

Trajectory-Aware Mobile Search

Shahriyar Amini^{1,2}, A.J. Bernheim Brush¹, John Krumm¹, Jaime Teevan¹, and Amy Karlson¹

¹Microsoft Research

One Microsoft Way, Redmond, WA

{ajbrush,jckrumm,teevan,karlson}@microsoft.com

²Carnegie Mellon University

5000 Forbes Avenue, Pittsburgh, PA

shahriyar@cmu.edu

ABSTRACT

Most location-aware mobile applications only make use of the user's current location, but there is an opportunity for them to infer the user's future locations. We present Trajectory-Aware Search (TAS), a mobile local search application that predicts the user's destination in real-time based on location data from the current trip and shows search results near the predicted location. TAS demonstrates the feasibility of destination prediction in an interactive mobile application. Our user study of TAS shows using predicted destinations to help select search results positively augments the local search experience.

Author Keywords

Mobile local search; mobile systems; location prediction;

ACM Classification Keywords

H.5.m [Information Interfaces and Presentation]: Misc.;

INTRODUCTION

Mobile applications often take advantage of user location to provide a localized experience. However, most only use a single snapshot of the user's current location. When users are moving (e.g., in a car), their current location can quickly end up behind them. Some tools, like GPS travel systems, are designed to support mobility, but require an explicit destination, which their users specify for only about 1% of their trips [6]. In this paper, we use location trajectory and destination prediction to implicitly incorporate information about where users are going, and not just where they are.

The application we explore is mobile local search. *Local searches* are searches for places with a geographic location; common examples include restaurants, gas stations, stores, or area attractions. A recent survey of 929 mobile searchers found that respondents often search while in motion (e.g., in a vehicle or walking), and that more than half of the time searchers wanted their results to be near their route or destination rather than their current location [9]. This suggests using the searcher's location trajectory could help find more relevant results than using current location alone.

Although destination prediction algorithms have been proposed and validated [e.g., 5], we are unaware of any

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI '12, May 5–10, 2012, Austin, Texas, USA.

Copyright 2012 ACM 978-1-4503-1015-4/12/05...\$10.00.

user-focused applications that incorporate them. Amin et al. [1] show users conduct searches most often along regularly used routes and near regularly visited places where users would not be likely to specify a destination. A few prior projects have proposed or looked at using the user's speed and direction. Gui et al. [3] mention that a user's speed could be used as a ranking feature for a context-aware proxy for mobile search, but do not describe exactly how this might be done. COMPASS, a mobile tourist guide prototype, uses user speed to present relevant points of interest [8], and Console et al. [2] consider speed and direction of travel to offer points of interest with respect to what can be reached and what is in the direction of the trip. These works focus on the use of location context for tour guide applications. Our work is unique in that we build and study an interactive trajectory-aware search prototype, and we use destination prediction within the search domain.

To demonstrate the value of incorporating destination prediction into mobile search applications, we implemented and studied the Trajectory Aware Search (TAS) prototype. TAS takes multiple samples of a user's GPS location to establish the user's trajectory and predicts a likely destination that is used to identify appropriately located search results. Our trajectory prediction method is based on an algorithm that is both user and history agnostic [5], meaning it is immediately useful to first-time users and in unfamiliar locations. It also mitigates privacy concerns by only using location data from the current trip.

We studied the use of TAS with a restaurant search task. Our studies in the lab (N=18) and car (N=6) found that restaurants suggested using destination prediction were well-located and relevant, in particular during the first few queries of a search session. We conclude with a discussion of how trajectory data could be used to further improve mobile search and other mobile applications and services.

TRAJECTORY-AWARE SEARCH PROTOTYPE

Given the popularity of restaurant searches (representing 45% of all local mobile searches [9]) we designed the TAS prototype to support local restaurant searches. Figure 1 shows screenshots of the TAS Mobile Client, implemented for Windows Phone 7. TAS has a textbox to input search queries (e.g., *pizza*), and an interactive map displaying results as flags. Results are fetched from the server once the user sends a query. Clicking on a particular flag provides more information about the associated restaurant at the bottom of the screen. Tapping on the description takes the user to another screen with restaurant reviews.

