

Activity-Based Serendipitous Recommendations with the Magitti Mobile Leisure Guide

Victoria Bellotti, Bo Begole, Ed H. Chi, Nicolas Ducheneaut, Ji Fang, Ellen Isaacs, Tracy King, Mark W. Newman, Kurt Partridge, Bob Price, Paul Rasmussen, Michael Roberts, Diane J. Schiano, Alan Walendowski
Palo Alto Research Center
3333 Coyote Hill Road, Palo Alto, CA
magitti@parc.com

ABSTRACT

This paper presents a context-aware mobile recommender system, codenamed *Magitti*. *Magitti* is unique in that it infers user activity from context and patterns of user behavior and, without its user having to issue a query, automatically generates recommendations for content matching. Extensive field studies of leisure time practices in an urban setting (Tokyo) motivated the idea, shaped the details of its design and provided data describing typical behavior patterns. The paper describes the fieldwork, user interface, system components and functionality, and an evaluation of the *Magitti* prototype.

Author Keywords

Field studies, user experience design, interaction, context-aware computing, mobile recommendation systems, leisure.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI); Miscellaneous.

INTRODUCTION

Finding out what's going on in a city has traditionally been supported by city guides such as "Time Out" in London and New York, and "Tokyo Walker" in Tokyo. These publications are gradually being displaced by electronic media; most recently by location-based services [29], which allow users to search for local restaurants, movies, stores and so on by entering a query. An alternative user experience, perhaps more in keeping with leisure, is *serendipitous, activity-based discovery* of activities and venues, made possible by advances in context-aware computing and machine learning.

This paper introduces a system code-named *Magitti*; an

activity-centered mobile leisure-time guide. *Magitti* delivers timely and personally relevant recommendations about nearby venues for pursuing activities in an urban environment. Unlike other guides, it is not designed for tourists, but for young urbanites interested in all kinds of leisure activities, emphasizing spontaneity rather than sightseeing. *Magitti* is also unique in using machine learning techniques to make a chain of predictions that target information to user interests. It first predicts ongoing and future activity based on the user's context and models of past behavior, and then predicts what information will be most useful within the predicted activity based on user preferences.

In this paper, we first review extensive field work that informed the design of *Magitti*. We then introduce *Magitti* along with related work in context-aware and guide systems. We provide an overview of the system and its user interface, and present findings from a user evaluation.

CONCEPTUALIZING MAGITTI

The *Magitti* project was a sponsored research engagement with Dai Nippon Printing Co., Ltd. (DNP), one of Japan's largest printing companies. DNP wishes to develop a service to replace printed city guides. To address that wish, we created a novel system that *detects current user context, infers likely current and future leisure activities and recommends content about suitable venues* (e.g., stores, restaurants, parks, movies.). The target demographic is 19-25-year-olds who go out more than most and quickly adopt new technologies. The target locations are Japanese cities that have so many venues that few people are familiar with more than a tiny fraction of them.

Designing such a system necessitated learning about the practices, problems, and needs of the prospective user category, urban Japanese young adults, and how they like to spend their leisure time. To this end, PARC researchers spent three two-week periods in Japan intensively gathering data, described next.

Understanding Leisure Time Priorities

There is a dearth of literature in English on the specifics of Japanese leisure time activities (but see [12, 24] for

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 2008, April 5-10, 2008, Florence, Italy.

Copyright 2008 ACM 978-1-60558-011-1/08/04...\$5.00.

examples of related work). A recent time-based survey [17] unfortunately classifies activities into categories that are too coarse for making specific recommendations (e.g., “Rest & Relaxation,” “Social Life”). We needed data about leisure practices, particularly out-of-home activities, which were fine-grained enough to make specific activity recommendations. So we conducted six types of field exercises focusing on the following questions:

- How do young Japanese spend their leisure time?
- What resources do they use to support leisure time?
- What needs exist for additional support that might be provided by a new kind of media technology?

Interviews and Mockups (IM): Twenty semi-structured interviews with 16-33-year-olds and a further 12 interviews with 19-25-year-olds examined routines, leisure activities, and resources used to support them. We first asked for accounts of recent outings and then for feedback on Magitti concept scenarios and a mock-up.

Online Survey: We conducted a survey on a market research web site to get statistical information on specific issues. We received 699 responses from 19-25-year-olds.

Focus Groups: We ran three focus groups of 6-10 participants each, concentrating on mobile phone use. In these we presented a walkthrough of the Magitti mock-up and its functions to gather detailed feedback on the concept.

Mobile Phone Diaries (MPD): To get a picture of the daily activities of 19-25-year-olds, we conducted two mobile phone diary studies, first with 12 people for one Sunday, and then with 21 participants for a seven-day week.

Street Activity Sampling (SAS): We conducted 367 short interviews with people who appeared to be in our target age range and at leisure in about 30 locations in Tokyo and surrounding areas at different times and days of the week. We asked people to report on three activities from their day, choose one as a *focal activity*, classify it into one of a number of pre-determined types and characterize it in terms of planning, transportation, companionship, information requirements, familiarity with the location, and so on.

Expert Interviews: We interviewed three experts on the youth market in the publishing industry to learn about youth trends in leisure, and information commonly published to inform and support their activities.

Informal observation: Finally, we “hung out” in popular Tokyo neighborhoods observing young adults at leisure.

Critical Findings from Field Exercises

The following findings were key to Magitti’s design.

How young people in Tokyo spend their leisure time

In IM interviews, shopping was the most commonly reported activity that requires an outing, followed by going out with friends, dining out, going on a date, and doing sport (some of these objectives could overlap). In SAS interviews, dining was the most frequent type (31.8%) of

activity, closely followed by shopping (24.6%), and then by browse/explore/look (7.5%). So dining and shopping appear to dominate activities that involve going out (cf [34]).

