

Enhancing In-Car Navigation Systems with Personal Experience

Benjamin B. Bederson, Aaron Clamage, and Catherine Plaisant

Computers are extremely powerful for data processing but less adept at handling problems that involve subjective reasoning. People, on the other hand, are good at such tasks. A framework is presented for adding subjective human experience to in-car navigation systems. People often rely on personal experience when planning trips, choosing the route that is fastest, prettiest, or recommended by a friend. A set of methods was developed to help people record personal driving history, add textual annotations describing subjective experiences, and share data with friends and family or even the broader community. Users can then learn from personal data or harness the multiplicity of individual experiences to enjoy new routes. This approach can be used in conjunction with traditional in-car navigation systems.

Designers have just begun to realize the potential of using human power to accomplish tasks that computers, automated sensors, and signal-analysis techniques cannot do in a timely and accurate fashion. For example, researchers at Carnegie Mellon University have created online games to attract users to annotate images or add facts to a knowledge base (1, 2). Google Co-op invites people to manually annotate web pages to improve Google's search capability. And search engines work so well in part because they analyze the result of human activity such as creating structured documents and linking among them.

Yet most current navigation solutions rely entirely on objective, automatically collected data and ignore human subjectivity. Personal experience plays a huge role in route choice. Traffic reports and websites do not always provide relevant, timely, and accurate information, and drivers often rely on personal knowledge—or impressions—when planning trips to the airport or train station or visiting friends and family. Most use a trial-and-error approach (3). People have different criteria for the “best” route—fastest, most scenic, most fuel-efficient, or even perceived to be most safe—so combining objective and subjective measures may enhance the route-selection process.

This approach differs from other route-planning systems by starting with a decentralized data-collection process (Figure 1). By recording their own driving histories, people can learn about the routes they actually use. By adding simple textual annotations to those routes, they will later benefit from data that cannot be recorded automatically,

such as the beauty of the scenery or feelings of insecurity. The aggregation of personal data can help drivers discover patterns and confirm (or not) impressions about route characteristics. Although all users get immediate benefit from collecting and analyzing their personal data, some will reap greater benefits from sharing with friends and family. They may choose to share it with larger communities—possibly making the aggregated data fully anonymous—and access more route data in return.

A design framework for collecting, aggregating, and sharing personal driving history is presented as a case study of taking advantage of untapped human experience. A survey is described that was conducted to determine the type of questions people have when selecting routes and how willing people are to share their personal driving experiences. Then, several scenarios—using an early prototype to illustrate the recording and viewing of routes and mockups—are discussed to show how annotations and sharing might take place. Methods of maintaining accuracy and anonymity also are discussed, as well as techniques for aggregating large collections of experiences and presenting that information to users in a meaningful way.

RELATED WORK

Harnessing Human Experience

Existing systems to collect human experiences can be grouped using a simple taxonomy (Figure 2). Some of these systems aggregate data; others do not. Either way, they motivate people to participate in four basic ways. Sometimes, no extra effort is involved (e.g., many e-commerce sites track when users purchase items to make recommendations to others). People may be motivated to spend some effort because it is fun (e.g., they play a game). Or, they may get some direct personal benefit from participating. Finally, people often are motivated if they feel they can help make the world a better place. This work focuses on the last two kinds of motivation (personal benefit and altruism), engaging people because they want to be involved.

Some of the best-known experience-collection systems are the www.amazon.com, www.netflix.com, and www.ebay.com websites, which use collaborative filtering. Amazon.com lets people write reviews and rate items from one to five stars. The ratings are aggregated, whereas reviews are simply listed individually. Readers also can rate the reviewers, and users can sort the individual reviews by those that are found to be the “most useful.” Netflix recognizes that many users favor restricting access to their comments and recommendations to a smaller circle of friends. eBay collect reviews to control malicious use; a seller's average rating may influence whether he or she gets more business in the future. People are willing to invest some effort to serve to the greater good and because they benefit personally when other people give positive feedback about them.

Human-Computer Interaction Lab, Computer Science Department, Institute for Advanced Computer Studies, 3171 A. V. Williams Building, University of Maryland, College Park, MD 20742. Corresponding author: B. B. Bederson, bederson@cs.umd.edu.

Transportation Research Record: Journal of the Transportation Research Board, No. 2064, Transportation Research Board of the National Academies, Washington, D.C., 2008, pp. 33–42.
DOI: 10.3141/2064-06

