CrossMark

Where's everybody? Comparing the use of heatmaps to uncover cities' tacit social context in smartphones and pervasive displays

Andreas Komninos[1] · Jeries Besharat[1,2] · Denzil Ferreira[3] · John Garofalakis[1,2] · Vassilis Kostakos[3]

Abstract We introduce HotCity, a city-wide social context crowdsourcing platform that utilises user's current location and geo-tagged social data (e.g., check-ins, "likes" and ratings) to autonomously obtain insight on a city's tacit social awareness (e.g., "when is best time and where to go out on a Saturday night?"). HotCity is available as a mobile application for Android and as an interactive application on pervasive large displays, showcasing a heatmap of social buzz. We present the results of an in-the-field evaluation with 30 volunteers, of which 27 are tourists of the mobile app, compare it to a previous evaluation of the pervasive display app and also present usage data of free use of the pervasive display app over 3 years in the city of Oulu, Finland. Our data demonstrate that HotCity can communicate effectively the city's current social buzz, without affecting digital maps' cartography information. Our empirical analysis highlights a change in tourists' foci when exploring the city using HotCity. We identify a transition from "individual [places]" to "good [areas]" and "people [choices]". Our contributions are threefold: a long-term deployment of a city-wide social context crowdsourcing platform; an in-the-field evaluation of HotCity on mobile devices and pervasive displays; and an evaluation of cities' tacit knowledge as social context as a denominator in city planning and for the development of future mobile social-aware applications.

Keywords Mobile maps · Social networks · Urban context · City dynamics · Heatmaps · Data visualisation

✉ Andreas Komninos
akomninos@ceid.upatras.gr

1 Department of Computer Engineering and Informatics, University of Patras, Patras, Greece

2 Computer Technology Institute and Press "Diophantus", Rion, Patras, Greece

3 Center for Ubiquitous Computing, University of Oulu, Oulu, Finland

🖄 Springer

1 Introduction

Travelers increasingly rely on mobile applications to access relevant city information. In particular, they will seek information on public transportation (e.g., buses, trains), landmarks they can visit and other venues such as restaurants, cafes, and shops. Many applications exist for popular big cities that attract global tourism. However, the rich yet user agnostic information that these services provide can be bleak and non-informative for users who simply try to decide what to do and where to go. On such occasions, more favorable information is obtained from friends' or even strangers' recommendations. To answer questions such as "which is the best restaurant on a Friday evening?" or "which are the famous and busy areas on Saturday?", we still rely on the concierge.

To overcome the scarcity of tacit knowledge in mobile information-finding applications, we created HotCity. HotCity is not just an application. HotCity consists of an infrastructure for sharing geo-tagged high-level and user-subjective place information harvested from locals' social networks (e.g., Foursquare). With a map and heatmap mashup, HotCity highlights mobility traces and preferences, all accessible through a mobile application. A heatmap visualization of social context, added as a middle layer between the spatial background (i.e. street map) and additional point of interest overlays provides users insight into locals' tacit preferences in the city.

Our work heavily draws on the concept of *venues* in social networks, which has become popular with social networks based exclusively on them (e.g., Foursquare), actively used by several million users. Other networks (e.g., Facebook, Google+) also integrate formal venue representations with their existing structure. These representations allow users to explore venues and allow content tagging (e.g., checking into them, rating, indication of a favorite spot). Research by Lindqvist et al. (2011) reports that this data is generated not randomly but under a very specific context: users will only tag venues they consider important, interesting or indicative of a social identity and lifestyle choice. Access to this personalised yet important data opens up significant new opportunities for the provision of services to a range of interested parties, in particular for tourism services. HotCity is one example of how such information can be exploited to provide city's tacit social context to everyone.

In this paper, we discuss how the use of heatmaps in the HotCity service affects the user experience and behaviour by providing tacit social context to tourists in two distinct use cases: via large, public pervasive displays, and via a mobile application for smartphones. The use case of pervasive displays was examined in a previous paper (Komninos et al. 2013a) so here we focus on discussing the evaluation of the smartphone use case and compare the outcomes between pervasive displays and smartphones, as appropriate. Our insights from both use cases show that despite the increased visual complexity, as compared to simple digital maps, the users' cognition in terms of mental demand and perceived effort is very modestly impacted. Use of the heatmap results in a shift of urban exploration, from a focus on individual venues and towards urban areas. We further report on the insights gained

through statistics of real in-the-wild use of the pervasive display application over 3 years in the city of Oulu where it was deployed. The paper begins with a review of relevant literature, introduces the HotCity platform architecture and discusses the two field experiments and in-the-wild use of the system in the next three sections. Finally, the paper concludes with a section discussing the outcomes and implications of our work.