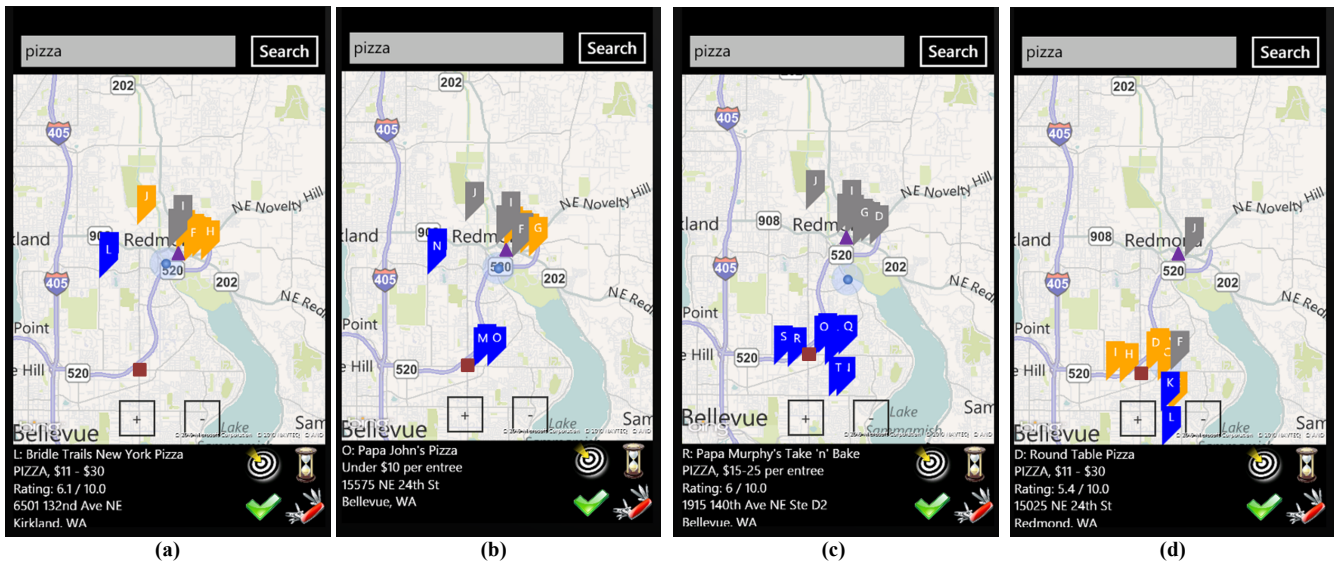


Figure 1. Trajectory-Aware Search Interface. User is traveling from the purple triangle to the red square with their current location shown by a blue dot. Trajectory results are shown in blue, conventional results are gray, and results returned by both are orange.

TAS discretizes space into square cells. The TAS application samples the GPS sensor every 5 seconds and determines the cell that the user is located in (see Figure 2). When the user issues a query, the mobile client sends both the query and the sequence of cells that the user has traversed to the TAS server, which performs destination prediction and retrieves two sets of search results from Bing™ Local: 1) *Conventional* results, based on the user’s current location, and 2) *Trajectory* results, based on the user’s predicted destination. By building on top of an existing search engine, we prevent the need to directly modify the search algorithm and provide a solution for how legacy location-aware services may be enhanced through destination prediction.

