

# Voting with Your Feet: An Investigative Study of the Relationship Between Place Visit Behavior and Preference\*

Jon Froehlich<sup>1</sup>, Mike Y. Chen<sup>2</sup>, Ian E. Smith<sup>2</sup>, and Fred Potter<sup>1</sup>

<sup>1</sup> Department of Computer Science and Engineering  
University of Washington, Seattle, WA, USA  
{jfroehli, fpotter}@cs.washington.edu

<sup>2</sup> Intel Research, Seattle, WA, USA  
{mike.y.chen, ian.e.smith}@intel.com

**Abstract.** Real world recommendation systems, personalized mobile search, and online city guides could all benefit from data on personal place preferences. However, collecting explicit rating data of locations as users travel from place to place is impractical. This paper investigates the relationship between explicit place ratings and implicit aspects of travel behavior such as visit frequency and travel time. We conducted a four-week study with 16 participants using a novel sensor-based experience sampling tool, called My Experience (Me), which we developed for mobile phones. Over the course of the study Me was used to collect 3,458 in-situ questionnaires on 1,981 place visits. Our results show that, first, sensor-triggered experience sampling is a useful methodology for collecting targeted information in situ. Second, despite the complexities underlying travel routines and visit behavior, there exist positive correlations between place preference and automatically detectable features like visit frequency and travel time. And, third, we found that when combined, visit frequency and travel time result in stronger correlations with place rating than when measured individually. Finally, we found no significant difference in place ratings due to the presence of others.

## 1 Introduction

Why do we travel to some places but not others? What do these places say about our interests? Could a person's movements to and from places in the physical world be an implicit form of expressing preference? We studied the travel routines of 16 participants over the course of four-weeks to determine what factors of visit behaviors, if any, could be used to infer preference for places. Using GSM-based location sensing and experience sampling on mobile phones (a technique to capture self-report data from participants in situ), participants provided explicit ratings for the places they visited. We used these ratings to explore the correlation between place preference and two implicit aspects of place visit behavior, *visit frequency* and *travel distance*. In

---

\* This work was approved by the University of Washington Human Subjects Division, application id: HSD# 05-6963-E/C 01.

addition, we looked at the impact of social effects on place ratings in an attempt to discover whether the presence of others makes determining place preference more difficult.

Understanding individual user preferences for items is critical to many mainstream commercial online systems, but has been under-explored in physical world settings. Online applications like the Amazon.com® book recommender, the Last.fm® music-based social network, and the TiVo® personalized TV guide depend heavily on observing human action to determine preference. Often, these inferences are augmented with explicit user ratings. With TiVo, for example, the only explicit feedback is the “thumbs-up/thumbs-down” button. However, when a user records a TV show, that show is automatically assigned a “thumbs-up” rating—thus, ordinary user activities are paired with explicit ratings to form a profile of likes and dislikes [1]. The benefit of explicit ratings is that they are precise and fairly well understood; however, they also interrupt the user experience and may not be consistently used across users [6]. Thus, recommenders and personalization engines typically rely on inferring preference from ordinary behaviors, as TiVo does when a user records a show, Last.fm does when a user listens to music and Amazon does when a user makes a purchase. These implicit indicators remove the burden of explicitly having to rate an item. In addition, as [1, 6] point out, nearly every interaction with the system becomes a potential indicator of interest. In this paper we seek to determine if this type of implicit preference inferring translates to behavior in the real world.

We explore how well *easily* detectable attributes of visit behavior correlate with explicit place ratings. If these attributes correlate positively, modern location sensing technologies (e.g., a mobile phone equipped with assisted GPS) could use inferred interests based on visit activity to build preference profiles of users. This has high value implications for a variety of mobile applications. Real world recommendation systems, online city guides, and personalized mobile search, could all potentially benefit from learning place preferences. For example, a recommendation system could be constructed to give personalized recommendations based on a user’s movement—a tourist who makes frequent visits to Italian restaurants in his or her hometown is given a list of the top Italian restaurants on a mobile device when traveling to a new city. These recommendations could even span the physical and virtual worlds. Frequent visits to a live jazz music club might indicate a musical preference that could be used by an online music store. Alternatively, user experience could be enriched through an online city guide like Citysearch via personalized portal pages, event information and special discounts based on inferred interests. Finally, mobile search engines could return personalized local search results based not just on current location but also on previous visit behaviors and inferred interests.

One difficulty with using place visits for prediction is that people move around for a variety of reasons, not all of which can be predicted or modeled. For example, people often defer their tastes for the sake of convenience (e.g., I don’t like fast food but was short on time) or the presence of others (e.g., all-day shoe shopping spree with spouse). To better understand these complexities and their impact on place ratings, two potential confounds are also investigated: social effects and the “convenience factor.” In particular, we examined how place ratings changed when the participant was close to a location rather than far away. We also looked at how ratings differ when the participant was alone versus when they were with others.

In our analysis, we found that visit frequency and travel time as single factors were only slightly correlated with place ratings. However, when combined they became more effective predictors of preference. In addition, the results of social effects were mixed—in the general case, no significant difference was found for places visited alone versus with others. However, when with others, participants rated a place significantly higher if they chose to go to that place themselves (i.e., the decision of where to go was *not* made by someone else in the group). We also show through the perspective of a small case study how simple, pragmatic splits of data can result in much higher correlations between visit activity and explicit place ratings.

The rest of the paper is organized as follows: Section 2 provides background on our study techniques; Section 3 describes our study and participant background; Section 4 outlines some high level results including the number of ESM surveys and places captured; Section 5 analyzes the results from our field study; Section 6 discusses implications; Section 7 describes our related work; Section 8 discusses potential for future work and Section 9 presents our conclusions.

## 2 Background on User Study Techniques

Before describing the details of our study, we offer brief backgrounds on the primary user study techniques that we employed: experience sampling and web diaries.