SAS interviewees reported going out on average 2-3 times a week. Average commutes to leisure took 20 to 30 minutes, but it was not unusual to commute for an hour or more.

What resources are used to support leisure time

In pre-IM questionnaires and the online survey, respondents reported using roughly the same top four resources to discover and plan leisure time: Friends and family, TV, Internet, and Magazines. While online survey respondents rated the Internet highest and the IM interviewees rated friends and family highest, both groups rated city-guide magazines as less important. This suggests that printed resources are not holding their ground in the electronic age.

We learned that information based on personal experiences of friends and family is trusted the most. “*I know that other people I know have tried out something before, that it becomes like trustworthy information and something that I can rely on.*” Because such experiences are not very extensive, people also rely on less trustworthy, commercial sources such as magazine articles, web sites, and advertisements.

At the time of the interviews (2005), the mobile Internet was only rarely mentioned as a resource for planning and engaging in leisure. A year later, when we polled people on the streets, mobile Internet use seemed to have increased, with 60% reporting using the mobile web, mainly for train schedules (45%) and weather (24%). There were still relatively few reported uses of the mobile web for restaurant, shops, and movie information (around 5% each).

What needs exist for additional support

When asked, 58.8% of SAS interviewees said they would have liked more information to support their focal activity. The most common requests were for maps and venue locations (14.6%), customers’ and friends’ opinions (8.2%), prices (7.8%) and store/venue contents (6.8%).

During IMs and focus groups we learned that young adults often go to places they already know or that their friends know because they are safe bets. However, they also like to explore places they don’t know: “*... Me and my friends are always looking to find fun places to go to ... Some places that will suit our style.*” Finding new places is not always easy, though, as SAS interviewees were generally unfamiliar with the neighborhood they were in, averaging just 2.7 on a scale of 1=not at all to 7=extremely well.

When looking for new places, people mentioned concern about crowds, since places in Tokyo often become packed (especially when it rains), and since people sometimes travel long distances, they do not want to be disappointed. Requests for information about interiors, ambience, and photos reflected the fact that many venues in Tokyo have no front window where one can peer in to get a sense of the

atmosphere. Indeed, in the city center where land is expensive, venues are often up or down many floor levels.

Design Requirements

The field work led to several core requirements for Magitti.

Relaxation, Serendipity and Spontaneity: One of two top leisure priorities for 19-25-year-olds was “Relaxation” (the other being “Companionship”). Young people have busy schedules, often with multiple occupations (e.g., student and part-time-worker). Our focus groups were attracted to the idea of very easy and serendipitous information access, even though planning and research do take place around leisure.

Avoidance of Information Overload: Related to the preceding requirement is the issue of information overload, with many publications and advertisements, all attempting to persuade. This problem is particularly acute in Tokyo’s physical environment. As one respondent said, *“it’s almost like the information is flowing over and there’s so much information that you’re crushed with information.”* This suggests a design imperative to reduce leisure information to only the most relevant.

Minimal size: Mobile phones, or “keitai” (handhelds), are fundamental to leisure in Japan. It is essential for many facets of life, particularly amongst the younger generations to have something as small as possible that can be carried in a pocket [16]. This requirement trumps desire for a large screen that can display lots of useful information.

One-handed operation: Anyone who has ridden a Tokyo train or subway during commute hours will be aware of passengers using their keitai, many of them standing up holding onto a rail or pole for support. People also often use their keitai while walking around with a bag in one hand. So it is not surprising that interviewees repeatedly voiced a strong requirement for one-handed operation.

In our first iteration, we chose to focus on relaxation, serendipity, and spontaneity by generating recommendations automatically using activity inferencing. But we also recognize that opinions from family and friends could be extremely valuable in the recommendation process and hope to capture this requirement in future design iterations.

Having laid out motivations for Magitti, we now discuss the system itself in the context of our own and related research.

Magitti

Magitti anticipates the coming age of augmented reality and GeoTagging services that will present city-guide-style information to users based on location or by sensing physical objects [5, 13]. For example GeoVector™ (geovector.com) provides a service that allows a user to point her mobile phone at items in the environment to get information about them. However, as more venues and services go online in this way, the augmented space will

become as cluttered as the physical environment; Tokyo streets bristle with a multitude of garish signs and Jumbotrons. One approach to culling such information is demonstrated by Socialight™ (socialight.com), a service that uses social networks as a means of sharing and filtering location-based “sticky-notes,” so the user sees only friends’ recommendations, although they can opt to see information shared with everyone.

Magitti takes a different approach, which is to use context filtering to narrow down the inevitable overload of leisure time offerings in dense urban areas. It can do so without the user having to explicitly define her profile or preferences. The system infers interests and activities from models that are learned over time implicitly, based on individual and aggregate user behavior, such as places visited, web browsing, and communications with friends. Of course, the user can also explicitly provide information to improve Magitti’s recommendations, but our objective is that this should not be necessary to benefit from the system’s ability to filter information.

Magitti is a leisure activity guide with three key features:

- **Context Awareness:** It knows about current time, location, weather, store hours, and user patterns. It also lets people specify a future context for planning.
- **Activity Awareness:** It filters items to recommend based on its user’s inferred or explicitly specified activity modes. Five modes were derived from observations in our field work: Eating, Shopping, Seeing, Doing, or Reading. Each item in the Magitti database is explicitly tagged as being associated with and therefore a possible candidate recommendation for one of these modes.
- **Serendipitous, relaxing experience:** Users need not enter profile, preferences, or queries. They can rely on context and activity inferencing for Magitti to continually and automatically update recommendations.