TAS presents the Trajectory and Conventional results using blue and gray flags. Occasionally the two result sets overlap, shown as orange flags. Figure 1 shows the series of results that occur as a user moves from an initial location (the purple triangle) towards the destination (the red square), with the blue circle indicating the user’s current location. Figure 1a represents the start of the trip, so the predicted destination is very similar to the current location and most of the results are in both sets (orange). As the user moves, the separation between the Conventional (gray) and Trajectory (blue) results becomes evident. Almost half-way through the trip (Figure 1c), the Trajectory results converge around the user’s true destination. The Conventional results, however, remain at the user’s starting location. This is because there are no restaurants near the halfway point of the trip, so the Conventional results continue to show restaurants from the starting location. Once the user is almost at the destination (Figure 1d), the Conventional search results again converge on the user’s destination.

Destination Prediction

Destination prediction is based on prior work by Krumm [5]. The approach assumes drivers choose efficient routes to

reach their destination. A discrete grid of cells is placed over a region of interest, such as the metropolitan area shown in Figure 2. For each trip, the user’s path up to the current time is recorded as a sequence of cells. Based on how efficiently the user has approached any particular cell in the grid, the algorithm computes the probability of that cell being the user’s destination. Cells to which the driver is taking an efficient route are considered more likely destinations than cells the driver has apparently decided to pass up. We limit prediction to locations reachable within 30 minutes from the trip’s start, since a majority of driving trips take less than 30 minutes [4]. Figure 2 shows a partial trajectory going south through four cells. The higher probability cells, shown with darker outlines, are those to which the partial trajectory is an efficient route.

TAS contributes a novel application of Krumm’s algorithm to generate a trajectory-aware local search experience in real-time. To conform to existing local search APIs, which expect a single GPS location and radius, the prediction

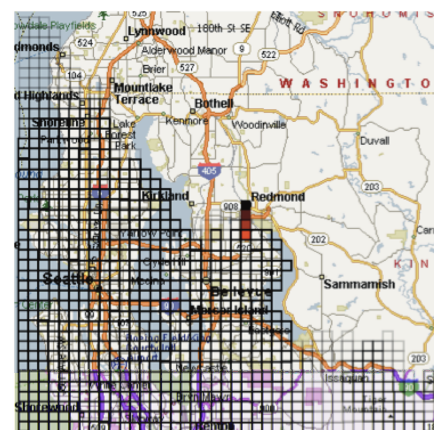


Figure 2. An example of the prediction grid after a driver has traversed south through the four filled cells. Cells with higher probabilities given that route have a darker outline.

algorithm computes the spatial centroid and standard deviation of the cell probabilities, essentially fitting a 2-D circularly symmetric Gaussian to the computed probability distribution. We use a circle with a radius of one standard deviation centered at the centroid of the Gaussian as the search region for the predicted destination.

USER STUDY METHODOLOGY

We evaluated TAS by asking participants to select a restaurant for lunch while en route to a destination. We provided participants Samsung Focus Windows Phones running the TAS Mobile client. Each participant received an introduction to the mobile client interface, and was told that we were comparing different local mobile search algorithms, the results of which were represented by different colored flags. We varied the mapping between flag color and search result type across participants to avoid a potential selection bias based on flag color. To motivate participants to choose a restaurant they actually wanted to eat at, they were told they would receive a gift certificate to one of the four restaurants they selected.

We tested TAS in a controlled laboratory environment with 18 participants (9 M, 9 F). We simulated four trips to nearby suburban centers using pre-recorded GPS data. During pilots we found the GPS playback felt slow, so we played it back at double speed. In addition to these lab studies where we could ensure each participant experienced the same “trip,” we also drove 6 participants (3 M, 3 F) to the same four suburban centers following the path of the pre-recorded GPS to evaluate the use of TAS in-situ. The trips studied were approximately 15 minutes long driving. In our analysis, C denotes an in-car participant, while L denotes lab. All participants worked at Microsoft and had previously used mobile search. We compensated participants with either a \$10 (lab) or \$25 (car) gift card.