### 2.1 Experience Sampling Background

In situ, self-report procedures such as the experience sampling method (ESM) have been used extensively in psychology and HCI to capture data on participants' thoughts, feelings, and behaviors as they are experienced [2, 8, 20]. Such procedures have a distinct methodological advantage over *ex situ* inquiries in that they do not rely on the reconstruction of information from memory, but rather involve reporting on experiences as they occur, thus minimizing recall bias. Traditionally this has been done with beepers and small booklets of paper-form questionnaires. The questionnaires would be carried and filled-out by participants when signaled by the beeper. This allowed researchers to get a sample of participants' experiences throughout the course of a study. Mobile computing has allowed this process to be computerized [2, 7]. One weakness of traditional ESM is that the beeper alerts may not always occur at relevant points of interest for the researchers.

In our study, we used a form of computerized ESM called “context-triggered sampling” to provide more targeted sampling of our participants. This technique, pioneered by MIT's Context-Aware Experience Sampling tool for the PDA [12], uses sensors to infer context to trigger a brief survey. It has several advantages when compared with traditional sampling methods, such as random or time-based triggering. For example, context-triggered surveys are much more likely to occur during events that are of interest to the researcher. This reduces disruption by decreasing the number of extraneous prompts on the participant. Additionally, context data can be continuously saved by the computer; allowing the researcher to cross-check answers with sensor data and perhaps uncover behavioral patterns not initially considered.

## 2.2 Web Diaries

Web diaries are often used in field studies [4, 21] to capture qualitative accounts of participants' thoughts and behaviors. Participants are asked to connect to a predetermined web site to fill out open-ended or semi-structured questionnaires at specific intervals. Sometimes the web site is used to upload pictures, recordings or other media captured by the participant [4]. We used web diaries in our study as a qualitative supplement to the quantitative data gathered through ESM (see Section 3.1).

## 3 Study Design

In our four-week, in situ study, we investigated the perceptions and feelings participants had about the places they visited. In this section, we describe our study design.

Our study consisted of three phases: Phase I: Participant backgrounds, Phase II: In situ experience sampling and Phase III: Study wrap-up. Phase I took place in our lab with one to five participants at a time and familiarized us with the participants' backgrounds and their general routines with respect to out-of-the-house activity. Participants were given an overview of the study and a training session on the technology they would use in Phase II. We administered two paper questionnaires: one to evaluate their visit behaviors and why they visited the places they did, and a second to determine their estimates of the number and variation of places they visited. Participants completed the second questionnaire (called the "My Places Questionnaire") on their own time and returned it at their second session in our lab.

In Phase II, sensor-triggered experience sampling was used to capture the places that the participants actually traveled and their subjective feelings for those places. Each participant was loaned an Audiovox® SMT 5600 mobile phone loaded with our novel sampling software, the My Experience (Me) Tool, which prompted participants to fill out a brief survey up to 11 times per day for four-weeks. In addition, participants kept a web diary to supplement the experience sampling surveys and validate features of our GSM place-tracking algorithm. Given the data input constraints of the mobile phone, we used web diaries to gather richer qualitative data on participants' places. At the midpoint of Phase II, participants returned to the lab for a one-on-one interview about their experiences using ESM so far.

In Phase III, participants returned to the lab for their final visit at the end of the four-week ESM period. A concluding interview was conducted that explored participants' attitudes and experiences with online web sites like Citysearch and Amazon and asked follow-up questions based on their ESM and web diary responses.

### 3.1 Experience Sampling and Web Diary Details

To implement experience sampling in our study, we built a generic, context-aware experience sampling tool called the "Me" (My Experience) Tool. This .NET based tool allows researchers to conduct in situ field studies using Windows Mobile® devices (e.g., PocketPC, SmartPhone, etc.). Similar to computerized sampling tools of the past [2, 7, 12], Me allows non-technical researchers to build computerized

self-report surveys simply by creating a textual input file. And, like CAES [12], Me supports context-triggered sampling. Unlike prior work, however, Me is the first tool to offer these capabilities on the mobile phone. Mobile phones are attractive data collection platforms for researchers because they offer real-time data connectivity, a small form factor, relatively long battery life, and an interaction paradigm familiar to most participants (e.g., keypad interactions). Of course, the mobile phone may not be suited for all studies as the small screens and buttons place certain physical demands on the participants. Me also runs on the PocketPC and PocketPC Phone platforms, which offer larger screens, higher resolutions, and touch panel interactions.

In our study, we built a software mobility sensor based on GSM signals received by the mobile phone [16]. ESM surveys were triggered by detecting when a participant's phone was stationary for a period of 10 minutes. Thus, once the phone shifted from "mobile" to "stationary" and remained in that state for the time threshold, a survey was triggered—see Fig. 1. If the phone remained stationary for longer than one hour, a random time-based survey would trigger. No two surveys were closer than 15 minutes apart. In addition to providing a fail safe against sensor failure, these time-based surveys allowed us to ensure that surveys were spaced throughout the day and that our quota of at least eight surveys per day was consistently met.



**Fig. 1.** Sample screenshots of the Me Tool on the Audiovox 5600 mobile phone used during the study

Our ESM surveys asked from one to twenty-three questions based on the participant's responses. A one-to-four question survey simply asked "Are you still at <last known place>?" If the participant responded "yes" that survey would end; if "no," at least three more questions were asked: "Place name," "Place category," and "Please rate how much you like this place." The rating scale was from 0.5 to 5 (5 is best), shown in Fig. 1 (right). To select a place name, participants could choose from previously entered names or enter a new one. This was the only ESM question that required text entry on the phone. The survey ended if the participant selected *Personal*, *Work*, or *Home*—all other responses would cause the survey to continue.

In addition to ESM surveys, participants were asked to fill out web diaries a few nights a week. These diaries were designed to capture information about place visits that were missed by the phone and augment the ESM responses with more qualitative data. Place visits were automatically uploaded to the website. Participants were then asked open-ended questions about a random subset of places marked as bars, cafes, restaurants or stores.