RELATED RESEARCH: COMPARATIVE OVERVIEW

Context aware information retrieval research has been underway for well over a decade. An early example of a context-aware guide system is the Personal Shopper Assistant [4]; an indoor RF infrastructure that communicated with a wireless handheld device to provide useful context-related information about offerings and their locations in a store. The information was location- and time-sensitive, and in this way resembled Magitti, but in 1994 much of the research was focused on infrastructure.

Related Activity-Detection Research

One early system related to Magitti was Lamming and Newman’s activity-based information-retrieval system [20], which presented information that was generated in contexts similar to the user’s current context. Although the intent was for activity detection to be a key element in system, the technology of the time was not able to infer activity with effective accuracy.

Activity detection is an active research topic with many promising results. Begole et al. [7] experimented with sensor-based availability detection. Various authors used sensing to infer human activities from use of objects with RFID tags on them [27, 37], from special purpose wearable sensing hardware [21, 30], or by using video and audio data analysis [25]. Froehlich et al. [10] found correlations between place preference and data from sensed locations, such as frequency of visits and distance traveled.

In our work we are interested in enabling mobile context and activity inferencing with no special infrastructure or hardware. Relying only on time and location makes the challenge of activity inferencing more daunting, but some research in this space exists. Closest to our own activity modeling research is that of Liao et al. [22], who use location-based sensing with Relational Markov Networks and other techniques to infer occurrence of certain activities: ‘AtHome’, ‘AtWork’, ‘Shopping’, ‘DiningOut’, and ‘Visiting.’ Shopping and DiningOut, two of the activity types we are interested in, were the most difficult to classify.

The previously described work relates to the *detection* of a person’s current activity, but the Magitti guide system needs to predict a person’s future activities. Begole et al. [6] modeled activity on the keyboard and mouse across multiple locations to predict a user’s future reachability for communication. Ashbrook & Starner [3] and Krumm & Horvitz [18] demonstrated the detection and prediction of movement to significant locations from traces of GPS data. Prior to Magitti, there have been no user applications that predict a person’s likelihood of being engaged in different types of leisure activity.

Related work on Mobile City Guide Applications

There is one class of location-based information recommendation system that is similar in spirit to Magitti: location-aware tourist guides. We survey just a few notable systems in this section (see [5, 36] for more extensive surveys). These systems are similar to Magitti in that they recommend venues based on the user’s stated or inferred preferences. However, none of the following systems base recommendations on prediction of the user’s activities.

Cyberguide [1] was a mobile tourist guide for the Georgia Tech campus (inside and out) and adjacent neighborhoods. It was aware of its time, location, and history and could match information on venues and special events to these data. However, it was not a true recommender system.

MobyRec [33] is a context-aware mobile tourist recommender system that lets users specify preferences for hotels, restaurants, etc. MobyRec improves its recommendations over time. However, unlike with Magitti, new users always have to specify queries.

GUIDE [9] offers tour routes and accesses ticket reservation services and related capabilities based on user-stated interests and queries. In its tours, *GUIDE* takes into account the hours and busy times of requested attractions

and the most scenic routes between them, and dynamically recomputes routes based on location and time. *GUIDE* is targeted for touring unfamiliar areas, whereas *Magitti* is optimized for city residents who are looking for new places to enjoy leisure activities. *GUIDE* is also not geared towards a serendipitous experience as it does not predict user interests; users must issue queries or create a tour for themselves.

COMPASS [38] is a tourist guide service covering a wide range of venue types. It uses profile and goal information entered by its user to constrain recommendations of venues. It takes account of location, speed, user profile, schedule, shopping list, and recency of previous visit. Non-user specific elements of context are also considered, such as weather and traffic conditions. Like *Magitti*, it allows the user to view items on a map interface, but *COMPASS* filters them by the user’s stated goal and preferences, rather than activity type. And again, *COMPASS* relies on a user to enter her initial interests.

CRUMPET [28] provides tips, tour suggestions, maps and other information on a range of tourist-related venues (restaurants, movies, shows, etc.). It aims to learn user preferences over time, but not to predict future activity.

MAGITTI SYSTEM DESIGN

We now present the *Magitti* system in more detail.

User Interface

Magitti’s Main Screen (Figure 1, left) shows a scrollable list of up to 20 recommended items that match the user’s current situation and profile. As the user walks around, the list updates automatically to show items relevant to new locations. Each recommendation is presented in summary form on the Main Screen, but users can tap each one to view its Detail Screen (Figure 1, right). This screen shows the initial texts of a description, a formal review, and user comments, and the user can view the full text of each component on separate screens. The Detail Screen also allows the user to rate the item on a 5-star scale.

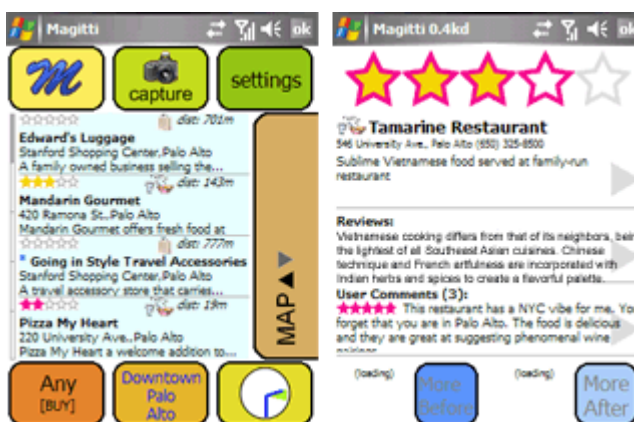


Figure 1. *Magitti*’s Main Screen (left) and Detail Screen (right).