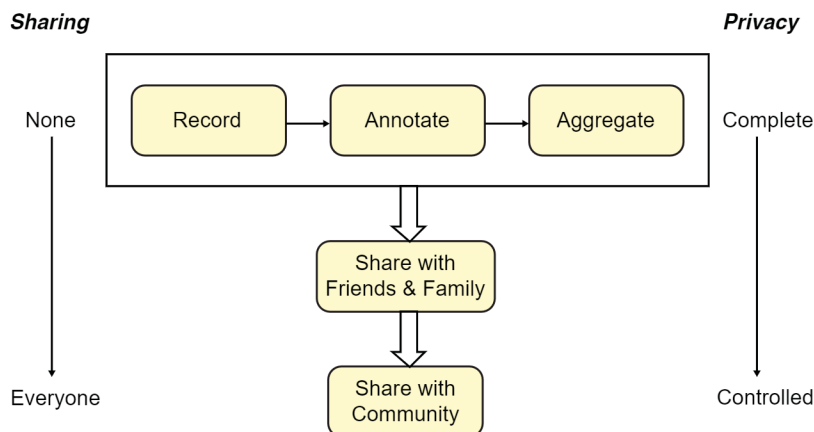


FIGURE 1 Personal experience aggregation and sharing: Drivers record, annotate, and aggregate personal data, then can choose to share with friends or a larger community.

Many systems appeal to a sense of altruism. Google Co-op invites users to manually annotate web pages to improve Google's search capability. Large fact-gathering communities, such as Wikipedia and many web forums, explicitly collect intellectual contributions from users. These efforts assume that many people will volunteer to contribute if they feel they can help make the world a better place.

Other systems motivate people to spend effort simply because it is fun to do so. Researchers at Carnegie Mellon have built games that can significantly improve computerized tools by enticing users to contribute knowledge. For example, to improve image search, von Ahn and Dabbish introduced the ESP game to collect and evaluate metadata about images (1). Building on that success, von Ahn et al. later developed Peekaboom, a game to determine the locations of objects in images (2).

Social networking websites such as www.myspace.com, www.linkedin.com, and www.facebook.com have built networks of millions of people that include personal experiences but do not analyze the data or aggregate it in meaningful ways. Systems like Google's

page rank, which estimates importance of web pages on the basis of others linking to them, use efforts previously spent. Other sites, such as Yahoo! Answers and Amazon Mechanical Turk, connect individual question askers with answerers, taking advantage of broad participation, but still do not aggregate the data or build structure from content.

Route-Planning Domain

For many years, researchers have developed systems that observe individual drivers to estimate road speeds. Cameras and inductive loop detectors embedded in the pavement routinely produce information about traffic speed, but the data are sparse and require significant infrastructure.

In the late 1990s, the Universities of Maryland and Virginia and several cellular phone companies pioneered efforts to use location data from cell phones to estimate travel time (4). Since, countless

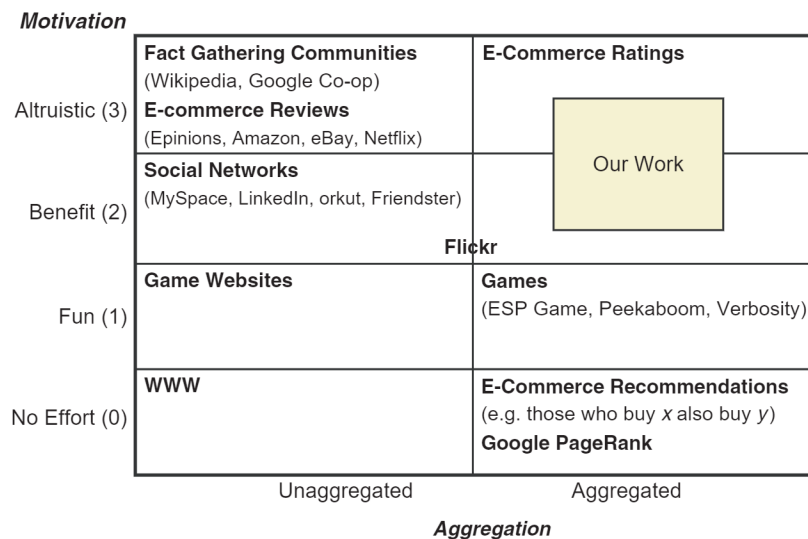


FIGURE 2 Taxonomy of existing systems: four levels of motivation to participate and two levels of aggregation.

studies have examined different ways to do this, using statistical techniques, artificial intelligence, and pattern matching (5). Fawcett and Robinson estimate road speeds for small time slices to predict the optimal route at a given time (6). Harrington and Cahill tag time records with weather and road condition information to predict travel times by context matching (7). However, these efforts tend to focus on improving traffic management and reducing traffic jams. They do not help users understand their personal experiences, aggregate those experiences, or add other useful information.

Some researchers have looked more closely at personal route planning. Konishi et al. described a scenario for selecting the best route by aggregating driving history and providing summary information, but there was no notion of annotation or subjective route attributes, and no interface was built (3). Letchner et al. developed the TRIP system to incorporate time-variant traffic data and user preferences into route planning (8). They collected Global Positioning System (GPS) traces from many drivers to estimate traffic speeds for small time slices, taking user preferences into account by suggesting routes similar to a driver's previous trips. However, they do not let users explicitly add information to the system or discuss how to combine data from multiple drivers.