### 3.2 Participant Profiles and Compensation

We recruited 16 participants, 8 male/8 female, from the Seattle area through flyers posted at local restaurants, cafés, and apartment buildings, and through online postings on Craigslist.org. Participants were screened based on their mobile phone experience (e.g., that they could add a number to their phone's contact list) and out-of-house behavior (e.g. employed away from home, bar and restaurant visit frequency, etc.) and internet connectivity at home or work. Ages ranged from 22-56 (median 29). Two participants were full-time students; the others included a furniture designer, political consultant, bookseller, translator, grant manager, artist, etc. Six were in a serious relationship; one had children.

Each participant was supplied with an Audiovox SMT 5600 SmartPhone, wall-charger and car charger. The phones were preconfigured to run the Me Tool at startup. Other SmartPhone programs and menu items were removed to simplify interaction with the phone. Participants were asked to carry the phone with them at all times; the study phones were used only for experience sampling, and were not used, for example, to make calls or replace the participants' personal phones.

Participants were compensated based on their level of participation. Participants could complete up to 11 ESM surveys a day, at \$1 USD per survey, regardless of their movement or place visit activity. This scheme was established to promote survey completion without artificially motivating behavior. The incentive was contingent on the participant regularly logging into the website and filling out the web diary—at a minimum of 3-4 times per week. Participants were also remunerated for their two interviews, paper questionnaires, and travel time to and from our lab. Participation was strictly voluntary and prorated compensation would be made if a participant dropped out prematurely.

## 4 High Level Results

On average, participants carried a study mobile phone for 28 days each<sup>1</sup> (median 29). The length of the study varied slightly per participant because of scheduling. The average completion rate for the ESM surveys was 80.5%. A total of 4,295 ESM surveys were administered, 3,458 were completed at a rate of 216 surveys per participant (median 211). A total of 19,865 questions were answered. The survey completion time ranged from 20 seconds to five minutes (1.5 minute average). Despite early technical issues with the web diary server, 368 web diary sessions were completed, averaging 23 per participant.

---

<sup>1</sup> We lost two weeks of data for one participant due to a multiple drive failure in a RAID5 storage system.

1,981 individual place visits were logged via ESM surveys. Of these, 862 were to a public place at a rate of 1.9 public place visits per day. The public place visits ranged from the usual—the grocery store, local park, Starbucks—to the rather unusual, an outdoor sausage festival, a wedding chapel and the state fair. Table 1 shows a breakdown of the types of places visited, the number of participants who logged at least one visit to that place, and its mean preference rating. The “Other” category is an amalgamation of 21 place categories including: Gas Stations, Shopping Malls, Movie Theatres, and Parks.

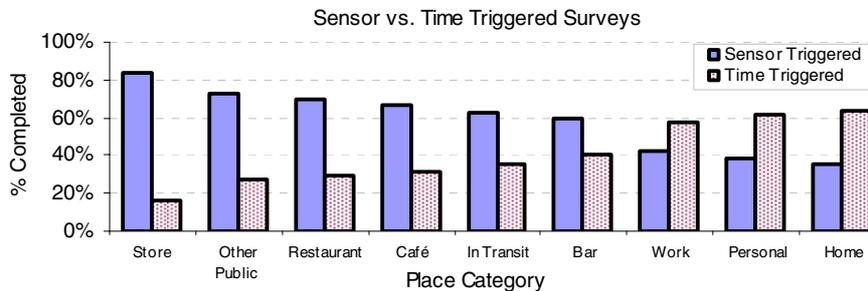
**Table 1.** Captured Places

Public and Private Places			
Place Category	# of Visits	# of Part.	Mean Place Rating (SD)
<b>PUBLIC</b>			
Bar	39	10	3.8 (0.76)
Café	122	15	3.8 (0.63)
Restaurant	251	16	3.8 (0.80)
Store	186	16	3.7 (0.91)
Other	266	16	3.6 (1.2)
<b>PRIVATE</b>			
Home	450	16	4.6 (0.85)
In Transit	354	16	2.9 (1.4)
Personal	253	16	4.4 (0.83)
Work	416	16	3.6 (0.99)

**4.1 ESM Effectiveness**

Given that we were interested in visit behaviors and travel routines, it was important to capture data about as many places as possible. We used context-triggered ESM to maximize the number of places surveyed and minimize participant disruption.

A majority of the 1,981 place visits were captured as a result of the mobility detection algorithm triggering a sensor-based survey. Fig. 2 reveals the effectiveness of the sensor-triggered ESM surveys in capturing public place visits versus private place visits. The graph is organized from left-to-right based on the percentage of sensor-versus time-triggered surveys completed at each place category. For example, over 80% of completed ESM surveys tagged as “at a store” were sensor-triggered vs. fewer than 40% for “at home.” The three right most place categories are all “private places” and received the greatest percentage of time-triggered surveys. This is to be expected if the mobility detection algorithm was behaving appropriately: each of those place categories (work, personal, and home) are characterized by long periods of sedentary activity. The ESM survey system would only invoke a time-triggered survey after failing to sense mobility for one hour. The algorithm was not perfect—it suffered from both false positives (i.e., detecting a phone was stationary when mobile) and false negatives (i.e., detecting a phone was mobile when stationary). However, it was quite successful in capturing ESM data as participants traveled from place to place.



**Fig. 2.** The sensor-triggered surveys captured a majority of public place visits

#### 4.2 The “My Places” Paper Questionnaire

The “My Places” questionnaire, administered at the beginning of the study, was an initial investigation into the complexities of visit behaviors. The purpose of this questionnaire was to establish a comparative point for the ESM data. It asked about travel routines to public places that are frequented two or more times a year. For each place the participant listed, the questionnaire asked for an explicit place rating, an approximate location, an approximate visit frequency per year, typical travel time, and their primary reasons for visiting that place. Place ratings were on a scale of 0.5 – to 5, where halves were allowed. A rating of 5 implied a strong liking for the place. The same scale was used in the ESM surveys.