2 Related work

When visiting a city, one traditionally seeks information on a city map, reads a guidebook, or visits a tourism office. Although useful to display cartographical information effectively, such city maps do not reflect citizen-generated information. By citizen-generated information, we mean information such as "Which are locals' favorite restaurants?", "Where to find best prices?". To overcome such limitation, in recent years, a new category of city guidebooks emerged, dubbed "personalized guides". These guides are quite often written by previous city visitors, locals, and can be, to some degree, autonomously and algorithmically deducted by counting visits to tourist points of interest (POI's) (Souffriau et al. 2008). However, these providers aim at the general public and are not tailored to the individual visitors' needs, tastes or interests. For our literature review, we focus on tourism-oriented system infrastructures including mobile and/or access to urban mobility patterns (Table 1).

2.1 Ubiquitous computing systems for tourism

The term "ubiquitous computing" refers to the use of mobile and pervasive computer systems to deliver information and services to user in a range of indoor or outdoor environments. Contrary to a tourism office, and more similar to the city maps and guides, a mobile device such as a smartphone can be carried around while visiting a city. For example, Chevrest et al.'s (2000) GUIDE mobile applications offered a digital version of a city guide. Context-aware mobile recommender systems, such as Magitti (Bellotti et al. 2008), inferred the user activity and displayed suggestions for related activities. COMPASS (van Setten et al. 2004) provided tourists with recommendations, centralized in a registry maintained by third party service providers. CRUMPET's (Poslad et al. 2001) "blackbox" approach allows service and content providers integration using dedicated interfaces, where the user can request information and recommendations based on his location. Using artificial intelligence and agent systems, Gullivers' Genie (O'Hare and O'Grady 2003) took a proactive stance to deliver content, displaying nearby POIs and multimedia (e.g., photos, videos) autonomously. Use of extended contextual information, including weather, location, time, social network data and user personal preference for POI recommendation, has been proposed in Meehan et al. (2013). Braunhofer et al. (2014) describe a solution to the cold-start problem in recommendations by acquiring a user's personality profile and asking them to rate a number of randomly selected venues, acting as a training set, in various lab-based sessions. The work was later extended to evaluate the system against further

Table 1 Summary of literature review on mobile and urban mobility systems

Paper	Research goal	Year	Participants	Duration
Alghamdi et al.	E-tourism: mobile dynamic trip planner	2016	Dataset	–
Alves et al.	Semantically describing a POI with Wordnet	2009	Dataset	–
Bellotti et al.	Context-aware mobile recommender by inferring current and predicting users' future activities	2008	Lab; urban	1 day; 1 week
Bergé et al.	Exploring smartphone-based interaction with overview + detail interfaces on 3D public displays	2014	Lab	2 sessions
Braunhofer and Ricci	Selective contextual information acquisition in travel recommender systems	2017	Dataset	–
Braunhofer et al.	Techniques for Cold-Starting Context-Aware Mobile Recommender Systems for Tourism	2014	Lab, dataset	4 sessions
Calabrese and di Lorenzo	Opportunistically use location data to estimate weekday and weekend travel patterns	2011	Dataset	–
Calabrese et al.	Using eigen decomposition on Wi-Fi data to identify locations	2010	Campus	14 weeks
Carter et al.	No app needed: enabling mobile phone communication with a tourist kiosk using cameras and screens	2017	–	–
Cheverst et al.	Understanding the user experience of a context-aware guide	2000	Urban	4 weeks
Churchill et al.	Using public displays to inform community building and maintenance	2003	Campus	6 months
Colomo-Palacios et al.	Towards a social and context-aware mobile recommendation system for tourism	2017	Lab	1 session
Crandall et al.	Using geotagged photos (Flickr) and content analysis to find popular landmarks	2009	Dataset	6 months
Fisher et al.	Using online maps' interaction to highlight geographic areas	2007	Dataset	1 year
Girardin et al.	Using cell phone network data and geotagged photos to analyse tourism activity	2008	Dataset	3 months
Grubert et al.	The utility of magic lens interfaces on handheld devices for touristic map navigation	2015	Urban, lab	2 sessions
Herzog	Recommending a sequence of points of interest to a group of users in a mobile context	2017	Lab	1 session
Jiang et al.	Modeling human mobility patterns in streets	2009	Urban	6 months
Kim and Kotz	Modeling mobility from Wi-Fi access points data	2005	Campus	2 months
Koch et al.	Information radiators: using large screens and small devices to support awareness in urban space	2017	–	–
Kostakos et al.	Using Bluetooth scans to sense, model and visualise urban mobility and copresence networks	2010	Urban	1 year
Meehan et al.	Context-aware intelligent recommendation system for tourism	2013	–	–
O'Hare and O'Grady	Intelligent agents for tourism guidance	2003	–	–