SURVEY AND USER NEEDS

To inform study theories and get a better picture of what a personal-experience-collection system might look like, a web survey was conducted to identify what people think about their current driving activities and needs (www.cs.umd.edu/hcil/routelens/survey). The survey asked about current driving and route-finding habits, focusing on repeated routes, the criteria used for choosing those routes, and the information sources used for determining which routes to take. It also asked about the relative importance of average route speed versus consistency. Finally, it asked about participants' willingness to share data about their driving and route-selection behaviors and ended with some basic demographic questions.

Terminology

Some basic domain terms are defined here to explain how they are used in the survey. Route means a particular path or set of roads from one location to another. Trip refers to each time a user drives along a route. Route collection is used to designate all the different ways to get from one location to another. Finally, locations are the named end points of a route.

Respondents

Survey participants were solicited by sending an e-mail to colleagues and asking them to forward that e-mail to anyone they knew. In all, 292 people responded, covering various fields, including computer science, information technology (IT), software design, and health care. Roughly 67% were female, ranging in age from 18 to 75 years old, of which 53% were between the ages of 46 and 60, and 92% were from the United States. Other regions represented included Asia, Australia, Europe, the Middle East, and Canada.

Although the survey participants represent reasonably diverse ages and professions, it should be noted that this sample is self-selected, relatively small, and unrepresentative (statistically speaking) of all

drivers in the United States or elsewhere; however, it is made up of technology-savvy users who are more likely to use this sort of approach. Although these data are believed to be broad enough to be informative, they are not definitive. Nonetheless, because the study goal is to show only that at least some nontrivial set of drivers would be interested in subjective data for route finding, the authors believe that the survey meets those needs. The study goal is not to show that this approach would support a viable business today, for example, but that the approach is interesting and worth continued and deeper understanding.

Domain Analysis

Many respondents admitted that they did not know the best route for their common trips, with nearly 80% agreeing that a system to track personal driving history would be useful. One participant even commented, "The number of reasons this would be useful to so many people are too numerous to count here." Nearly 60% of respondents said they argue at least once in a while about which route is best. About 80% could think of at least three trips with alternate routes that they drive regularly, and 33% listed more than five. Finally, most respondents said they usually choose routes by considering alternatives and selecting the one that seems best, apparently by using their own judgment to make a decision on the basis of all the available information. As suspected, many people seem to think that personal routing is an important problem, and they often use personal experience to make better route choices.

Questions and Tasks

Choosing the right problem domain is not enough to build a successful personal experience-aggregation system. To motivate people to participate, the tasks supported by the system must be personally beneficial to them. Users were asked to rate how useful the initial tasks were and to suggest their own tasks.

The final study task list can be classified into several groups: personal, collaborative, and other. Personal tasks use an individual's own driving history alone, collaborative tasks involve information from other users, and other tasks involve some external information. The following are examples of personal tasks:

- Find the best route for given trip and departure time.
- Find the best departure time for a given trip.
- Find routes to places I have been but do not remember exactly where they are.
- View summaries (e.g., average speed, number of stops, and total time) for my routes.
- Find a route based on my personal driving style.
- Add my own information (e.g., road conditions, scenery, safety, and points of interest) to routes.
- Find a route based on my annotations (e.g., minimize highways and maximize safety).
- Find points of interest near a route.
- Find a route near points of interest.
- View a map showing only roads I take frequently.

Of course, many of these tasks also could be performed with data from family and friends or from other community users. However,

some tasks are more suited for collaboration than others. The following are examples of collaborative tasks:

- Get better statistics by sharing routes with others.
- View others' routes to find a good alternative.
- View others' route annotations.
- Determine the validity of others' annotations.
- Find routes to places I have not been before.
- Find routes similar to a particular route I like.
- Find people with similar routes for carpooling.
- Compare differences between drivers.
- Generate detailed directions to share with others.
- Provide interactive communication between drivers.

Finally, some tasks are neither personal nor collaborative. For example, 12 respondents independently commented that they would want to view real-time traffic information to choose alternate routes in response to exceptional circumstances such as accidents or construction. Although this information might come from annotations, it more likely would be drawn from external resources maintained by transportation authorities but customized to personal routes.

Route Criteria

When collecting personal experiences, it also is important to distinguish between data that can be tracked automatically and data that must be manually added to the system by users. The most common task was to find the best route according to some user-specified criteria. Thus, it is important to know what kind of data people care about when choosing a route. Respondents were asked to rank the importance of several route criteria and suggest their own. Results show that although people care a lot about the measurable attributes of a route, they also care about many more subjective attributes, including

- Distance,
- Travel time,
- Time stopped,
- Number of stops, and
- Driving speed.

Other attributes deemed important cannot be easily recorded automatically and need to be added as annotations. These kinds of attributes include

- Scenery,
- Safety,
- Fuel consumption,
- Road conditions (e.g., potholes),
- Nearby points of interest (e.g., cheap gas station), and
- Types of roads (e.g., back roads versus busy highways).