The total number of places listed was 634—roughly 40 per participant. The average place rating was 3.8 (SD=0.79). While on the surface, the visit behaviors seemed routine (e.g. grocery shopping, eating out, going to the park) participants included stories that conveyed underlying complexities. One participant visited a particular coffee shop once a day “for the caffeine” despite conflicting feelings about the company’s economic and political policies: the “Coffee is OK but they are too corporate and they give to democrats.” Another participant patronized a restaurant 12 times a year, yet he rated it 2/5 stars because he “didn’t like the food”—“[I go there] because my friends like it.” Most often, proximity played a critical role—either because a service or item was only available in a certain area (“[this grocery store has] good selection—some foods I cannot get anywhere else”) or because of sheer convenience (“It’s nearby but the food is bad—it’s cheap and easy”).

### 5 Analysis

We present statistical analyses exploring the relationship between place visit behavior and a person’s explicit place rating. The focus is on exploring factors that could be automatically detected by emerging location technologies, for example, with assisted GPS or beacon based location [14]. We explore two implicit factors in detail: *visit frequency* and *travel effort*. We hypothesize that the number of visits a person makes to a place and the required travel time to get there reflects a corresponding interest. We also investigate social effects to determine whether place ratings differ for those visited alone versus with others.

In the analysis below, Likert-scale responses are often categorized in two or three nominal groups (e.g. a “disagree” group and an “agree” group) to partition data for significance tests—in these cases we will use non-parametric tests for frequency distributions and t-tests to compare equality of means. For correlative analysis we will be using Spearman’s correlation coefficient analysis (abbreviated  $\rho$ )<sup>2</sup>. Significance will be denoted by one star (\*) for  $P < 0.05$  and two stars (\*\*) for  $P < 0.01$ .

Place ratings were on a scale of 0.5 to 5, where halves were allowed, resulting in ten discrete rating points. To better understand the rating variable as an expression of preference, we occasionally asked one or two additional follow-up “rating” questions during the course of an ESM survey: “I really like this place” and “This place is

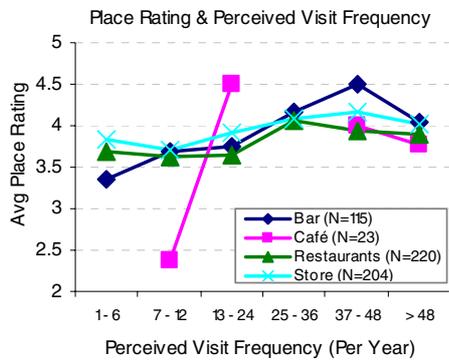
<sup>2</sup>  $\rho$  measures association between ordinal data without making assumptions about the frequency distribution of the variables.

important to me.” Each had answer choices on a 5-point Likert scale. As one would hope, both were positively correlated with rating. “I really like this place” was found to be highly correlated with the explicit place rating ( $\rho=0.68^{**}$ ). The second question, which inquired about importance rather than preference, resulted in a lower correlation, though still positive ( $\rho=0.47^{**}$ ).

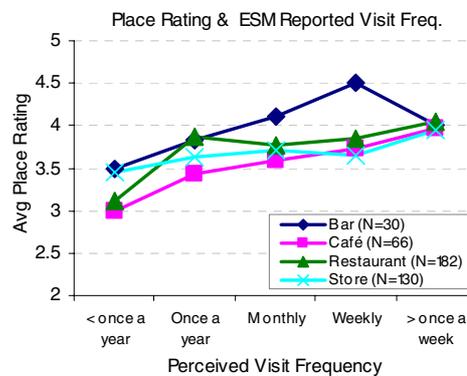
### 5.1 Visit Frequency vs. Explicit Rating

Our first hypothesis is that the number of visits a person has to a place is a strong indicator of their preference for that place. We will investigate this in two ways: first, by examining participant responses to a paper questionnaire about place routines and, second, by looking at the participants’ in situ, self-reported visit frequency to places as answered on ESM surveys.

**My Places Questionnaires.** For each place listed in the “My Places” paper questionnaire (see Section 0), the participant was asked to list their estimated visit frequency to that place. Using this data, we found a positive correlation between a person’s preference for a place and their respective visit frequency ( $\rho=0.20^{**}$ ). When broken down into subcategories as shown in Fig. 3, only bar ratings and visit frequency are significantly correlated ( $\rho=0.29^{**}$ ). The correlation is weaker than we expected due to the unequal distribution of highly rated places amongst the visit frequency categories. That is, there is a distribution of places that people really like but only visit a few times a year. We note this occurrence in our ESM data as well.



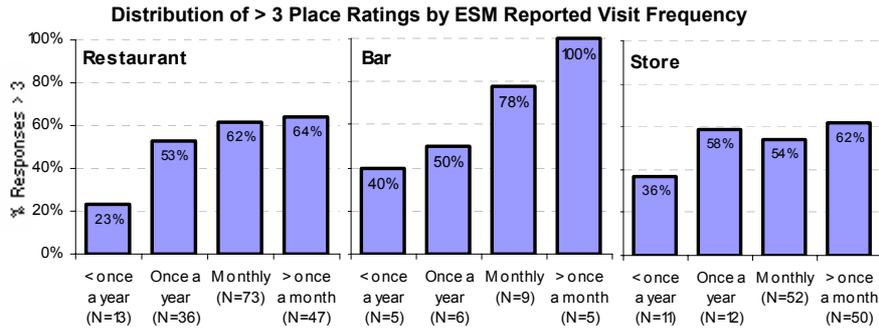
**Fig. 3.** “My Places” Questionnaire. Shows a correlation between bar ratings and visit frequency ( $\rho=0.29^{**}$ ). Other categories show similar trend but failed significance.



**Fig. 4.** ESM Data. Though all show a positive trend, only café and bar ratings were correlated with visit frequency ( $\rho=0.39^{*}$  &  $\rho=0.27^{*}$ ).