We collected qualitative data about the participants’ experiences and logged interactions with TAS. Following each task, participants filled out a brief questionnaire asking them about the restaurant they chose (e.g., had they heard of or been to it before?), the location of the search (e.g., was it familiar and were the results nearby?), and their search experience (e.g., was the experience satisfying or frustrating and did they have enough choice?). All participants chose 4 restaurants in 4 locations, for a total of 96 restaurants across 24 participants. Logging failed for two participants (1 car male, 1 lab female). The logged participants issued a total of 464 queries.

USER STUDY RESULTS

Overall, we found that participants found a restaurant that they would like to eat at faster with Trajectory search than they did with Conventional search, and that they were more satisfied with the Trajectory search results. The results identified via destination prediction appeared particularly useful early in the search process, when the destination was some distance away, and for participants who employed exploratory search strategies by entering short, vague queries, versus long, targeted queries.

Target Restaurant Found Sooner with Trajectory Search

The process of selecting a specific restaurant to eat at (referred to as the *Selected restaurant*) took most participants about five minutes, and, on average, 5.3 queries. When we look at the first query within the search session to return the Selected restaurant somewhere in its result list, we observe that 56% of the time the restaurant appeared in both the Trajectory and Conventional results sets. However, when the Selected restaurant was initially found by only one approach, the restaurant was more likely to be found in the Trajectory results (32%) than the Conventional results (13%). Participants were equally familiar with the areas where the Selected restaurant was first found in Conventional or Trajectory results (Median response “Familiar”). However, they were more familiar when the restaurant was returned in both sets (Median response “Very Familiar”). This was particularly true for the in-car studies. Of the 20 search tasks logged in the car studies, 50% of the time the Selected restaurant was first returned only in the Trajectory results. The rest of the time the Selected restaurant was returned by both the algorithms. We conjecture the two search result sets being distinct longer in slower-paced car studies account for this fact.

Participants More Satisfied with Trajectory Search

Trajectory search also led to a more satisfying experience. When a participant’s Selected result was first found in the Trajectory results, they reported having a more satisfying search experience (median 3 out of 5, 1 = least satisfying) than when the Selected restaurant was first returned in the Conventional results (median 2). Participants’ comments supported the relative lack of value they found in Conventional results. For example, C2 said she “never saw the [conventional] results,” and C4 said “I never picked [conventional], I did not see many.” In contrast, C5 commented after finding a restaurant returned in Trajectory results: “It gives plenty of options at [my] destination vs. where we are presently.”

Trajectory Search Useful Early in the Search Process

During the study we observed that Trajectory search was particularly useful to our participants early in the search process, while they were still far enough from their destination that Conventional search could not retrieve results near where they were heading. Our search logs corroborate this observation, revealing that the Conventional and Trajectory results were most distinct during the first query and became more similar as participants’ tasks progressed. C5, who selected restaurants returned first in Trajectory results for all four suburbs, reflected, “I chose while we were not that close to destination,” indicating his use of Trajectory results to select restaurants ahead of him.

Trajectory Search Supports Exploration

We also found that the participants who successfully used the Trajectory approach often adopted an exploratory approach, issuing short, general queries. The average query length when the Selected restaurant was first found in the

Trajectory results was much shorter than those issued when results were first found with only the Conventional algorithm (1.2 words vs. 2.5). L2's comment, "I began to trust that I didn't have to specify location because I was 'close enough' even if we were going between towns," suggests users also noticed that the correct prediction of trajectory allowed them to type shorter queries.

In order to examine how exploratory the queries participants used were in different conditions, we manually coded the 147 queries that returned the Selected restaurant in the results by whether the query contained:

- location information (e.g., *[city name] pizza*),
- the name of a specific restaurant (e.g., *pizza hut*),
- a restaurant genre (e.g., *italian*), or
- generic (e.g. *restaurant* or *food*).

