Focussed palmtop information access through starfield displays and profile matching

Mark Dunlop[1], Alison Morrison[2], Stephen McCallum[1], Piotr Ptaskinski[1], Chris Risbey[2], Fraser Stewart[1]

[1] Computer and Information Sciences University of Strathclyde, Glasgow G1 1XH, Scotland
Mark.Dunlop@cis.strath.ac.uk
[2] Scottish Hotel School, University of Strathclyde, Glasgow G1 1XH, Scotland
Alison.J.Morrison@strath.ac.uk

Abstract. This paper presents two palmtop applications: Taeneb CityGuide and Taeneb ConferenceGuide. Both applications are centred around Starfield displays on palmtop computers – this provides fast, dynamic access to information on a small platform. The paper describes the applications focussing on this novel palmtop information access method and on the user-profiling aspect of the CityGuide, where restaurants are recommended to users based on both the match of restaurant type to the users' observed previous interactions and the rating given by reviewers with similar observed preferences.

1 Introduction

Starfield display technology has been proven to provide quick, dynamic and easy access to large amounts of complex data through use of scatter-plot displays and dynamic queries [1]. These techniques have been shown to be of great benefit in searching in many domains, e.g. house purchases [2], movies [1] and musical pieces guide[3]. However, starfield technology has traditionally been used only for large colour screens and is thus not widely considered suitable for small mobile devices.

In a previous project we showed that a palm-top computer based starfield display was a successful access method for a movie database despite being used on very small, monochrome, low resolution screens [4]. In that work, we compared using traditional palmtop-style access and starfield access to a collection of movies using two axes – year of release (x) and popularity of movie (y) together with direct on screen filters for movie genre (e.g. comedy, thriller…) and film classification certificate (e.g. U, PG…). The results confirmed our belief that starfield displays could be used on such small screens. Figure 1 shows an example search for all non-18 certificate movies excluding comedies. As well as providing fast searching, Starfield displays provide two main benefits over traditional data access methods: dynamic feedback and intuitive transitions from data overviews to focussed searching.

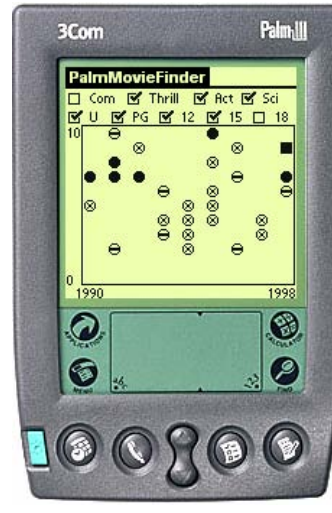

Fig 1: PalmMovieFinder

Dynamic feedback is supported through controls and filters over the dataset. These controls support users searching the database rapidly and provide easy correction for many traditional database problems, e.g. when no data matches a query. In traditional database searching, null queries are notoriously difficult for users to correct – it is very difficult to slightly weaken a complex database query. In contrast, with starfield displays the query is built up in stages and the user knows precisely what (s)he did to cause the null query, thus (s)he can quickly undo that operation. For example, in the PalmMovieFinder a user looking for the lowest certificate thriller would deselect all genres bar thriller and then deselect 18, 15, 12,… in decreasing order. Once there are no matches turning that certificate back on shows the lowest certificate thrillers.

As highlighted in the HomeFinder [2] users can use the searching method to get an overview of the data and rapidly focus in on areas of interest. In the HomeFinder, users are given filters including type of property and house price on a geographic map of an area. When lowering the maximum price filter for a selected type of house, users rapidly learn the expensive areas of a city because they are the first to disappear. This kind of clustering or data overview is very hard to achieve with non-visual interfaces. If a user identifies a suitably priced area near his/her office (s)he can then zoom into that area and naturally restrict further queries to that geographic area.

This paper discusses our development of two starfield displays on palmtops and goes on to propose a combination of starfield displays with recommendation systems as a natural extension to starfield displays.