Table 1 continued

Paper	Research goal	Year	Participants	Duration
Quercia et al.	Measuring audiences outdoors in circumvented areas	2011	Dataset	1.5 months
Souffriau et al.	Use applied machine learning for creating mobile guides based on categorized POIs	2008	Dataset	–
Tammet et al.	Crowdsourcing geo-tagged databases to locate, describe and rate potential POIs	2013	Dataset	–
Umanets et al.	GuideMe-ATourist guide with a recommender system and social interaction	2014	Lab	1 session
van Setten et al.	Context-aware mobile recommender by leveraging users' profile and current location	2004	Lab	1 session

datasets to determine the contextual factors most relevant for recommending items to users (Braunhofer and Ricci 2017). More recently, Colomo-Palacios et al. (2017) describe the successful implementation of a POI recommendation system based on location and a biologically-inspired AI recommendation engine that includes POI rating and opinion mining. Umanets et al. (2014) demonstrate that collaborative filtering (via social networks) and the user's own previous visits and ratings benefits a recommendation system for POIs. In a system that recommends routes between two points, Herzog (2017) show how the qualitative properties of the POIs in an area (e.g. type, suitability during various times in the day, presence or absence of other similar POIs in the area) can be used, briefly evaluate the generated routes in the lab.

Public information displays across a city can also be used to provide information captured from several sources (Hinrichs et al. 2013; Linden et al. 2012). Early research on public displays was mostly conducted on single-purpose bespoke public displays, for example Plasma Posters (Churchill et al. 2003) or GroupCast (McCarthy and Costa 2001), while more recently. Recent advances in public display technology have enabled increasing numbers of displays to be deployed and installed in public locations. These deployments are increasingly making a transition from static "broadcast" displays to interactive ones (e.g. Koch et al. 2017 introduce the hybrid concept of information "radiators"—broadcasting to, and interacting with users). This transition to interactive displays, where members of the public are empowered to control and use the display, has opened a range of new research challenges and at the same time has broadened the design space for public displays. Whereas on "broadcast" displays the primary challenge is designing for the effective sharing of information with the public, interactive displays' main design requirement is that of interaction: designing and implementing a mechanism for members of the public to browse, navigate and identify information that the display makes available. The ad-hoc synergy between public displays and user devices is still an ongoing subject of research, with some work focusing on supporting this synergetic approach (Grubert et al. 2015; Bergé et al. 2014) while others focusing on eliminating it (Carter et al. 2017).