Some of these attributes may become automatable (e.g., fuel consumption could be calculated if the route tracking system were integrated with a car's on-board computer). However, there will always be subjective attributes that cannot be measured—like beauty or feelings of safety—which the survey indicates are important.

Sharing

Finally, it was important to know whether people were willing to share their personal data with others. Data sharing is crucial for collecting and aggregating individual driving experiences. Surprisingly, most respondents (~75%) said they would be willing to share some or even all of their data. Privacy is of course an issue, with one respondent stating, "I would not want criminals to have access to when I would not be home."

Most people wanted the ability to choose who to share with and whether to remain anonymous. About 35% said they would share with a group they could choose, and nearly one-half said they would share with anyone if they could hide their identity.

Respondents also wanted to control which data they shared with others. Most preferred to choose a subset of their favorite routes, annotations, and statistics; one participant said, "I would share any information that I did not classify as personal." Seventeen respondents independently commented that they would want to clip the end points of their routes, so others could not see where they came from or exactly where they were going. Many people were hesitant to tell everyone their tricks. One person said, "I would not share too many of my secret methods. . . . If everyone knew the exit chute, then it would not be quick." Another noted, "Sharing my secrets for alternate routes may adversely affect that route for me in the future—selfish, probably, but a survival technique in a heavily congested area." However, people wanted to share these tricks with their friends and family.

Finally, comments suggested that people were aware of the public benefit of sharing elements of personal experiences (e.g., road condition annotations or wait times at intersections), and had some altruistic motives to share knowledge of the routes (e.g., local history information).

In summary, the survey results indicate that people are willing to share if they can reap some personal benefit, help friends and family, or even serve the greater good. However, they want control over what to share, who to share with, and whether to remain anonymous.

A DRIVER-CENTRIC APPROACH

Following the personal experience aggregation process outlined earlier and informed by the survey, a set of methods and interaction scenarios were designed to motivate individuals to collect their personal driving experiences and share those experiences with others. An early interface prototype called RouteLens was built to illustrate how drivers might record, aggregate, and learn from their personal data. For annotation and sharing, mockups were presented.

Decentralized Data Collection

Current efforts to collect traffic information typically are started by transportation authorities using cameras, sensors embedded in the road, or cell phone information, relying on a great deal of supporting infrastructure. Soon, some cars will also automatically collect and report information about road conditions (e.g., speed, rain, and ice). In contrast, the study approach uses personal information collected by the drivers themselves on personal devices, maintaining complete control of privacy and requiring no infrastructure aside from GPS



FIGURE 3 Mobile interface: Simple buttons start and stop recording, labels display position, and map shows route (updating route view only partially implemented).

services. Each time an individual drives along a route, his or her personal mobile device (illustrated in Figure 3) tracks the location and computes trip distance, time, and speed as long as it is purposefully turned on. Note that the actual GPS tracking to create valid road-based routes is a well-studied problem but beyond the scope of this paper and is not discussed here.

Personal Data Analysis

View Routes

After users have recorded some data, they can view all their destinations and routes in a desktop visualization (Figure 4). The authors’ experience suggests that even viewing the shape of well-known routes can be useful and surprising.

Because route information is personal, users get other immediate benefits, even before aggregation or adding a single annotation. Consider this scenario: “John is going to a doctor’s office, where he has not been for months. He knows the general area of that office but does not remember the exact location. He asks to see where he has driven in the area and immediately recognizes the doctor’s office location.”

Uploading Multiple Trips

To aggregate statistics properly, the system must determine whether each trip recorded by the mobile device was driven along an existing route or a new route altogether. It could be done completely automatically with pattern-matching techniques, but the result is not accurate. Instead, the system makes a suggestion for each trip and lets users accept or correct those suggestions. When uploading trips, the interface suggests an existing route if a close match exists. The new trip appears as a dotted orange line and the suggested route blue, making visual matching easier. Users can hit a key to add the trip to the suggested route, select an alternate route, or create a new route altogether. Practically, many trips can be entered in a matter of seconds after common locations and routes have been uploaded. A “secret” trip or location can be erased if recorded by mistake. The authors’ impression is that most people do not seem to have locations to hide, and mobile devices can simply be left off. Stops along