**Self-reported ESM Visit Frequency.** Each “public place” ESM survey asked “How often do you go to this place?” with six answer choices, which ranged from “This is my first time” to “More than once a week.” The purpose of this question was to collect visit frequency data in the event that the study was not long enough to capture a sufficient amount of *observed* repeat visits for statistical testing. We found a slight

positive correlation ( $\rho=0.14^{**}$ ) between ESM reported visit frequency<sup>3</sup> and explicit place ratings. This is similar to what we found with the paper questionnaire. When divided into the categories shown in Fig. 4, we found that visit frequency is a modest indicator of preference for bars ( $\rho=0.39^*$ ) and cafes ( $\rho=0.27^*$ ) but not for restaurants and stores. We will discuss possible reasons for this in the next section.



**Fig. 5.** Restaurant and bar visits show a clear monotonically increasing trend between visit frequency and preference. This trend is not evident in store visits.

To better understand the distribution of places that are liked, but only visited a few times a year, place rating data was split into two nominal groups: those places rated as less than or equal to 3 (“ $\leq 3$ ”) and those places rated as greater than 3 (“ $> 3$ ”). For restaurants, bars, and cafes the visit frequency distributions for the two categories were shown to be significantly different. For example, over 64% of the restaurants visited more than once a month were rated positively ( $> 3$ ) while only 23.1% of the restaurants visited less than once a year were rated positively. A similar monotonic upward trend occurred for bars but not for stores—see Fig. 5.

Finally, looking at the ESM question “I plan on returning to this place” we see additional evidence that repeated place visits implies preference. This question uses a 5-point Likert-scale from “Strongly Disagree” to “Strongly Agree.” We found a positive correlation between these responses and explicit place ratings ( $0.31^{**}$ ). When broken into sub-categories, the correlation for bars increases to  $\rho=0.91^*$ —the other place types do not change significantly from  $0.31^{**}$ . The key takeaway here is that planned returned visits seem to indicate preference for a place. These correlations, although a bit stronger, are consistent with the correlations between ESM reported visit frequency and place ratings explained above.

We conclude that visit frequency is a modest indicator of preference for bars and cafes but not for restaurants and stores. We will explore how combining visit frequency with other factors can boost these correlation coefficients later in Section 0 by reducing “noise” in the data.

<sup>3</sup> The correlations were run on five answer choices instead of six. As we were only interested in judging visit frequency, “This is my first time” was removed, leaving the ordinal scale: “less than once a year” to “more once a week” as shown in Fig. 4.

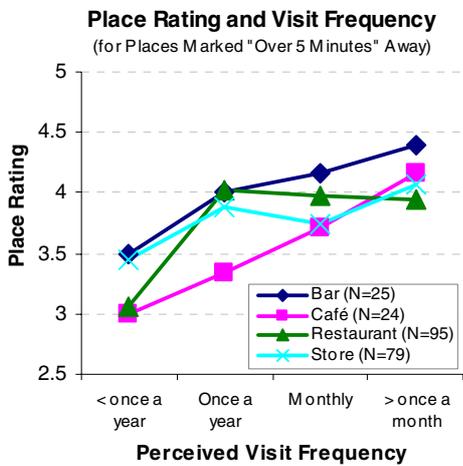


rated positively while 76.6% of the visits made over 15 minutes away are rated positively ( $P < 0.04$  for chi-square).

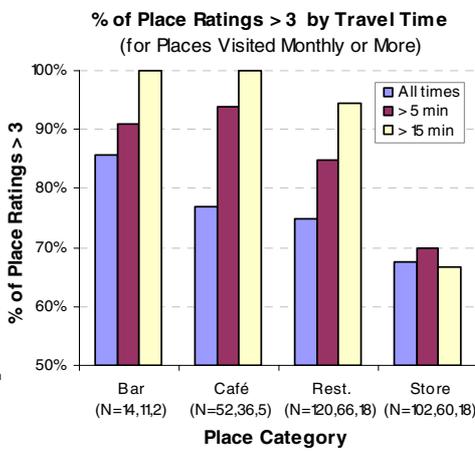
We conclude that travel time alone is a weak indicator of preference.

### 5.3 Combining Visit Frequency and Travel Effort

To examine whether visit frequency and travel time perform better together, we split all place visits into two nominal groups: places marked “0 – 5 minutes” and “over 5 minutes.” We would expect that filtering out visits to places within 0 – 5 minutes would reduce the amount of noise in the data by removing those trips highly motivated by convenience.



**Fig. 8.** ESM Data. A stronger upward trend is evident once removing “0-5 min” trips. Café ratings correlation w/visit frequency increases to  $\rho = 0.56^{**}$ .



**Fig. 9.** ESM Data. The % of bars, cafes, and restaurants rated >3 increases steadily with travel time

Running correlative analysis between place ratings and visit frequency with the “over 5 minutes” group increased the previously calculated correlation  $\rho = 0.14^{**}$  to  $\rho = 0.21^{**}$  (see Fig. 8). The correlation increased significantly for cafes to  $\rho = 0.56^{**}$ . Splitting at “0 – 15 minutes” and “over 15 minutes” instead results in even stronger correlations— $\rho = 0.37^{**}$  for the “public place” general case and  $\rho = 0.38^*$  for restaurants. In addition, for places visited monthly or more, the percentage of places that are liked (>3) goes up significantly for bars, cafes, and restaurants as their travel time increases—see Fig. 10.

As a result, we conclude that combining visit frequency and travel effort result in better indicators of preference than treating each factor separately.