2 Palmtop Starfield Displays

We have developed two starfield displays on palms: the CityGuide tourist information application based around a geographic map using starfield display to show tourist attractions around a city and ConferenceGuide based around a timetable visualisation

of a conference. Both applications have been developed for high resolution (320x320) PalmOS devices (the CityGuide in colour, the ConferenceGuide in greyscale) and were developed using a combination of PalmOS C and Sybase database storage.

2.1 CityGuide

The CityGuide application is designed around a map-based starfield (c.f. HomeFinder [2]) display to help tourists find attractions around a city. Our current implementation is based on a guide to Glasgow and contains an extensive restaurant guide with some information on cinemas, theatres and pubs. Brown and Chalmers [5] state that "tourists deliberately make plans that are not highly structured and specific, so that they can take advantage of changing circumstances". Our aim in developing this application of mobile starfield technology is to support tourists' unstructured searching of a city centre.

Fig. 2 shows a map of Glasgow city centre with an overlay of all restaurants, represented as squares. The map interface offers typical electronic map features such as zooming and panning: a user tapping on the display in Fig. 2 over, say, Central Station would zoom the map into that location, a further tap zooms further into that location and then the user can tap on a square attraction icon to see the name and then again for more details on that attraction. Users can pan the display by dragging with their stylus and zoom out by clicking on the '-' zoom icon.

Fig. 2: Restaurant guide to Glasgow[1]

On the top right of the display are a set of dynamic filters for controlling what points are shown on the starfield display, here showing type of attractions as restaurant (), the restaurant type (menu) filter () and the restaurant price filter ().

[1] Colour images are available at www.taeneb.com

Single choice filters are controlled by a pop-up menu, for example price in Fig. 3B. Due to limitations in PalmOS, multiple choice filters are controlled via a pop-up window, for example restaurants food type filter is shown in Fig. 3A. The results of a query are displayed immediately on the starfield display as a revised set of attraction icons that match the current set of filters.

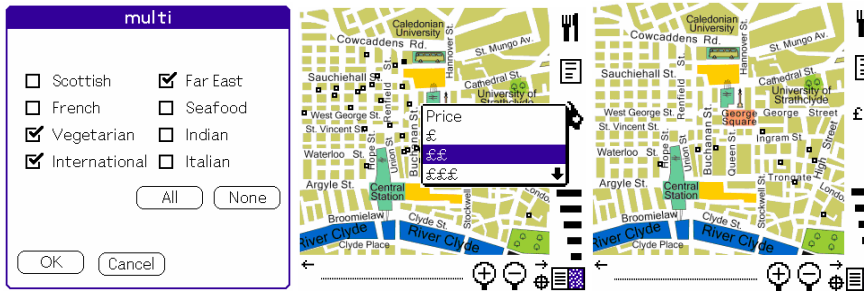

Fig. 3. Restaurant query filters (A: restaurant filter, B: price band filter, C: result of filters)

Brown and Chalmers [5] state that "when choosing where to go to, it is often safer to pick an area with more than one potential facility". Providing this kind of clustering information is one of the traditional strengths of starfield displays. Fig. 4 shows a brief interaction after applying the filters shown in Fig. 3 – here the user zooms into a promising looking area of the map, clicks on one restaurant then clicks on the restaurant name to get full details.