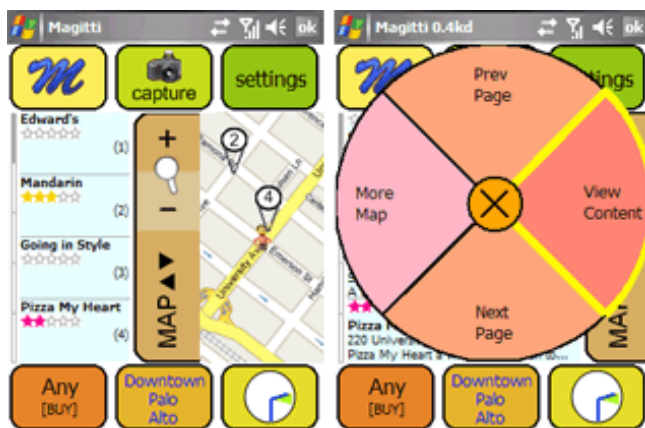


Figure 2: Main Screen with partial map view (left) and Main Screen marking menu (right).

To locate recommended items on the Main Screen, users can pull out the Map tab to see a partial map (Figure 2, left), which shows the four items currently visible in the list. A second tap slides the map out to full screen.

The minimal size and one-handed operation requirements have a clear impact on the UI. As can be seen from Figures 1 and 2, large buttons dominate the screen to enable the user to operate Magitti with a thumb [26] while holding the device in one hand. Our design utilizes marking menus [19] on touch screens to operate the interface, as shown in the right side of Figure 2. The user taps on an item and holds for 400ms to view the menu; then drags her thumb from the center X and releases over the menu item. As the user learns commands and their gestures, she can simply sweep her thumb in that direction without waiting for the menu to appear. Over time, she learns to operate the device without the menus, although they are available whenever needed.

Menu buttons at the bottom of the Main Screen allow the user to adjust the recommendation list if needed. By default, the system is in “Any” mode, meaning it will offer recommendations based on its predictions about the likelihood of each of the five modes of user activity; Eat, Buy, See, Do, or Read. But users can ask to see recommendations from just one category (and in so doing, informing the device of their activity interests, thus providing data for learning). Users can also ask for recommendations within a certain distance or time range. For planning purposes, they can ask for items in another location and/or at another time. Additionally, users can indicate their general preferences for cuisines, shops, activities, and events, which influence the types of recommendations offered. Users can also specify preferences for attributes of places, such as price range, noise level, availability of parking, smoking, and so on. And, finally, they can bookmark recommended items and perform keyword searches.

System Architecture

Figure 3 is a basic representation of Magitti’s overall client-server architecture. The mobile client UI runs on a handheld

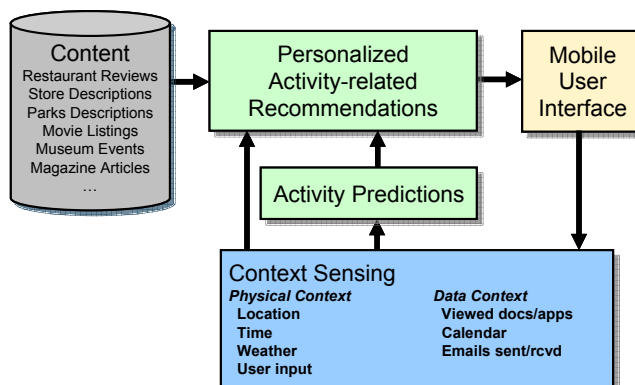


Figure 3. Magitti architecture.

device¹ that provides data for the *Context Sensing Module*. This module gathers data about the user’s *physical context* (GPS, time of day, user inputs, weather) and *data context* (content of emails sent/received, calendar, web pages and documents viewed, applications used). The text in the users’ data context is analyzed for information about their tastes and preferences as well as current and planned activities. The context data are sent to the Activity Prediction module and to the Recommender module.

A defining characteristic of Magitti is the chaining of predictions to filter information to the user’s likely current interests. Magitti first *predicts* the probability of the user being engaged in any of five activity modes during a specified period of time. The default is now plus two hours. The probability distribution determines the number of slots to use for activity-related information in the interface, as illustrated in Figure 4. Second, Magitti *predicts* the user’s interest in particular pieces of content for each activity mode using a combination of recommendation models. The two prediction modules are described next.

Activity Prediction Module

Activity probability prediction is based on a combination of patterns observed across the user’s demographic population and individual behavior pattern. The population patterns were derived from the data collected on Magitti’s target demographic in our fieldwork and from a Japanese Survey on Time Use and Leisure Activities [17]. Individual user behavior models are learned over time by associating each venue in Magitti’s database with one of the five activity modes that were observed in our fieldwork, and modeling the frequency of each mode by tracking user behavior. For example, if a user visits a retail store, the system records that as being in the Buy activity mode; similarly for visiting a restaurant or café, (Eat), theater or museum (See), gym or park (Do), and reading of content on Magitti itself (Read). These recorded location visits create models of the user’s individual activity behavior and preferences (e.g., cuisine

¹ We used a Mio A701 model with touch-screen, GPS, WiFi, running Windows Mobile OS.

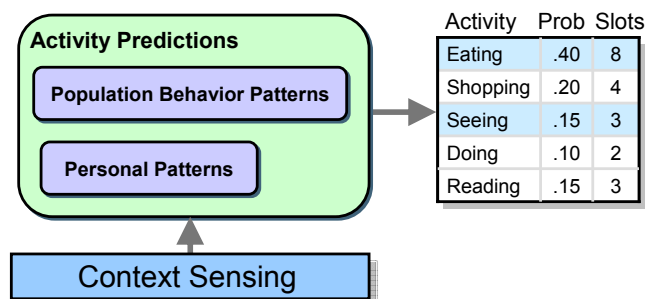


Figure 4. The Activity prediction process.

types, product, type, gallery genre, expensiveness, etc.). Because of error rates in GPS and other data, the behavior model tracks the uncertainty of visits and constructs a probabilistic model, rather than a precise definite record, of visited venues.