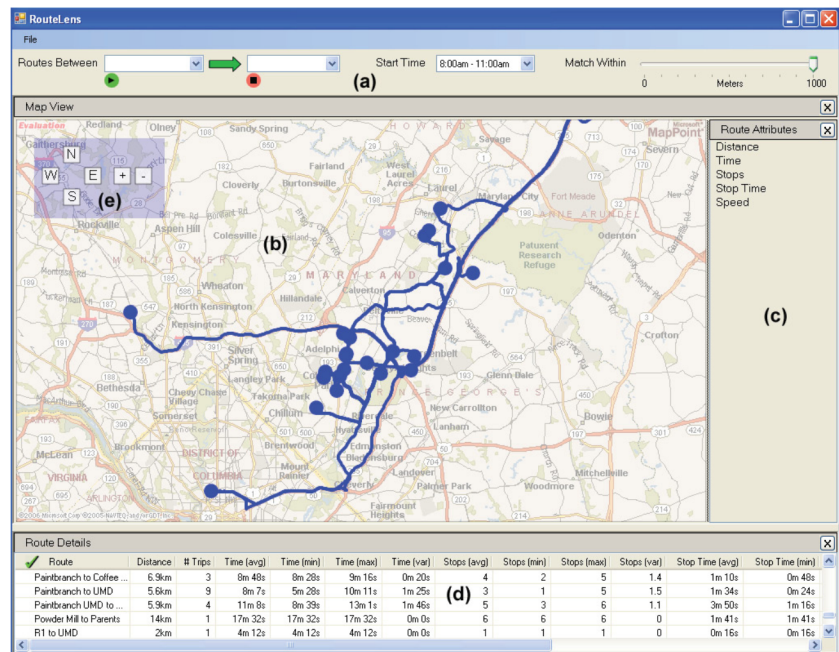


FIGURE 4 Desktop interface consists of components to (a) select routes, (b) browse data, (c) compare routes, (d) view detailed statistics, and (e) zoom and pan.

routes (e.g., for gas) can be removed automatically, or a whole trip can be ignored.

Learning from Aggregation

After users have collected enough trips, they can review summary information to compare differences in distance, travel time, number of stops, total stopped time, and speed. Locations are selected, and when an attribute is chosen, the map highlights the best route (Figure 5). Users can also make more subtle comparisons using variance: “Mark frequently drives to the coffee shop and wants to reduce his time stuck at traffic lights. He has used three alternative routes and selects Stop Time to compare them. He sees that the route that takes 495 to the coffee shop has the shortest average stop time but the most variance. On the other hand, the route that takes Paintbranch Road has only a few more seconds of stopped time on average and less variance.”

Other Information

One of the popular tasks from the study survey was to view real-time accident information. Although personal data (e.g., road conditions and locations of snow drifts and downed trees) might help answer this question, the desktop interface adds real-time traffic incidents from public sources that are located close to their personal routes. Potentially, statistics about accidents also could be added to the route information.

However, the display of real-time accident information as well as any other kind of data display that would be useful in the car for use while driving brings up the crucial issue of safety. As with GPS data interpretation, it is recognized that the safety issues are important—

but they have been studied extensively elsewhere and are beyond the scope of this paper because the kinds of solutions offered for other in-car systems will apply to the subjective kind of data suggested in this paper as well.

Annotation

Although annotations are not required for the system to be useful, the survey findings suggest that annotations increase benefits to users. First, users want to compare routes on the basis of subjective criteria such as beautiful scenery, ease of driving, or neighborhood safety. By adding such details to the data themselves, drivers will benefit from having the information they actually care about: “Wendy needs to drive to the train station and knows several routes. Viewing her annotations, she sees one route has ‘cheap gas’ and a ‘fast decent Mexican restaurant.’ Figuring she can stop for a quick bite and save some money, Wendy chooses that route.”

People might also record annotations to help them remember important information they want to share with friends or family. Consider the situation of explaining directions to a friend: “Jim is driving to the Greenbelt Community Center, where he will meet his friends for an art exhibit next week. While approaching a tricky turn he knows well, he records a note: ‘Curve right onto the exit ramp for the parkway, but as the ramp straightens, avoid exiting by merging one lane to the left.’ ”

There are several other reasons to record annotations. People may want to share their knowledge of local sites (“That green house is the oldest one in College Park”), record personal memories with family (“This is the hill where we took the kids sledding each year”), or share important notes with the community (“I wouldn’t drive here at night”).

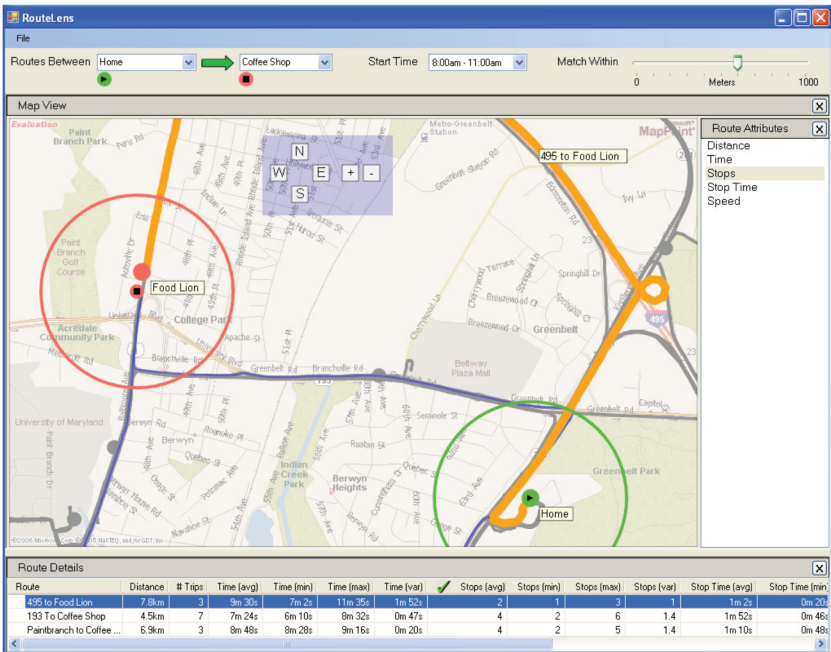


FIGURE 5 Attribute queries: The user selected “number of stops” to compare routes between two locations; the best route is displayed as a thick orange line.