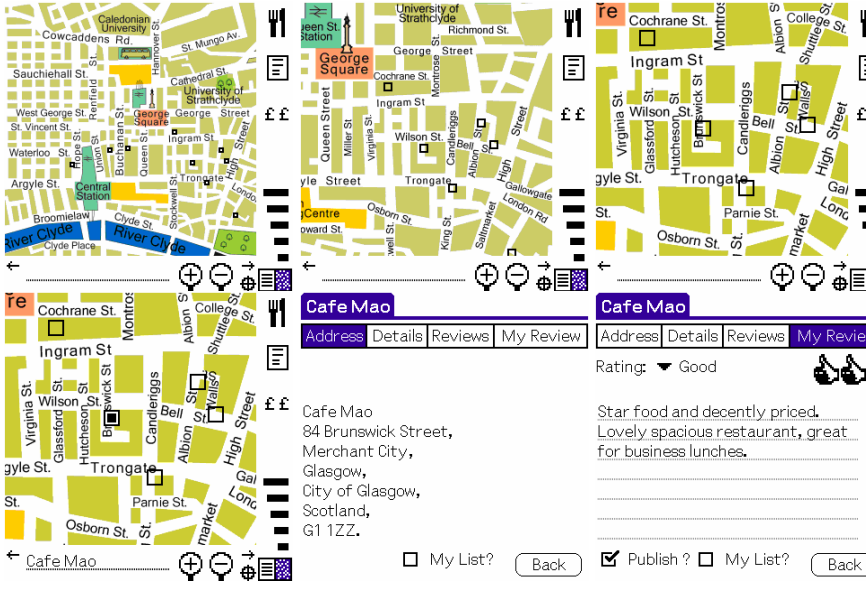

Fig. 4: Zooming into full details on Café Mao after applying fig. 2 filters

Users can mark an attraction as being on "My List" (similar to favourites/bookmarks in web browsers, see last image in Fig. 4) and later filter to show only attractions that have been added to this list. User can also write their own reviews (see last image in Fig. 4), and have these published for others to read. On the bottom right of the map display is a scale bar for relevance and a list view – both of which are discussed later.

2.2 ConferenceGuide

Based around a timetable starfield display the ConferenceGuide initially shows users an overview of a day at a conference (see Fig. 5 for a sample day from our trial conference – EMAC 2003). Here parallel streams are shown as vertical columns with plenary sessions (in case of Fig 5, only breaks) being horizontal bars across all columns – closely reflecting the standard PalmOS Date Book application. Clicking on a session shows its name in the info box at the bottom of the screen with a further click giving full session details (e.g. session name and theme together with a list of talk titles, speakers and abstracts).

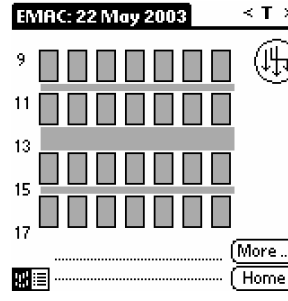

Fig. 5: EMAC conference timetable overview

Filters are provided on session theme and expected audience (not shown in fig 5). The session theme filter is initially set for all themes and users can limit the filter to show only themes they are interested in (e.g. "Consumer Behaviour" or "Social Responsibility"). As with the CityGuide, users can add a session to their "My Sessions" list for later filtering to show only these sessions (a helpful tool for planning time at a multi-parallel conference) and view a textual version of a day's events. The ConferenceGuide also supports inter-delegate communication through messaging and discussion forum services associated with each conference session.

2.3 User Trial

A prototype of the CityGuide and ConferencePlanner were distributed to delegates at the EMAC 2003 conference on "Marketing: Responsible And Relevant?". Twenty delegates were selected by the conference organisers for the trial and were given a Sony PalmOS Clié greyscale device for the duration of the conference. The city guide was populated in advance with restaurant, pub and cinema guide (including "what's

on" information). Some restaurants were populated with reviews but both trial users and other delegates (through combined paper review forms and prize draw entries) were encouraged to write new reviews. At the conference venue the participants were provided with access to synchronise the software and were encouraged to do so at least daily. The feedback was gathered after two days of trial in a form of informal interviews.

A lot of interest in the application was shown by people who had experience of palmtops and by those interested in high-tech applications. People who had other than PalmOS devices, mostly PocketPC-based, expressed considerable disappointment that they could not use their own palmtop. One user found our PDA device in general too small and difficult to read and gave up using the system after the first day of trial. However, in general the users found the interface easy to use and intuitive – even those who had never used a palmtop before. They used the CityGuide mostly for searching for restaurants, they found the starfield interface an easy way of finding a restaurant and the review system very helpful during their selection of a restaurant (in particular delegates found it helpful to be able to read the opinion of other delegates attending the conference and most users added their own reviews).

The overall feedback was positive with many suggestions of extending the data content (such as adding museums, galleries, train timetables etc.) and connectivity of the device to give live update features.