Past context-aware mobile information-retrieval systems such as Jiminy [32], stick-e documents [8], and Lamming and Newman's activity-based information-retrieval system [20], among others, were based on calculating the distance between the user's current context (location, time, people, etc.) and corresponding features in the metadata of documents. Rhodes found that recommending documents based on context similarity had only marginal value in the domain of memory assistance [32], but he did not study leisure activity. These systems did not try to classify users' likely current or future activity modes, as Magitti does.

Prior research has shown that the activity predictions will not be completely accurate, but complete accuracy is not necessary for our system. Magitti presents a list containing a mix of content related to the activity probabilities, allowing the user to decide which items are of interest.

Recommender System

For a given user and context, the *Recommender* computes the utility of each content item by combining results from a variety of recommendation models. When all items have been scored, the top results are returned to fill the slots allocated by the activity model.

Recommender systems use a variety of techniques. Two of the most common are collaborative filtering [11, 31], based on identifying clusters of people with similar interest, and content similarity [2], which calculates the similarity of various attributes of content (genre, date, etc.). There are tradeoffs in the techniques and some systems use a combination [23]. Magitti takes this hybrid approach.

In Magitti, the final score for an item is computed based on the results of a large number of models (see Figure 5). A Set Generator maintains a list of all models available, and combines them in an ad hoc fashion depending on the input it receives from other system components. The way models are combined can be specified in a set of rules, or inferred from the user's context. The Set Generator can also learn which models are most appropriate for a user in a given

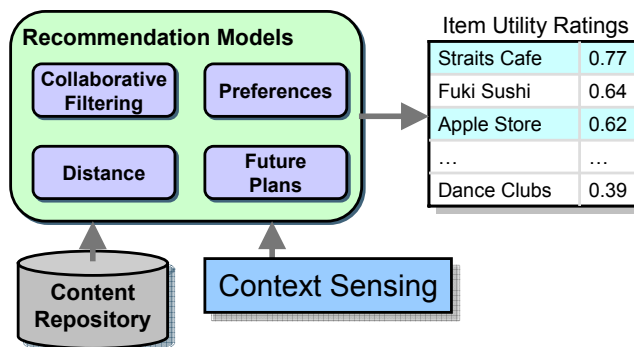


Figure 5. For each activity type, the Recommendation system uses different combinations and weightings of utility models to score recommendable content items.

context. In the current incarnation of Magitti we combine eight models:

- **Collaborative filtering:** This model uses ratings to compute similarities between users and scores each item based on how other similar users rated it.
- **Stated Preferences:** This model scores items according to how closely they match the user's stated preferences (cuisines, noise level, price range, product types, etc.).
- **Learned Preferences:** This model works similarly to the Stated Preferences, but is learned from observed behavior rather than explicitly stated preferences.
- **Content preference:** This model measures the similarity of an item's content to a profile of the user's previously viewed content in web pages and documents.
- **Distance:** This system gives maximum weight to items within a distance range (either entered or inferred from location traces) and uses an exponential decay function to rate the others.
- **Reading:** The system uses a model of when users are most likely to read according to data from the fieldwork.
- **Boredom Buster:** This model reduces the scores of items that have previously been seen, providing diversity to the set of recommended items.
- **Future Plans:** This model temporarily raises scores based on evidence of future plans derived from the Content Analysis, described below.

Content Repository

The *Content Repository* contains the items to be rated and returned by the Recommender for presentation to the user. Each activity type can have one or more content types associated with it. Spatial data structures and caching mechanisms are used to hold the content in memory, avoiding repeated trips to the database and providing for fast location-specific queries.

The recommendable content is indexed according to contextual metadata that indicates the physical situations in which a piece of content might be useful, such as a venue's

business hours and location. The metadata were extracted using finite-state-based entity finders and perl scripts.

Data Context Detection

In addition to detecting the user's physical context, Magitti detects the context of the user relative to personal data from the device. Magitti includes a Content Analysis module that analyzes the content of calendar appointments, viewed documents, and messages to extract information about the user's plans. The fieldwork results indicated that people often discuss leisure activity plans with friends using mobile email and Short Message Service (SMS). To test the potential usefulness of this source of information, we examined the content of a corpus of approximately 10,000 SMS messages generated by students at the National University of Singapore, similar to the Magitti target demographic [15]. Approximately 11% of the messages in the corpus contain information related to leisure activities. Our prototype Content Analysis module currently targets only Eat and See activity planning, with other activities planned for future work. Figure 6 shows an example message along with the extracted information. The extracted information is used by the Activity Model to infer the probability of the user's interest in activities at current or future times.

```
xy:tomorrow what time you be in school? think me and
shuhui meeting in school around 4. then duno still can
see movie or not because duno if a rest want meet for
dinner. how?
```

```
ACTCAT=MOVIE, EAT :: ACTTIME=2007/05/26 16:00 ::
UNCERTAINTY=10 minutes :: TENSE=FUTURE
```

Figure 6. Example SMS message and extracted activity-related information.

FIELD EVALUATION

After the majority of Magitti's features were implemented, we conducted a field evaluation to learn how people would use Magitti in a real world context. Here, we report on our findings regarding Magitti's three primary characteristics, namely supporting serendipitous discovery of leisure activity, predicting user activity, and offering context-aware recommendations.

Eleven volunteers went out with Magitti in the Palo Alto, California area between one and four times each over several days. They visited a total of 60 places over 32 outings, averaging 1.9 places per outing. About half the outings (16) were accompanied by a family member or friend. Participants, who were company employees not working on the project, ranged in age from mid-20s to late-50s, and averaged 37. A further evaluation with Japanese youth is planned, but we felt that an initial study was warranted and would also allow us to identify usability issues, some of which are reported below.