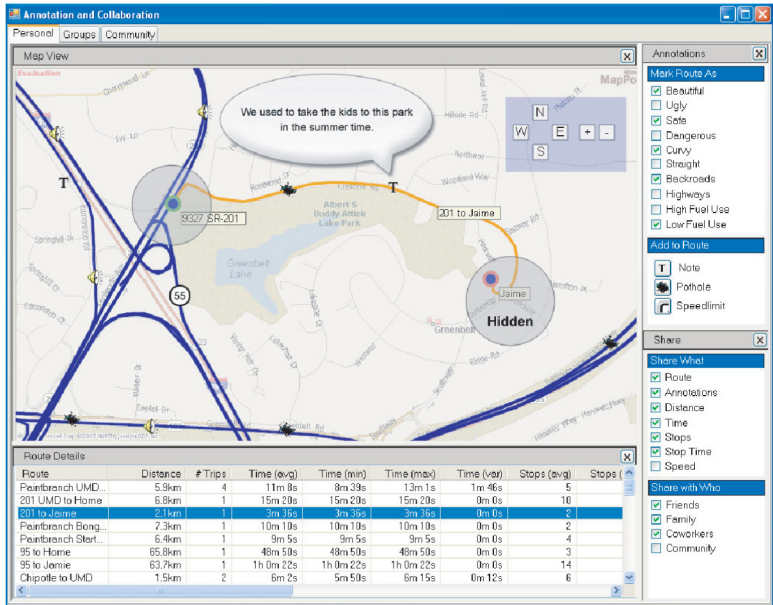
Classification

The annotations described here can be classified along several dimensions, including type, level, and localization. Two basic types of annotation are identified: structured and unstructured. For some route attributes, such as safety, scenery, or road type, it makes sense to provide a structured way for users to quickly rate or categorize their routes. Although these kinds of annotations are easy to aggregate, they are limiting. Users who enjoy annotating might add less-structured

notes, involving voice, pictures, sound, or free text: “After passing an area with dozens of trees down, Jerry presses the Add Comment button and says, ‘Whoa, check the damage of the July 4th tornado.’ It will be replayed to his friends and family ahead of the location when they drive in the area (and choose to hear annotations.)” Two additional levels of annotation are identified: en route (Figure 6a) and after traveling (Figure 6b). The most natural time to add notes often is en route. As with any device widely used in the car (phones, PDAs, or direction finders), interaction should be limited



(a)



(b)

FIGURE 6 Adding annotations: (a) en route, drivers—or, better, passengers—can add structured notes and voice comments on the mobile device and (b) after traveling, users can add structured notes and free text, edit annotations, and share data.

or done by a passenger while driving, or should wait until the vehicle has been parked. While at home, users might add or edit annotations (e.g., adjust locations of comments or add more detailed information to give directions to others).

Finally, annotations may have different degrees of localization, pertaining to an entire route (“Always lots of trucks”), a segment of a route (“Last mile is unpaved”), or a specific point along a route (“Drive 35—speed trap”).

Presentation

Just as there are multiple ways for users to record annotations, there are multiple ways to view and interact with recorded data. People might view annotations on a personal computer before driving or with a mobile device while on the road. The presentation will be different in each case and will depend on how localized the data are. Users might listen to general route annotations at home but want to be notified when approaching a tricky turn or a pothole while driving. The mobile interface must be restricted because of safety concerns and limited screen space. While driving, users should view only data and listen to or add voice annotations; when the car is stopped or if a passenger is present, users can perform more complex tasks.

Sharing

Survey results indicate that people are willing to share, but privacy is a key issue. The proposed solution is to make collaboration a multi-level opt-in task, letting people choose who to share with, what to share, and whether to remain anonymous. The system can guide users through a natural progression, from individual to community use. Drivers will benefit from recording their own personal information, but the more they are willing to share, the greater the benefit they will receive. They may start by interacting with just friends and family and later discover that they can learn more by collaborating with the rest of the community.

Who to Share With

Some people may be more likely to ask friends and family for advice than to consult public data repositories when planning unusual trips because of trust issues and shared knowledge of driving preferences. Recent work related to social network analysis shows that including information about the reliability of data sources can increase not only the perceived but also the actual quality of the information for a particular individual (9). One usage example is, “Mary needs to drive to Annapolis for a concert after work. She could MapQuest it, but instead she trusts that her coworker Anne, who commutes from Annapolis every day, will know the best route. Anne sends her recommended route and shares her historical data so Mary can decide when to leave.”