3 Palmtop Collaborative Filtering

Collaborative filtering has proven to be a very successful tool in many domains for helping users select appropriate items from large collections [6] and has been used in tourism, for example, to calculate guided tours [7]. Malone et al [8] describe three forms of information filtering: cognitive (often known as content-based), social (or collaborative) and economic. In line with Balabanović's and Shoham's [9] work on recommendation systems for internet pages, we take a view that combination of content-based and social recommendations are likely to be most effective for a tourism application. Balabanović and Shoham discuss the relative merits of both approaches, primarily that content-based approaches are less prone to individuals with unusual tastes or to a small number of ratings, while social recommendations naturally take more account of significant non-content information that is likely to be missed by content-only recommendations.

In this project we investigated combining a recommendation system with starfield displays to provide a filter on "relevance" in addition to the more traditional database style filters. Our recommendation system is based partly on content-based matching between user profiles and attraction profiles and partly on a social element from similar users' ratings of, say, restaurants. To build each user's profile we use a combination of implicit and explicit ratings. Nichols [10] highlights the problems of achieving a satisfactory number and quality of explicit ratings, where users are requested to explicitly *score* each, in his case, document: "the act of rating alters a user's behaviour from their normal pattern of reading" and that "unless the user perceives some benefit for participating in the system then they have an incentive for leaving". Within the tourism domain, local residents have a clear incentive for writing

reviews of restaurants that their friends/colleagues may benefit from and thus they will benefit from their colleague's and friend's reviews. However, for visitors to a city there is little incentive for a user to write reviews as there is now little direct link between gain and effort. In contrast, implicit ratings are developed by simply monitoring a user's behaviour with the system. Nichols identifies the following potential types of implicit information: purchase, assess, repeated use, save/print, delete, refer/cite, reply, bookmark, examine/read, consider, glimpse, associate, and query. Most of these categories are used to build the user's profile in the CityGuide, as discussed below.

This section reviews our model for combining explicit information with implicit monitoring of the user's interaction and discusses how these are used to drive a relevance filter on a starfield display.

3.1 User profile building and direct matching

For each filter on the city guide we keep an individual user weight for each filter-value (e.g. Italian for the food-type filter value). When our prototype system is started users are asked to fill in a brief questionnaire for each food type (e.g. how much they like Italian food on a 5-point likert scale from "hate" to "love"). These initial scores are given a weight of 0 (hate) to 50 (love) and are then adjusted based on implicit ratings.

Following a scheme similar to Nichols and inspired by relevance feedback techniques in information retrieval [e.g. 11], scores for each matching criteria (e.g. Far-eastern and sea-food) are adjusted for many user actions in the interface. Currently the weights are adjusted as follows (Nichols's categories shown in parenthesis) when a user:

- writes a very good review of restaurant in that type: score +5 (assess)
- writes a good review of restaurant in that type: +3 (assess)
- writes a medium review of restaurant in that type: +1 (assess)
- writes a very bad review of restaurant in that type: -1 (assess)
- filters with this type turned on: +2 (c.f. query)
- add to "my attractions": +3 (bookmark)
- gets more details of a restaurant in this type: +1 (examine)
- read reviews of a restaurant in this type: +1 (examine)

In Nichols scheme reply, examine and glimpse were considered to be time based – the longer a user spent examining, the more important the document. For a mobile setting we felt this to be unreliable as levels of interruption were likely to be much higher, thus more frequently giving misleading measures of, say, how long a user spent reading a restaurant review. As such a fixed increment was used instead of a time-based measure, future experimentation is needed to investigate this decision.

Fig. 6 shows a sample rolling user profile after using the system for some time. Users would not normally see this information but it highlights how the system has developed a simple model of the user's tastes. Here the user has viewed/reviewed more on *International* and *Far Eastern* food than other categories thus we assume (s)he has a preference for these categories and has a relative dislike of sea-food.

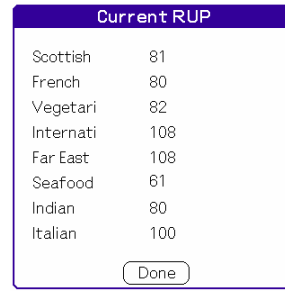

Fig. 6: Sample Rolling User Profile

3.2 End user reviews

As shown earlier (Fig. 4), one of the core elements of the CityGuide is community reviewing where all users can read and write reviews. When a review is submitted the author's current Rolling User Profile is submitted with that review. This allows reviews to be measured for closeness to the current user's views (for example someone who hates expensive Indian food will have different view on an Indian restaurant to someone who loves all Indian food)[2].