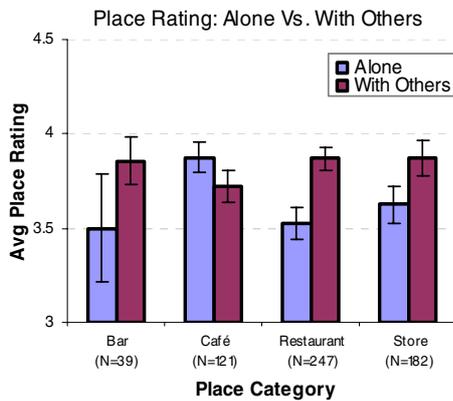
### 5.4 Exploring the Social Effect

The places we go are often affected by the presence of others. This is true both at the macro level as one must deal with commuter traffic, long lines at the supermarket, or

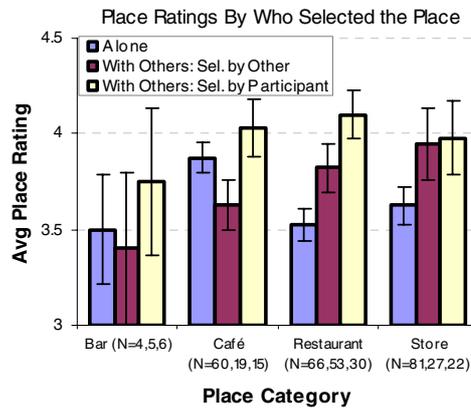
a crowded public square during lunch as well as at the micro level as we are obliged to run errands for others, occasionally defer our tastes to those of our partners, or forgo our typical preferences for a special social event. We were interested in studying the more micro social effects to try and determine their impact on place ratings.

Of our 16 participants, 10 were single and 6 were involved in steady romantic relationships. Most of our participants lived with others (Median=1 housemate, SD=1.14) though three of our participants lived alone. A majority of the 862 public place visits logged by our participants were with others (62.2%); however, different levels of social activity exist between place categories. Bars, for example, were the most “social” place—nearly 90% of those visits were with a group of one or more. In contrast, café visits were split nearly 50/50 as participants would often pick up coffee on their route to work in the morning or, conversely, with co-workers during afternoon break. If there is a significant difference in place ratings when a participant is alone versus when they are with others, this may make it more difficult to infer a person’s place preference based on their visit behavior as it is difficult to automatically sense when people are with others (though perhaps [15, 19] is a start).

Each ESM survey administered for a public place included the question “How many people are you with?” We used participant responses to this question to divide visits into two categories: those that occurred while alone and those that occurred while with others. We found no significant difference in mean ratings for public places as an aggregate variable when alone versus with others. When broken down by category, only restaurants had a significant difference. This was surprising; we expected a much larger disparity. However, the correlation between visit frequency and explicit rating did change slightly when a participant was with others versus when they were alone ( $\rho=0.14^{**}$  vs.  $\rho=0.22^{**}$ ). Similar results were found for travel time and place rating. In the general case, it does not appear that the presence of others serves as a strong confounding variable.



**Fig. 10.** ESM Data. Bars, restaurants and stores visited with others were rated higher than if alone. However, only restaurants pass significance tests.



**Fig. 11.** ESM Data. Exploring how the selection of place changes ratings.

To explore social effects in more detail, we asked follow-up 5-point Likert-scale questions after the participant indicated that they were with others. We explore two of these follow-up questions here: “I would not have gone here if it wasn't for the group” and “Someone in the group besides me selected this place.” The responses were broken up into three categories: “Disagree,” “Undecided,” and “Agree.” Average place ratings were then compared for the “Disagree” and “Agree” groups using t-tests for equality of means. For the first question, “I would not have gone here if it wasn't for the group,” participants agreed 51.6% of the time. We would expect a place rating to be higher when the participant disagreed—meaning that they would go to that place regardless of the group. Our results concur with these expectations. Places were rated higher if the participant disagreed, but only slightly (4.0 versus 3.7,  $SD=0.9$ ,  $0.9^{**}$  t-test). For the second question, “Someone in the group besides me selected this place,” participants felt as though they selected the place a minority of the time (46.4%). We would expect a higher place rating when the participant themselves selected the place. A statistically significant increase was found, from 3.8 to 4.1 ( $SD=0.9$ ,  $0.8^*$  t-test)—see Fig. 11.

We conclude that the presence of others in itself does not significantly change the explicit place ratings. It is the decision process that matters—that is, who selected the place, but even then the difference in rating is minimal.

## 6 Discussion

Based on our findings, we believe that creating a place-based preference inferencer is possible but not straightforward. Visit frequency and travel time treated separately were positively correlated with place ratings in our data. However, the magnitude of these correlations were far below our expectations ( $\rho=0.14^{**}$  and  $\rho=0.11^{**}$  respectively). Pairing these implicit factors together lead to better results. For example, places visited more than once a month and over five minutes away were rated significantly higher on average than places in general. This suggests that combinations of implicit factors will likely be the best indicator of preference. However, it's still unclear whether these correlation coefficients are strong enough to generalize a rule. A real application could also use the hybrid approach by combining both implicit and explicit ratings to correct and augment the inferred preferences.

We were surprised that, in general, places were not rated significantly differently when alone versus with others. However, the trend does suggest that when with others, a place is rated higher—particularly for restaurants and bars. We believe this is due to the strongly associated social component of those places (e.g., visiting a bar is a social activity). The question of alone versus others was found to be less significant than the question of who actually selected the place. When in a group, participants tended to rate a place higher if they themselves selected it rather than someone else in the group. Automatically detecting such nuances is probably not realistic. An opportunity for future research is whether such an effect is detrimental to actually determining preference.

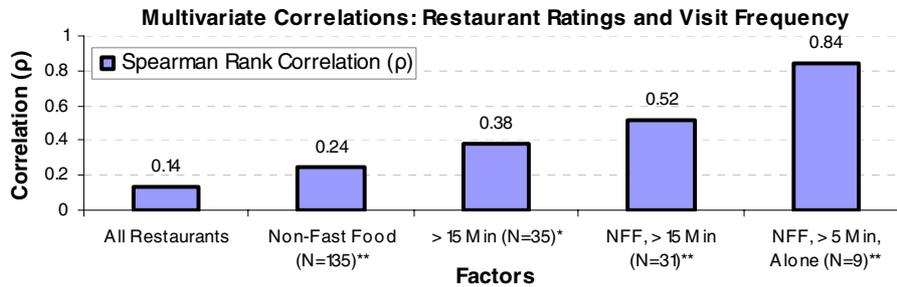
We believe that a real system will likely need to control for noise in the data. The following describes a brief case study of restaurants to demonstrate how simple data filters can be applied to increase correlations.

