t), and number of steps (s) of route r.

$$C(r) = \alpha_d(d_u + \beta_d d_f) + \alpha_t(t_u + \beta_t t_f) + \alpha_s(s_u + \beta_s s_f)$$

Each factor has familiar (e.g., d_f) and unfamiliar components (e.g., d_u). The α coefficient represents the weight of each factor in relation to the other factor. The tradeoff between taking unfamiliar routes and familiar routes is expressed by the β coefficient. Since we are trying to minimize the cost, β values below one indicate a preference towards familiar routes. Depending on the values of α and β, the route returned may or may not go through a known landmark.

When looking for a familiar route, we find the shortest path from the departure point to the landmarks in the landmark graph. We then compute the path from each landmark to the destination. For our cost function, the path from departure to landmark is considered familiar and the path from landmark to destination is considered unfamiliar. MyRoute computes the cost of all the possible rerouted paths and chooses the route with the lowest cost. The cost of rerouting is compared to the cost of compression and the lowest cost route is chosen.

INITIAL USER EXPERIENCE

To get initial feedback on MyRoute, we conducted an informal study with three participants comparing directions generated by MyRoute to those generated by an online routing service. Each participant owned a car and had lived in the local area for less than a year. We used a fixed cost function that worked well in initial tests. Function parameters were set to the following: $\beta_d = \beta_t = \beta_s = 0.5$ and $\alpha_d = 0.02, \alpha_t = 0.5, \alpha_s = 1$.

Participants generated user profiles and searched for routes from their home to ten new local destinations of their choice. Each search returned two distinctly labeled routes: the original route returned by MapPoint and the personalized route generated by MyRoute. Participants rated the quality of each route on a five point scale and commented on the benefits and failures of our system.

Step count gives an indication of the number of different pieces of information a user must keep track of and the number of interruptions a user faces while driving. For all thirty queries, the step counts for MyRoute routes were at worst the same as MapPoint routes, and on average MyRoute routes had fewer steps than MapPoint routes.

For most routes, participants preferred directions generated by MyRoute, especially those where MyRoute rerouted them though a familiar landmark. When asked why he preferred rerouting, one participant pointed to the tradeoff between familiarity and speed stating, "The psychological value of my knowing where a place is approximately is worth a lot more than five minutes of driving time."

For most route compression examples, participants either liked the compressed route or felt indifferently, indicating that they "already ignore the first few steps." While compressed routes may not have large benefits for directions generated in advance, they may be extremely useful when using in car navigation systems, where distractions and interruptions can be reduced.

In the few cases where participants did not prefer MyRoute, they pointed to the following two reasons: First, small compressions of one or two steps, instead of being useful, were often confusing. Second, participants were often familiar with a landmark but not the surrounding area. Creating heuristics for minimal compression and gathering more detailed landmark information may help mitigate usability problems.

FUTURE WORK

Our informal study uncovered a few design flaws but reinforced the idea that, when traveling, users value traveling on familiar routes. We plan to conduct a full user study to quantitatively assess user preferences and expand MyRoute into an end-to-end system that gathers data via in-car GPS, infers likely landmarks and routes, and provides familiarity-based routing in the contexts of in car navigation systems and online routing services.

In addition to creating and studying a more comprehensive system, we believe that additional features could greatly increase the benefit of familiarity-based routing. For example, learning rates and the tradeoffs between familiarity versus speed and interruptions versus reassurance vary between people, making it impossible to create a solution that works for all users. Future work should focus on creating a better model of the user and applying existing machine learning techniques to fit the parameters to each user.

Our cost function assumes a boolean notion of familiarity. Future work can focus on creating better cost functions by including more varied notions of familiarity or including familiar areas (e.g., Chinatown) rather than specific landmarks. In addition, we can include more variables (e.g., time of day, traffic) and use non-linear cost functions to better capture the users true utility.

Our notion of landmarks is user specific; however, there are universal landmarks that are easy to see, remember, and use while navigating [3, 9]. Since universal landmarks are often visible from long distances, we can exploit real world cues to provide additional context when orienting users (e.g., "turn left at the third traffic light") [4]. Universal landmarks can also help users avoid driving off route by using landmarks to reassure (e.g., "you will pass by a grocery store") and redirect lost users (e.g., "you have gone to far if you see the gas station") [2]. Future work will focus on automatically gathering these landmarks from existing local search APIs and integrating them into MyRoute's personalized directions.

CONCLUSION

We presented MyRoute, a routing service for generating personalized driving directions. MyRoute uses familiar, personal landmarks to reduce route complexity and provide users with more concise and more familiar directions. We found MyRoute reduced the number of steps in a route, and routes generated by MyRoute were preferred by users over non-personalized routes.

REFERENCES

1. Maneesh Agrawala and Chris Stolte. Rendering effective route maps: Improving usability through generalization. *Proceedings of SIGGRAPH 2001*, pages 241–250, August 2001. ISBN 1-58113-292-1.
2. Jorg Baus, Antonio Kruger, and Wolfgang Wahlster. A resource-adaptive mobile navigation system. In *Proceedings of IUI '02*, pages 15–22, New York, NY, USA, 2002. ACM Press.
3. G E Burnett, D Smith, and A J May. Supporting the navigation task: Characteristics of 'good' landmarks. In *Proceedings of the Annual Conference of the Ergonomics Society*, Contemporary Ergonomics 2001, pages 441–446, April, 2001.
4. Andreas Butz, J Baus, Antonio Kruger, and Marco Lohse. A hybrid indoor navigation system. In *Proceedings of IUI '01*, pages 25–32, New York, NY, USA, 2001. ACM Press.
5. Krzysztof Gajos and Daniel S. Weld. Preference elicitation for interface optimization. In *UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 173–182, New York, NY, USA, 2005. ACM Press.
6. iCrossing. How america searches. Technical report, `http://www.icrossing.com/rdm/article_1-1.html`, June 2005.
7. Privilege Insurance and YouGov. Unsafe use of navigation equipment. Technical report, `http://www.privilege.com/aboutus/Unsafeusenavigation.htm`, Febuary 2006.
8. Alexander Klippel, Heike Tappe, and Christopher Habel. Pictorial representations of routes: Chunking route segments during comprehension. In *Spatial Cognition III*, pages 11–33. Springer; Berlin, 2003.
9. Paul U. Lee, Heike Tappe, and Alexander Klippel. Acquisition of landmark knowledge from static and dynamic presentation of route maps. *KI Special Issue Spatial Cognition*, 02(4):32–34, 2002.
10. Lin Liao, Dieter Fox, and Henry A. Kautz. Learning and inferring transportation routines. In Deborah L. McGuinness and George Ferguson, editors, *Proceedings of AAAI*, pages 348–353, San Jose, CA, 2004. AAAI Press / The MIT Press.
11. Clemens Nothegger, Stephan Winter, and Martin Raubal. Selection of salient features for route directions. *Spatial Cognition Computation*, 4(2):113–136, 2004.
12. M. A. Pettit and G. E. Burnett. Defining driver distraction. In *Proceedings of World Congress on Intelligent Transport Systems*, November 2005.
13. Barbara Tversky and Paul U. Lee. Pictorial and verbal tools for conveying routes. In *COSIT '99: Proceedings of the International Conference on Spatial Information Theory*, pages 51–64, London, UK, 1999. Springer-Verlag.