Inspired by free text information retrieval techniques [e.g. 12] we calculate user P_u a personalised rating, $PARR_i$', for restaurant R_i as follows:

$$PARR_i' = \frac{\sum \cos(P_{ai}, P_u) * R_{ai}}{\sum \cos(P_{ai}, P_u)}$$

where

P_u = user's current profile for whom we are personalising
P_{ai} = author's rolling user profile at time of submitting review i
R_{ai} = author's rating for restaurant R_i (scaled to between 0 and 1)
cos = cosine function for matching document vectors (see, e.g. [12])
$PARR_i$' is a value between 0 and 1
summations are carried out over all reviews for restaurant R_i

The final personalised review value is given as follows to reduce the effect of only one review and manage zero reviews:

$PARR_i$ =

- 0.5 if there are no reviews
- ($PARR_i$' + 0.5) / 2 if there is only one review
- $PARR_i$' if there are two or more reviews.

Given the user ratings shown in Fig. 6, Fig. 7 shows Glasgow Restaurants rated by PARR given the current database of community reviews. This textual view is a useful complement to the starfield display for when location is not a prime issue in the user's

[2] The current implementation only supports food-type

selection – all filters work identically on the list view and map view and the user can rapidly flip between the two views.

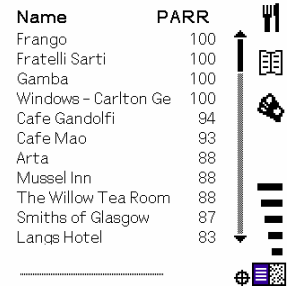

Fig. 7: Glasgow Restaurants rated by PARR

3.3 Combining review ratings and user profile

The Personalised Rating (PARR) uses explicit review ratings of restaurants which are biased towards reviews from people with similar profiles to the user. However, the user's rolling profile does not directly impact these scores (it simply impacts the belief given to others' reviews). In contrast, the Rolling User Profile (RUP) does not take into account restaurant reviews.

These two scores are simply combined into the Combined Attraction Score (CAS): $CAS_i = RUP_i * PARR_i$.

3.4 Combination filtering

Bringing together collaborative community reviews and starfield filtering, we have added a "relevance filter" to the starfield map display. On the bottom right of the display a set of bars represent the openness of the relevance filter (from very open to very restrictive). Fig. 8 shows all restaurants in Glasgow on the left, reasonably tightly relevance filtered in the centre and only the best match on the right.

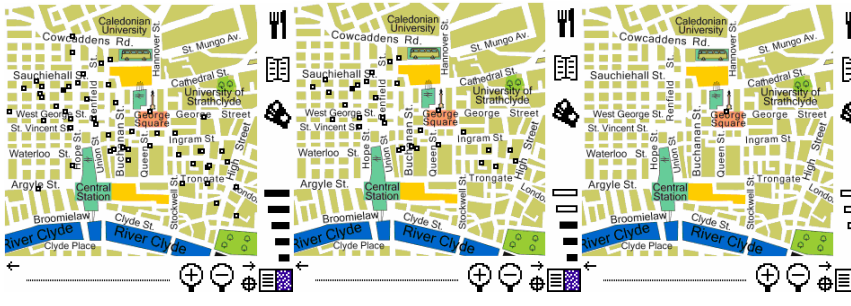

Fig. 8: Relevance Filter on wide, medium and tight settings

The relevance filter is driven by the Combined Attraction Score (CAS) to recommend restaurants based on both their review ratings and their match to the user's rolling profile. This filter works in combination with other filters so, for example, the tightest relevance filter will show the highest CAS ranked attraction that matches the current restaurant-type and price filter settings.

4 Discussion

In this paper we have presented two novel starfield display implementations on palmtop computers: the CityGuide, based around a tourist attraction guide on a geographic plan of Glasgow, and the ConferenceGuide, based around a timetable for a multiple parallel session conference. Both interfaces have proven easy to use in a user trial of conference delegates visiting Glasgow.