### 6.1 A Case Study: Restaurants

This section offers insight into how multiple factors can be combined to increase the correlation between visit frequency and explicit place ratings. We believe that designers of location-aware applications could make certain adjustments for each place type based on expectations of travel behavior and a few basic intuitions. We will explore restaurants as an example.

As noted in Section 0, visit frequency and restaurant ratings failed the significance test for rank correlation. Given our suspicion that convenience would be a confounding variable in detecting place preference relationships, our first insight was to split restaurants into two groups, fast food and non-fast food. As a result of this split, the non-fast food group resulted in a slight, but statistically significant correlation between visit frequency and rating ( $\rho=0.24^{**}$ ). If we split the non-fast food group further into two groups based on travel time “0 – 15 minutes” and “over 15 minutes” the correlation coefficient increases to  $\rho=0.52$  for the “over 15 minutes” group. So, simply by filtering restaurant visit behavior by fast food and distance, our correlation coefficient increases from  $\rho=0.14$  (not significant) to  $\rho=0.52^{**}$ . Similar steps could be taken with other place types to improve the correlations between visit behaviors and explicit place rating.

If we divide our data further along social lines, we see a very high correlation ( $\rho=0.84^{**}$ ) between visit frequency and place rating when our participants were alone, the restaurant was non-fast food and over 5 minutes away—see Fig. 12. This last analysis should be taken with caution due to the relatively small sample size.



**Fig. 12.** Using intuitive splits in the data, correlation coefficients can be increased for specific place types. Here, splitting data across fast food and travel time result in an increase correlation between visit frequency and place rating for restaurants. Bars with stars are statistically significant. NFF=non-fast food.

## 7 Related Work

The study of human spatial and temporal behavior in the environment is a subset of human geography [9], which intersects with research in tourism, urban and transportation planning, and the study of travel behavior. Low-level research of individual travel patterns has historically been rare in these fields as past techniques such as paper diaries or direct observation were costly and required high personnel resources

[17, 18]. Quite recently, however, advances in location-sensing technology have begun to dramatically change the methodology employed in these areas [17, 18]. Schönfelder, for example, recently looked at the relationship between routines and variety seeking with respect to the characteristics of location choice in daily travel by studying GPS data streams. He found that location choice is strongly routinized (e.g. the top four “leisure” locations received 40-50% of all visits). Our study also found a high level of repeat visits (e.g., 40% of restaurant visits and 80% of café visits). Other researchers in this area have studied the spatial distribution of places that individuals come into contact with in their daily living [10]. Called “activity spaces,” their size and structure depends on three factors: the individual’s location of home (and duration of residence), their regular activities (e.g. work, school, working out at the gym) and the travel between and around these “place anchors.” This has practical implications for customizing algorithms per individual based on their activity space—e.g. variety of places visited, typical travel time, etc.

The use of implicit features or “implicit interest indicators” [6] to infer preference is also an active research area, though to our knowledge we are the first to actively study it with respect to real world behaviors like visit patterns. One of the original systems, called GroupLens [13], studied the correlation between time spent reading an article and explicit user ratings. Mainstream commercial systems such as Last.fm, Amazon.com, and TiVo have successfully employed implicit interest indicators to make recommendations, improve the user experience through personalization and build stronger online community. Moving from the virtual to the real world, Chen [5] proposes a context-aware collaborative filtering system for the ubicomp environment to recommend activities based on what others have done in similar contexts (e.g., based on location, weather, group proximity, etc.) but did not investigate how well those context features could perform.

Brown et al. [3] describe a mobile system for sharing data amongst tourists. A collaborative filtering algorithm was used to recommend photos, web pages, and places to visit based on historical data (including GPS traces). Their focus was not on investigating the relationship between place preferences and visit behaviors but rather on exploring how a visit to a city could be shared across the Internet and, crucially, how physical and online visitors to an area could interact (e.g., an online visitor could “piggyback” on the experiences of a physical visitor). In this way, they did not evaluate the effectiveness of their recommender algorithm nor provide details on its function. However, they do show how the physical and virtual space can rather seamlessly converge in a location-aware mobile application as well as the potential of mobile recommenders for filtering and suggesting content.

Others have explored the detectability of place [11, 14, 22]. Zhou et al., for example, looked at the relationship between place discovery and importance and found that those places judged meaningful by the subject were much easier to detect. In our work, we found that 43.7% of all logged place visits were to home and work—places reasonably inferred to be “meaningful” by our participants. Given the large number of these visits compared with other locations, we would expect place discovery algorithms to do much better discovering them (particularly when tied to temporal patterns, e.g., work during day, home during night). Participants in our study were also asked about the meaningfulness of places via ESM surveys. We found evidence that an explicit judgment of preference for a place is correlated positively with an explicit

judgment of meaningfulness ( $\rho=0.47^{**}$ ). Future place discovery systems may want to investigate both.

## 8 Future Work

This paper presents the first investigation of implicit interest indicators derived from visit behaviors and explicit place ratings gathered in the field. As such, there are many opportunities for future work. First, this work only considered two implicit features of place visit behavior, visit frequency and travel time. We believe that other factors such as dwell time (e.g., how long one spends in a place), temporal patterns (e.g., weekday vs. weekend, lunch vs. dinner, season) and mode of transportation may also contribute to inferring preference. In addition, we did not look at negative interest indicators—those features that would correlate negatively with rating. For example, a restaurant that is across the street from work but never visited might indicate negative interest. As our GSM-based sensors did not provide us with high-resolution location data, an interesting follow-up study could correlate actual location data streams (e.g., assisted GPS) with explicit place ratings. A longitudinal study (i.e., 6-12 months) could be used to further investigate visit behaviors as well as to look at long-term temporal patterns and their relationship to preference. Future work could also explore the potential of generalizing place preferences to general interests. For example, frequent visits to a snowboarding shop may indicate a general interest in downhill snow sports.