We interviewed each participant about their leisure time and how they typically get recommendations, and then gave

them a demonstration of Magitti. After each outing, participants filled out a questionnaire about their activities. In addition, we logged all their Magitti actions and collected map traces of outings. An experimenter accompanied each participant on one of their outings to observe their use of the system. Finally, we interviewed each participant after they completed all their outings.

Supporting Serendipity

Most people told us they usually go to the same few places, particularly restaurants, and only occasionally try a new one because of the difficulty of finding good recommendations. These comments echoed those in our Japanese field studies. While using Magitti, however, they were very successful at discovering new places. Over half the places they visited (53%) were new to them, including 38% that they had never heard of and 15% they had heard of but never been to. The rest they had been to once or twice (25%) (sometimes long ago) or many times (23%). In 67% of the outings, they visited at least one place that was new to them. Also, during 69% of the outings, people noticed unfamiliar places that they planned to visit in the future.

People expressed delight at finding new places in areas they had been to many times before. Upon finding a new restaurant, one person said, *"Cool! I like that. I would never have found that place if it wasn't for this."* One woman found a new restaurant in the downtown near her home and explained this was a nice contrast, since *"we usually spend 30 minutes walking up and down the street looking at places we've seen a million times before, and always going to ones we know."* One long-time resident said, *"I think it makes life more interesting. It allows you to get out of your daily routine, almost as if you're going to a different city."*

Overall, people were particularly enthusiastic about this aspect of Magitti, namely its effectiveness in helping them discover new places to go. When asked to rate Magitti's overall usefulness, they rated it 4.1 on a scale of 1-5 (5=very helpful). This reassured us that Magitti would be useful for residents and not just tourists or newcomers.

Predicting User Activity

We encouraged users to try different activity types, but the most common activities were in the Eat and Buy categories (consistent with population patterns from our fieldwork and [17]). They visited 30 places to Eat, 27 to Buy, and 3 to Do. Participants were also asked to rate items and to enter preferences to improve the quality of the recommendations. Still, with relatively few outings per person, Magitti's recommendations were not as customized as they would be after more consistent use. In addition, some of the models that inform its recommendations were not available when we ran the study, so the findings from this evaluation represent a conservative estimate of its usefulness.

Magitti started up in "Any" mode, meaning the number of recommendations of each activity type reflected its prediction of the user's interests at that time and place. In

many cases, users quickly switched the list to show only items related to their current activity. They changed activity type an average of 5.1 times per outing, mostly to Eat (1.8 times per outing) and Buy (1.4), and to a lesser degree to Do (0.7), See (0.5), and Read (0.1). They switched it back to Any mode an average of 0.7 times per outing.

We observed that even when Magitti did accurately reflect the user's current activity, people still filtered the list to Eat or Buy, etc. They seemed to do so because, with only four recommendations visible at a time, they wanted *all* of the items to be relevant to their activity. It wasn't helpful to see recommendations for other activity types interspersed in the list. This useful finding suggests that Magitti should have the courage of its convictions and show only items related to its predicted activity. If it is wrong, people seem to find it easy to switch to a different activity. And each time they do, Magitti learns more about their activity patterns, enabling it to improve its predictions for them.

Context-Aware Recommendations

When asked how often the recommendations were relevant and of interest, users gave Magitti an average rating of 3.8 (1=rarely, 5=almost always), a little less than "usually." As one person put it, "*Most of the time, the list contained a mix of useful and not so useful recommendations.*"

Given that Magitti had only a few outings in which to learn from people's behavior, that we had a small set of ratings for collaborative filtering, and that not all of the intelligence in Magitti was fully online, we took this as a good start. However, during the shadowing sessions we noted several factors that affected people's confidence in the system.

Omission: The most common problem occurred when someone walked near a place they liked and did not see it in the list, or at least in the top screen or two. When this happened, people either concluded "*the list did not represent what downtown has to offer,*" i.e., that our content was incomplete, or, more often, that the system didn't know their tastes well enough. This echoes the findings of others who found that small omissions or inaccuracies reduced people's trust in the system [9, 14, 35, 38].

Distance: As designed, Magitti used distance from the user as just one of several factors affecting recommendations – it sometimes showed a "better" recommendation around the corner higher in the list than a "lesser" one next door. The thought was that people would want to choose from the best options near them, but instead they seemed to expect that the closest places would be at the top of the list. Similarly, people remarked on recommendations that were too far away. Even if an item were close as the crow flies, people considered it a poor recommendation if it required driving.

Although people wanted distance to be a strong if not the primary factor in making recommendations, doing so could lead to other problems. The See and Do venues for activities were much further apart than those for Eat and Buy. So when users asked for See or Do recommendations

and had limited the distance setting, they got few or no recommendations, which again led them to think the system content was incomplete.

These findings highlight a need for a recommender to be sensitive to venue density, taking into account whether a user is browsing along the street or searching for the best place to go, and whether they are in a vehicle. COMPASS [38] zooms out its map if the user is traveling quickly, but does not filter content based on speed.

First Item: We observed that people put more weight on the first item recommended than we expected. If it was reasonable, they considered the recommendations good, even if some of the others were inappropriate. One person searched for "home furnishings," and Magitti listed Williams-Sonoma (a furnishings chain store) first and McDonald's second (likely because of a hit on "home" in a review or user comment). Still, the user discounted the second item because the first one was on target. Conversely, if the first item didn't make sense, people were more skeptical of the rest of the list.

Guide vs. Recommender: Surprisingly, people lost relatively less confidence when Magitti recommended a place that turned out to be closed (this should never have happened, but our data about the hours of operation were occasionally incorrect). They were disappointed but they didn't seem to judge Magitti as harshly, apparently because they didn't expect it to take hours of operation into account. Also, although more may have noticed it, only one person mentioned the problem that Magitti continued to recommend a place they had just visited.