Given the TAS focus on location, if the query referenced a location it was categorized as such, regardless of other content. One author coded all of the queries and another coded 40%, with 100% agreement. We observed that 87% of the time when the Selected restaurant was returned only in the Trajectory result set, the participant used a generic or restaurant genre query to find it, compared with only 40% of the time when it was returned in the Conventional result set. In contrast, the queries used to find the Selected restaurant in Conventional results were much more likely to include a location or the name of a specific restaurant (doing so 60% of the time) than those used to find it in the Trajectory results (doing so only 12% of the time).

Overall, the search patterns we observed by participants whose Selected restaurant was first found in the Trajectory results exemplify the scenarios for which destination prediction may have highest value: when the issuer does not have a particular restaurant in mind, and when the destination is not specified (because the user does not want to bother, cannot name the destination, or would accept several destinations along the route). In Teevan et al.'s survey, 31% of mobile-issued local searches had the characteristics that the searcher is both moving and searching for a non-specific result [9]. We expect trajectory-augmented search results to have the highest potential value for these searches.

CONCLUDING REMARKS

By building and studying TAS, we demonstrated destination prediction can usefully augment mobile local search. Identifying results near a mobile user's predicted destination appears to be a simple means of augmenting and personalizing conventional local search. Our experience also suggested to us other ways destination prediction could enhance the mobile search experience. For example, trajectory information could be used to inform auto-complete query suggestions, thus reducing the amount of user input. Destination prediction could also help determine the most appropriate focus and zoom level of the map to display results (e.g., perhaps centering the map so that the predicted destination is visible).

We believe using trajectory prediction to enhance mobile search works well because the destination prediction does not have to be perfect to usefully restrict the search region. A range of other mobile applications could likewise benefit from prediction; for example, location based reminders could be presented in advance of reaching a location [7] and traffic warnings could enable drivers to an alternate routes.

More generally, our design and implementation of TAS demonstrate a feasible approach for evolving a legacy location-based service to become trajectory-aware. Specifically, we did not make any changes to the underlying search engine to produce a trajectory-aware search engine. With respect to the vast number of location-based services available, we envisage other services that could easily provide trajectory-awareness without modifications to the underlying service (e.g., a friend finding application might show friends near a predicted destination rather than a current location). Based on our experiences with TAS, we see tremendous potential in using richer location information in mobile applications.

REFERENCES

1. Amin, A., Townsend, S., Ossenbruggen, J. and Hardman, L. Fancy a Drink in Canary Wharf?: A User Study on Location-Based Mobile Search. In *Proc. INTERACT 2009*, 736-749.
2. Console, L., Gioria, S., Lombardi, I., Surano, V. and Torre, I. Adaptation and personalization on board cars: A framework and its application to tourist services. In *Proc. Adaptive Hypermedia and Adaptive Web-Based Systems 2002*, 112-121.
3. Gui, F., Adjouadi, M. and Rische, N. A contextualized and personalized approach for mobile search. *Advanced Information Networking and Applications Workshops*, 2009, 966-971.
4. Hu, P.S. and Reuscher, T.R. Summary of travel trends: 2001 national household transportation survey. *Oak Ridge National Laboratory Technical Report ORNL/TM-2004/297*, 2004.
5. Krumm, J. Real time destination prediction based on efficient routes. Paper 2006-01-0811. *SAE International*, 2006.
6. Krumm, J. Where will they turn: Predicting turn proportions at intersections. *Personal and Ubiquitous Computing*, 2009.
7. Ludford, P., Frankowski, D., Reily, K., Wilms, K., and Terveen, L. Because I carry My Cell Phone Anyway: Functional Location-Based Reminder Applications. In *Proc. CHI 2006*.
8. Setten, M.V., Pokraev, S. and Koolwaaij, J. Context-aware recommendations in the mobile tourist application COMPASS. In *Proc. Adaptive Hypermedia and Adaptive Web-Based Systems*, 2004, 235-244.
9. Teevan, J., Karlson, A., Amini, S., Brush, A.J. and Krumm, J. Understanding the importance of location, time, and people in mobile local search behavior. In *Proc. MobileHCI 2011*.