The paper has also proposed a combined content and social recommender system for restaurant reviews based on a hybrid explicit/implicit rating system. This rating system is then used to drive a novel interaction tool, the relevance filter, on a starfield display so that users can directly control how much the system recommendations are taken into account when looking for items on the starfield display.

We are currently planning more formal user and technical evaluation of the algorithms and interface. Other context such as weather, user's current context, c.f. [13], and distance of attraction from current location, c.f. [14], are also being investigated as possible inputs to the recommendation system for general tourism attractions (e.g. walks in beautiful but distant botanic gardens tend to lose their appeal in heavy rain). Distance is likely to be more useful in textual lists, e.g. Fig. 7, as starfield displays naturally support zooming into a sub-area of the map, the interaction between these two views is also being investigated. We are also investigating possible improvements to the interface through using recommendations to guide the application of labels to some attractions (c.f. [15]).

In conclusion starfield displays on small devices have been shown to be successful on small devices and combining these with a recommendation system provides a powerful information access interface for small handheld devices.

5 Acknowledgements

The work presented here was mostly funded through a grant from The Scottish Enterprise Proof of Concept Fund – to whom we are most grateful. We also extend our thanks to our trial users and to the organisers of EMAC03 conference for allowing us to experiment on their attendees!

6 References

1 C. Ahlberg and B. Shneiderman, "Visual Information Seeking: Tight Coupling of dynamic query filters with Starfield displays", *Proceedings of CHI '94*, 313-317 & 479-480, 1994.

2 C. Williamson and B. Shneiderman, "The Dynamic HomeFinder: evaluating dynamic queries in a real estate information exploration system", *Proceedings of ACM SIGIR 92*, 339-346, 1992.
3 H. Hochheiser, "Browsers with changing parts: a catalog explorer for Philip Glass' website", *Proceedings of the ACM conference on Designing interactive systems*, 105-115, New York, 2000.
4 M. D. Dunlop and N. Davidson, "Visual information seeking on palmtop devices", *Proceedings of HCI2000*, vol2, 19-20, 2000.
5 B. Brown and M. Chalmers, "Tourism and mobile technology". In: K. Kuutti, E. H. Karsten et al (Eds.), *ECSCW 2003: Proceedings of the eigth european conference on computer supported cooperative work*, Helsinki, Finland, p335-355, Dordrecht: Klewer Academic Press, 2003.
6 P. Resnick , H.R. Varian, "Recommender systems", *Communications of the ACM*, 40(3), 56-58, March 1997
7 J. Fink and A. Kobsa: "User Modeling for Personalized City Tours". *Artificial Intelligence Review*, 18(1): 33-74, 2002.
8 T.W. Malone, K.R. Grant, F.A. Turbak, S.A. Brobst, and M.D. Cohen, (1987), Intelligent information sharing systems, *Communications of the ACM*, 30(5), 390-402.
9 M. Balabanović and Y. Shoham, "Fab: Context-based, collaborative recommendation", *Communication of the ACM,* 40(3), 66-72, March 1997.
10 D.M. Nichols, "Implicit Rating and Filtering", *Proceedings of 5th DELOS Workshop on Filtering and Collaborative Filtering*, 31-36, Budapest, Hungary, 1998.
11 D. Harman, Relevance feedback and other query modification techniques. In Frakes and Baeza-Yates *Information Retrieval: Data Structures and Algorithms*, Ch 11, 241-263, 1992.
12 C.J. Van Rijsbergen, *Information Retrieval (second edition).* Butterworths, 1979.
13 K. Cheverst, N. Davies, K. Mitchell, A. Friday and C. Efstratiou, "Developing Context-Aware Electronic Tourist Guide: Some Issues and Experiences", *Proceedings of CHI'2000*, Netherlands, 17-24, 2000.
14 M. Brunato and R. Battiti, "PILGRIM: A location broker and mobility aware recommendation system", in *Proceedings of IEEE PerCom2003*, 2003.
15 J. Tatemura, "Dynamic Label Sampling on Fisheye Maps for Information Exploration", *Proceedings of the ACM Working Conference on Advanced Visual Interfaces*, 238-241, Palermo, Italy, 2000.