It appears that some users were treating Magitti as a location-based information guide rather than a recommender, expecting it to offer information about the closest places but not expecting it to know about hours of operation. Others, though, did expect it to behave more like an expert advisor. In one case, for example, a participant went to lunch with a friend, choosing a Vietnamese Pho restaurant. After lunch, they happened upon a nicer looking Pho place. It didn't appear in the list, but they found it by searching. Sure enough, it was rated four stars to the other one's three stars. His companion griped, "*It would have been good if Magitti had let us know that there was another similar but nicer restaurant right around the corner.*"

Transparency: We noticed that many users spent a fair amount of effort speculating on how Magitti decided which activities and venues to list. Magitti used a complex set of algorithms based on many factors (location, time, preferences, similar users' opinions, prior behavior), and the lack of transparency of the algorithm sometimes confused or even frustrated users. This suggests that Magitti could be improved by offering more cues to help users develop an appropriate user model of its behavior and by allowing users to turn context filtering on and off as desired.

Usability Issues

The evaluation uncovered a wide range of usability issues that have since helped us improve Magitti. We discuss two of the most prominent types of issue here.

User Control: Even when recommendations were seen as appropriate, people wanted more control in managing the recommendation list itself. They especially wanted the ability to sort the items by factors such as rating, price, or distance. They also wanted to be able to remove items from the list for a given session. Even though they might generally like a certain Thai restaurant, for example, if they weren't in the mood for Thai they didn't want to keep seeing it in the list. They also wanted to be able to filter the list not just by general activity (Eat, Shop), but also by subtype (Pizza, Shoe Stores). This last ability has since been added.

Social Use: Over half the Magitti outings involved two or more people, and yet people found it challenging to incorporate Magitti into a social setting. Users tended to walk with their head down staring at the screen or manipulating the interface. Even when they read the information aloud, they didn't always engage the other person. Of the eight who used it with a companion, five reported that the other person got annoyed at them at some point for not paying enough attention to them or their surroundings. The more successful social uses occurred when one person used Magitti ahead of time to narrow down the choices and then asked for input once they were near the venues. Another useful strategy was to stop and look at it together to make a decision and then put it away while walking, except perhaps to check the map. Only one pair shared the device itself, passing it back and forth and reading aloud from it as they walked, although both said they preferred to be the one not holding the device so they could look around. This finding reminds us of the inherent contrast in using a personal device for a social experience and challenges us to find more graceful ways to design small portable devices to fit into inherently social settings.

CONCLUSIONS

Based on extensive field work in Tokyo, we developed a novel context- and activity-aware mobile leisure guide system, code-named Magitti. It predicts the user's current and future leisure activity and uses this, along with models of the user's preferences, to filter and recommend relevant content. The interface uses a novel one-handed, thumb-based interaction. Despite expected technical problems with network delays and lack of comprehensive content, our evaluation suggests that Magitti is useful as a guide for leisure activities in a city and represents a new way to deliver content that is targeted to a mobile user's activity.

ACKNOWLEDGMENTS

Dai Nippon Printing Co., Ltd. 'Media Technology Research Center' and 'Corporate R&D Division' sponsored this research. Ame Elliott participated in early fieldwork. Prof. Fumitoshi Kato and students of Keio University assisted

with street interviews. Pratik Rathod and Tian Luo developed the client context-sensing capabilities.

REFERENCES

1. Abowd, G., Atkeson, C., Hong, J., Long, S., Kooper, R., Pinkerton, M. Cyberguide: a mobile context-aware tour guide. *Wirel. Netw.* 3, 5 (1997), 421-433.
2. Adomavicius, G. & Tuzhilin, A. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Trans. Knowledge and Data Eng.* 17, 6 (2005), 734-749.
3. Ashbrook, D. & Starner, T. 2003. Using GPS to learn significant locations and predict movement across multiple users. *Personal Ubiquitous Comput.* 7, 5 (2003), 275-286.
4. Asthana, A., Cravatts, M., Krzyzanowski, P. An indoor wireless system for personalized shopping assistance. In Cabrera, L. & Sattyanarayanan, M. (Eds). *Workshop on Mobile Computing Systems and Applications*. IEEE CS Press, 1994, 69-74.
5. Baus, J., Cheverst, K., Kray, C. A survey of map-based mobile guides. In Zipf, A., Meng, L. & Reichenbacher, T. (Eds). *Map based mobile services - Theories, Methods and Implementations*. Springer-Verlag, 2005.
6. Begole, J., Tang, J. C., Hill, R. Rhythm modeling, visualizations and applications. *Proc. UIST'03*, ACM Press, 2003, 11-20.
7. Begole, J., Matsakis, N. E., and Tang, J. C. Lilsys: Sensing Unavailability. In *Proc. CSCW '04*. ACM Press, 2004, 511-514.
8. Brown, P and Jones, G., Context-aware Retrieval: Exploring a New Environment for Information Retrieval and Information Filtering, *Personal and Ubiquitous Computing*, (5) 4, Springer-Verlag, 2001, 253-263.
9. Cheverst, K., Davies, N., Mitchell, K., Friday, A., & Efstratiou, C. Developing a context-aware electronic tourist guide: some issues and experiences. *Proc. CHI'00*. ACM Press, New York, NY, 2000, 17-24.
10. Froehlich, J., Chen, M., Smith, I., Potter, F. Voting With Your Feet: An investigative study of the relationship between place visit behavior and preference. *Proc. Ubicomp*, 2006, 17-21.
11. Goldberg, D., Nichols, D., Oki, B. M., and Terry, D. 1992. Using collaborative filtering to weave an information tapestry. *Commun. ACM* 35, 12 (Dec. 1992), 61-70.
12. Harada, M. Japan. In Cushman, G. & Veal, A. (Eds). *Free time and leisure participation: International Perspectives*. CABI Publishing, 2006.
13. Hazas, M. Scott, J. & Krumm, J. Location-aware computing comes of age. *IEEE Computer*, 37, 2, 2004, 95-97.