A driver also might want to leave messages only to a spouse (“Honey, monster pothole in the intersection”) or add annotations from a neighbor who has unique knowledge of landmarks: “Jane is driving with out-of-town family but can never remember which buildings are significant. She has turned on commentaries of her friend Isabelle, an architect-historian. As they pass by an old house,

everyone hears, ‘Montpelier Mansion, a good example of Georgian architecture that has hosted many famous people, from George Washington to Abigail Adams.’”

Although sharing route data with friends and family alone is useful and offers good control over privacy, drivers also will benefit from data aggregated from multiple users: “Sara is meeting a friend for dinner at a restaurant she has never been to. She searches the community data, telling the system to ‘minimize highways’ because she prefers back roads and to ‘maximize safety’ to avoid driving through any dangerous neighborhoods. Finally, she sorts the results by ‘number of ratings’ to choose the route with the highest confidence level.”

When more users share annotations with the community, public benefit increases. Transportation agencies can use aggregated reports of road conditions to spot problems. During an emergency, thousands of users could decide to share their reports of the conditions at their locations (e.g., downed trees, power outages, and fires), thereby dramatically increasing the information available to first responders.

What to Share

Drivers may not want to share all of their routes, statistics, and annotations. Someone who has a secret trick (e.g., where to change lanes) will share that information only with friends. People usually do not want to give away personal information or reveal their identity (even though there are many exceptions). It is important to provide mechanisms to allow users to limit what they share or to share data anonymously: “John wants to share his favorite routes and recorded statistics with his small-town community but does not want people to see how fast he drives or where he goes exactly. He chooses to share everything but removes travel speeds and clips all end points of the routes to hide them.”

Anonymity

Drivers also may want to conceal identity. They should be allowed to choose whether to remain completely anonymous, show some personal information, or reveal who they are. Again, the more information they share, the greater the benefit they will receive. For example, revealing personal information may help users identify people they have something in common with: “Bob is very anxious about privacy. He shares nothing at first. . . . Later on he decides to share aggregated data about route segments (i.e., between two intersections), never even revealing entire routes, to access similar data from others.”

In contrast, “Harry spills details of his life on his blog and loves to share all his notes on the routes he uses. He now wants to find someone to carpool with. He enters start and end locations and preferred departure time to search for people with a similar commute. After sorting the results by users with similar driving style and schedule to his, a message is sent to the few people who chose to identify themselves to their neighbors.”

These two scenarios bring up one challenge with this approach, which is the statistical validity of the aggregated data. Given that individual users are not likely to generate much subjective data, it is by nature going to be idiosyncratic. Surprisingly, this does not invalidate the utility of the data. If individuals are sharing with other known individuals, then the personal and idiosyncratic nature of

the data is precisely what is going to make the data trustworthy. Only when the data contributors are not as well known by the users will statistical validity be more valuable. As more people contribute data, however, the data become less personal and more easily statistically aggregated. So, these two potentially conflicting factors (trustworthiness of individuals vs. reliability of data) are likely work hand in hand with different use cases (i.e., personal recommendations vs. collective knowledge).

Accuracy

Because participants are self-selected and the data self-contributed, issues of selection and content bias—unintentional or malicious—could arise. Several methods are proposed to mitigate these problems:

- Minimize impact by having many participants,
- Detect and discard outliers,
- Add community ratings of users and content, and
- Automatically detect personal and group consistency.

People's perceptions of subjective attributes vary greatly. What one person considers frightening, another might find exciting. Therefore, collaborative filtering can provide a way to focus on data from people with similar sensibilities.

Presentation

Users should be able to search any subset of their personal and shared data to find a good route. As shown in Figure 6b, they can view routes, statistics, and annotations from other groups they share with or from the broader community. For the aggregated community data, the interface would show an overview first and let users filter down by entering criteria they care about.

Aggregation

Because the aggregation of different kinds of data from potentially thousands of users is particularly challenging, a range of approaches is suggested. For the route data and recorded statistics, the process should be fairly straightforward. After many users share their routes, there will be much overlap. One user may drive along a route from his house to the coffee shop that takes 495, whereas his neighbor may have a route to the Chinese restaurant next to the coffee shop that also takes 495. There probably is no need to show both of these routes in the community data. Pattern-matching techniques could be used to find similar routes and then combine the statistics to get more accurate summary information. After aggregated routes are found, aggregating structured annotations (e.g., safety or beauty) is fairly straightforward.

Unstructured annotations, such as free text, are more difficult to aggregate. Linguistic analysis techniques could be used to group similar notes by searching for keywords in context. For example, Cesarano et al. built the Opinion Analysis System (OASYS) to track worldwide opinions on matters of national security by searching news feeds for positive or negative keywords (10). However, such approaches can be problematic because a similar statement can be phrased in many ways (e.g., "This route is really dangerous" and "I

drove through several bad neighborhoods"). To improve aggregation, human skills could be used to supplement computational methods.