We are currently in the process of collaborating with the Department of Statistics at the University of Washington regarding a more sophisticated analysis of this data. A linear mixed effects model has been created to take into account the variation between subjects and within subjects. Preliminary results from this analysis are in accordance with the findings above. For example, a slight, but statistically significant correlation was found between visit frequency and preference. When visit frequency and travel time were combined, the positive correlation strengthens.

## 9 Conclusion

This paper examines to what degree do automatically detectable visit behaviors indicate preference for a place. We explored two implicit factors in particular, visit frequency and travel time. We found that both features have a slight, but statistically significant positive correlation with explicit place ratings. When combined, however, they become better measures of preference. In general, we found that splitting public place visits into sub-categories based on place type resulted in significantly higher correlations, particularly for bars, cafes, and restaurants. This finding implies that studying travel routines at an aggregate level split simply between “private” and “public” is only somewhat effective at determining preference—there may be too many differing motivations for visiting a place at this resolution. Further, we found that the presence of others itself is not a confounding variable. Finally, our four-week study is the first to explore the use of context-triggered ESM on mobile phones. We believe that the Me Tool is a promising technology for studying human behavior on mobile platforms as well as for validating ubicomp technology in the field.

## References

1. Ali, K. and van Stam, W. (2004), TiVo: Making Show Recommendations Using a Distributed Collaborative Filtering Architecture. In *Proc of KDD'04*. Aug 22-25. Seattle, WA.
2. Barrett, L.F. and Barrett, D.J. (2001), An Introduction to Computerized Experience Sampling in Psychology. *Social Science Computer Review.*, V. 19, No. 2, S01, pp. 175-185.
3. Brown B., M. Chalmers, M. Bell, I. et al. (2005) Sharing the square: collaborative leisure in the city streets. Proceedings of ECSCW 2005, Paris, France. pp. 427-429.
4. Carter, S., Mankoff, J. (2005). When Participants do the Capturing: The Role of Media in Diary Studies. In *Proc of CHI'05*. April 2- 7, Portland, OR.
5. Chen, A. (2005), Context-Aware Collaborative Filtering System: Predicting the User's Preferences in the Ubiquitous Computing Environment. In *Proc of CHI'05*.
6. Claypool, M., Phong, L., Waseda, M., Brown, D. (2001), Implicit Interest Indicators. In *Proceedings of Intelligent User Interfaces '01*, Jan 14-17 '01, Santa Fe, NM, pp 33-40.
7. Consolvo, S., Walker, M. (2003). Using the Experience Sampling Method to Evaluate Ubicomp Applications. *IEEE Pervasive: The Human Experience*, pp. 24-31.
8. Consolvo, S., Smith, I., Matthews, T., et al. (2005). Location Disclosure to Social Relations: Why, When & What People Want to Share. In *Proc of CHI'05*. Portland, OR.
9. Fellman, J.D., Getis, A & Getis, J (1999) Human Geography, WCB/McGraw-Hill, Boston.
10. Golledge, R.G. & Stimson, R.J. (1997) Spatial Behavior, Guilford Press, New York.
11. Hightower, J., Consolvo, S., LaMarca, A., Smith, I., and Hughes, J. (2005) Learning and Recognizing the Places We Go. In *Proc. UbiComp 2005*. Tokyo, Japan.
12. Intille, S. S., Rondoni, J., Kukla, C., Anaconda, I., and Bao, L. (2003), A context-aware experience sampling tool. In *Proc of CHI '03* NY, NY: ACM Press, 2003, pp. 972-973.
13. Konstan, J., Miller, B., Maltz, D., Herlocker, J., et al. (1997), GroupLens: Applying Collaborative Filtering to Usenet News. *Communications of the ACM*, 40(3):77-87, 1997.
14. LaMarca A., Chawathe Y., Consolvo S., et al. (2005). Place Lab: Device Positioning Using Radio Beacons in the Wild. In *Proc Pervasive'05*, Munich, Germany.
15. Paulos, E., Goodman, E. (2004) The Familiar Stranger: Anxiety, Comfort, and Play in Public Places. In *Proc of CHI'04*. April 24-29, Vienna, Austria
16. Smith, I., Chen, M., Varshavsky, A., Sohn, T., and Tang, K. (2005), Algorithms for Detecting Motion of a GSM Mobile Phone. In *ECSCW 2005*, Workshop Paper.
17. Schönfelder, S. (2003), Between routines and variety seeking: The characteristics of locational choice in daily travel. In *10th Intl Conf on Travel Behaviour*, Lucerne, Aug. '03.
18. Shoval, N., Isaacson, M. (2006). Tracking Tourists in the Digital Age. In *The Annals of Tourism Research*.
19. Terry, M., Mynatt, E. D., et al. (2002) Social Net: Using Patterns of Physical Proximity Over Time to Infer Shared Interests. In *Proc. CHI 2002*, Short Paper. pp., 816-817.
20. Wheeler, L., & Reis, H. (1991), "Self-recording of Everyday Life Events: Origins, Types, and Uses," *Journal of Personality*, 59, pp. 339-354.
21. Zhang, D. and Adipat, B (2005). Challenges, methodologies, and issues in the usability testing of mobile applications. *International Journal of HCI*. v18.3, pp. 293-308.
22. Zhou, C., Ludford, P., Frankowski, D., and Terveen, L. (2005) An Experiment in Discovering Personally Meaningful Places from Location Data. In *Proc. CHI 2005*.