14. Herlocker, J, Konstan, J., Terveen, L., and Riedl, J. Evaluating Collaborative Filtering Recommender Systems. *ACM Transactions on Information Systems* 22, 2004, ACM Press, 5-53.
15. How, Y. & Kan, M.Y. Optimizing predictive text entry for short message service on mobile phones. In *Proc. Human Computer Interfaces International (HCII 05)*. Lawrence Erlbaum Associates. Las Vegas, July 2005. <http://www.comp.nus.edu.sg/~rpnlpir/downloads/corpora/smsCorpus/>
16. Ito, M., Okabe, D. Matsuda, M. (Eds). *Personal, portable, pedestrian: Mobile phones in Japanese life*. MIT, 2005.
17. Japan Statistics Bureau. Survey on Time Use and Leisure Activities, 2006. <http://www.stat.go.jp/english/data/shakai/index.htm>
18. Krumm, J. & Horvitz, E. Predestination: Inferring destinations from partial trajectories. *Proc. Ubicomp* ACM Press, 2006, 243-260.
19. Kurtenbach, G., Sellen, A. Buxton, W. An empirical evaluation of some articulatory and cognitive aspects of "marking menus". *Human Computer Interaction*, 8, 1, 1993, 1-23.
20. Lamming, M. and Newman, W. Activity-based Information Retrieval: Technology in Support of Personal Memory. In *Proceedings of the IFIP 12th Congress on Personal Computers and intelligent Systems '92* - IFIP Trans., vol. A-14, 1992, 68-81.
21. Lester, J., Choudhury, T., Kern, N., Borriello, G. Hannaford B. A hybrid discriminative/generative approach for modeling human activities. *Proc. 19th International Joint Conference on AI*, 2005, 766 - 722.
22. Liao, L., Fox, D., & Kautz, H. Location-based activity recognition using relational Markov networks. *Proc. International Joint Conference on AI (IJCAI)*, 2005.
23. Middleton, S., De Roure, D., & Shadbolt, N. Capturing knowledge of user preferences: ontologies in recommender systems. *Proc. 1st international Conference on Knowledge Capture. K-CAP '01*. ACM Press, New York, NY, 2001, 100-107.
24. Nishino, H. & Larson, R. Japanese adolescents' free time: Juku Bukatsu, and Government efforts to create more meaningful leisure. *New Directions for Child and Adolescent Development*, 99, Spring 2003, 23-35.
25. Oliver, N., Garg, A., & Horvitz, E. Layered representations for learning and inferring office activity from multiple sensory channels. *Comput. Vis. Image Underst.* 96, 2 (2004), 163-180.
26. Parhi, P., Karlson, A., Bederson, B. Target size study for one-handed thumb use on small touchscreen devices. *Proc. MobileHCI*, ACM Press, 2006, 203-210.
27. Philipose, M., Fishkin, K. P., Perkowitz, M., Patterson, D. J., Fox, D., Kautz, H., Hahnel, D. Inferring activities from interactions with objects. *IEEE Pervasive Computing* 3, 4 (2004), 50-57.
28. Poslad, S., Laamanen, H., Malaka R., Nick, A., Buckle, P., Zipf, A. CRUMPET: Creation of user-friendly mobile services personalised for tourism. *Proc. 3G*, 2001 26-28.
29. Rao, B. & Minakakis, L. Evolution of mobile location-based services. *Commun. ACM* 46, 12 (2003), 61-65.
30. Raj, A., Subramanya, A., Fox, D., Bilmes J. Rao-Blackwellized particle filters for recognizing activities and spatial context from wearable sensors. *Proc. Experimental Robotics (ISER)*, Springer Tracts in Advanced Robotics (STAR), Springer-Verlag, 2006.
31. Resnick, P., Iacovou, N., Sushak, M., Bergstrom, P., Riedl, J. GroupLens: An open architecture for collaborative filtering of netnews. *Proc. CSCW*, ACM Press, 1994, 175-186.
32. Rhodes, B. Using Physical Context for Just-in-Time Information Retrieval, *IEEE Transactions on Computers*, (52) 8, August 2003, 1011-1014.
33. Ricci, F. & Nguyen, Q. N. Acquiring and revising preferences in a critique-based mobile recommender system. *IEEE Intelligent Systems* 22, 3 (2007), 22-29.
34. Schiano, D. J., Elliott, A., Bellotti, V. Tokyo youth at leisure: towards the design of new media to support leisure planning and practice. In *CHI '06 Extended Abstracts*. ACM Press, 2006, 309-314.
35. Schmidt-Belz, B., Aspects of user trust in mobile guides. *Proc. Mobile HCI*. Udine, Italy, 2003.
36. Schwinger, W., Grün, Ch., Pröll, B., Retschitzegger, W., Schauerhuber, A. Context-awareness in mobile tourism guides – A comprehensive survey. Technical Report, Johannes Kepler University Linz, IFS/TK, 2002. <ftp://ftp.ifs.uni-linz.ac.at/pub/publications/2005/0405.pdf>.
37. Smith, J. R., Fishkin, K. P., Jiang, B., Mamishev, A., Philipose, M., Rea, A. D., Roy, S., Sundara-Rajan, K. 2005. RFID-based techniques for human-activity detection. *Commun. ACM* 48, 9 (2005), 39-44.
38. van Setten, M., Pokraev, S., Koolwaaij J.: Context-Aware Recommendations in the Mobile Tourist Application COMPASS. In Nejdil, W. & De Bra, P. (Eds). *AH 2004*, LNCS 3137, Springer-Verlag, 2004.