CONCLUSIONS

A framework for harnessing human experience has been presented to help individual drivers make better route choices. The framework was motivated by a survey and developed by implementing a partial prototype and mockups. The study approach differs from other route-planning systems by focusing on subjective human experience. Results have shown how driver experiences can be aggregated from the ground up, first by helping people record personal history, then by providing people with a way to share data with family and friends or even the broader community. Issues of maintaining anonymity and accuracy, the technical challenges of aggregation, and data presentation to the user have been discussed.

Although limited in generality, the survey collected data from a self-selected sample of individuals believed to represent those technology-savvy users most likely to be initially interested in the sort of the proposed system. Survey participants mentioned that they care a lot about subjective route criteria, whereas previous work clearly indicates that they also can benefit from automatically recorded information. They also want to have access to external data, such as real-time traffic information. This finding indicates that the best system would provide a combination of automatic, annotated, and other supplemental information, when available.

Survey results also indicate that many people are willing to share all data with friends and family and some things with the community. This finding provides strong evidence that experience-aggregation systems would be useful. Even if people share their data with only a circle of friends, they will reap significant benefits. If they share anything with the broader community, then the overall benefit to everyone will be even greater.

The general challenge of personal experience-aggregation systems seem to be to provide

- Immediate benefit, even without sharing information;
- Increased benefit through annotating and sharing information with selected users; and
- The potential to serve the greater good through global sharing of anonymous information.

However, the nature of this kind of system is that there is a "chicken and egg" problem in deployment: it does not become useful until there are many participants, but people are less likely to participate until it can provide value to them. Even so, plenty of systems whose primary value comes from network effects (e.g., the popular social network websites) have overcome this problem. If individually contributed subjective data can provide value when many people participate, then it is worth pursuing how to motivate them to do so.

A framework was described for collecting individual subjective data, but much more remains to be studied to ensure the excellence of such a system. Future work might include a more complete implementation of the proposed route-planning system. The development of a general infrastructure to support a set of common tasks also could be envisioned. It is hoped that the design investigation reported in this paper will inspire others to collect, aggregate, and communicate human experiences.

ACKNOWLEDGMENTS

The authors thank John SanGiovanni for input and support, Ben Shneiderman for contributing the original idea of recording and comparing routes, and Makoto Konishi for original development of that idea. This work was supported in part by Microsoft Research.

REFERENCES

1. von Ahn, L., and L. Dabbish. Labeling Images with a Computer Game. In *Proc., SIGCHI Conference on Human Factors in Computing Systems (CHI 2004)*. ACM Press, New York, 2004, pp. 319–326.
2. von Ahn, L., R. Liu, and M. Blum. Peekaboomb: A Game for Locating Objects in Images. In *Proc., SIGCHI Conference on Human Factors in Computing Systems (CHI 2006)*. ACM Press, New York, 2006, pp. 55–64.
3. Konishi, M., C. Plaisant, and B. Shneiderman. Enabling Commuters to Find the Best Route: An Interface for Analyzing Driving History Logs. In *Proc., IFIP Conference on Human-Computer Interaction (Interact2001)*. International Federation for Information Processing IOS Press, Amsterdam, Netherlands, 2001, pp. 799–800.
4. *Probe-Based Traffic Monitoring State-of-the-Practice Report*. NCHRP Project 70-01. Virginia Transportation Research Council, University of Virginia Center for Transportation Studies, Charlottesville, Va., 2005.
5. Chung, E. Classification of Traffic Pattern. In *Proc., 10th ITS World Congress, 2003* (CD-ROM), IBEC, Birmingham, United Kingdom, 2003.
6. Fawcett, J., and P. Robinson. Adaptive Routing for Road Traffic. *IEEE Computer Graphics and Applications*, Vol. 20, No. 3, 2000, pp. 46–53.
7. Harrington, A., and V. Cahill. Route Profiling: Putting Context to Work. In *Proc., Symposium on Applied Computing (SAC 2004)*. ACM Press, New York, 2004, pp. 1567–1573.
8. Letchner, J., J. Krumm, and E. Horvitz. Trip Router With Individualized Preferences: Incorporating Personalization Into Route Planning. In *Proc., Conference on Innovative Applications of Artificial Intelligence (IAAI-06)*. AAAI Press, Menlo Park, Calif., 2006.
9. Golbeck, J., B. Parsia, and J. Hendler. Trust Networks on the Semantic Web. In *Cooperative Information Agents VII* (M. Klusch et al., eds.). Springer Verlag, New York, 2003, pp. 238–249.
10. Cesarano, C., B. Dorr, A. Picariello, D. Reforgiato, A. Sagoff, and V. S. Subrahmanian. Opinion Analysis in Document Databases. In *Proc., AAAI Spring Symposium on Computational Approaches to Analyzing Weblogs*. AAAI Press, Menlo Park, Calif., 2006, pp. 21–26.

The Geographic Information Science and Applications Committee sponsored publication of this paper.