2.2 Urban dynamics for tourism applications

To improve the recommendations provided to tourists from ubiquitous computing systems, researchers have attempted to capture and quantify urban dynamics, i.e. to use sensing methods for the capture of urban movement across a city. The premise behind this idea is that by identifying the mobility patterns in an environment, a system can recommend to a tourist the venues that are worth visiting. These methods include geo-tagged photos, social network data, mobile phone logs, smart card records, taxi/bus GPS traces, and Bluetooth sensing (Calabrese and di Lorenzo 2011; Girardin et al. 2011; Jiang and Yin 2008; Kostakos et al. 2010; Bao et al. 2012; Yang et al. 2013; Majid et al. 2013; Quercia et al. 2011). These demonstrate that it is possible to develop a better understanding of city-dwellers' space use over time, and subsequently inform important decisions about development, growth, and investment across a city, as well as tourism. In other words, understanding how various groups of people move in a particular area, and when, provides better context for understanding the types of potential audiences for services in those areas, but also in terms of long-term investment and development decisions (Quercia et al. 2011). As seen in Table 1, only 4 out of 24 of the reviewed literature actually had an urban experiment or evaluation (from a 1-week to a 1-year long deployment). Most of previous work use existing datasets to test and evaluate methods of obtaining urban context. This is understandable, especially considering the challenges of collecting longitudinal and widespread sensor data from cities. However, we argue that such approaches only capture a snapshot of what happened, and not what is happening right now. It does not account or reflect the dynamics of a city (e.g., new business, roads, events). Obtaining a wide-net citizens' data is challenging and some researchers have employed special equipment or infrastructure access to obtain it, i.e., phone providers (Girardin et al. 2011; Quercia et al. 2011), satellite access (Fisher 2007; Jiang and Yin 2008), router firmware (Calabrese et al. 2010; Kim and Kotz 2005) or Bluetooth scanners (Kostakos et al. 2010).

To achieve higher granularity, researchers are increasingly turning to alternative datasets. Analysis of user-generated content is becoming increasingly popular, for example using geotagged photos to extract "place" and "event" information from Flickr (Rattenbury et al. 2007). This approach was adopted by one of the first attempts at identifying tourists and visitors in a city (Girardin et al. 2011) by analyzing geo-tagged photographs from a 3-year period, focused on identifying locations visited by individuals exhibiting short and focused activity in terms of photographs taken. Alghamdi et al. (2016) discuss a system for recommending optimal routes for trip planning, with the aim being to provide a route between two points, that would provide the most satisfaction to the users. In their work, POIs that for intermediate parts of the route are ranked by mining Flickr for volume of images, as a proxy to POI importance.

Previous work has used granular Wi-Fi data (Calabrese et al. 2010; Kim and Kotz 2005), but so far has been limited to campus scale. Often, mobility analysis attempts to cluster locations based on similarity to each other in terms of volume of visitors. For instance, researchers have demonstrated a bottom-up approach to grouping locations into clusters that exhibit similar temporal mobility patterns in terms of

volume of visits (Kim and Kotz 2005), and subsequently labels these clusters according to a tacit understanding of both the locations as well as the mobility patterns there (Calabrese et al. 2010).

2.3 Recommending areas instead of single venues

Venue recommendation has benefited greatly from research such as that, mentioned in the previous section. Typically, venue recommendations are provided as lists, or displayed as markers on digital maps. A large concentration of map markers may hint to the user that a particular area of the city is populated with highly recommended venues, however, this approach is not without problems. Large marker densities can result in occlusion and selection problems, especially on small mobile screens. Furthermore, map markers can only present spatial and not qualitative information—just because a certain area contains fewer markers than another, doesn't mean that that area is less popular (for example, its venues might be heavily popular or it might have many venues that are simply not displayed in digital datasets, e.g. streetfood vans). In this work, our goal is to evaluate the user usability aspects of highlighting popular *spaces*, i.e. geographical areas, instead of popular points of interest. This is a critical aspect of differentiating our work from previous literature. To illustrate the point, let us consider an example, whereby the user gets a recommendation for a highly popular venue to visit. Upon arrival, the venue might be fully booked (as is likely, given its highly recommended status). If the venue is located in an area where other alternatives cannot be easily found, then the visit will result in disappointment and frustration for the users. On the other hand, as can be seen in many tourist guide books, experienced guides will recommend not just single venues but whole areas and also the temporal context appropriate for a visit (e.g. "During the evening, take a trip to New York's Little Italy"). Quite often, the entire area (instead of venues) can be the focus of a tourist's visit, for example, taking a stroll through the market district, a historic neighbourhood, a pedestrian street, along the river etc. Such recommendations are entirely missed in literature, where the focus remains on recommending individual POIs. Our hypothesis is that, providing a highlight as a background layer, instead of a foreground layer, informs the user of location importance without obfuscating their understanding of the underlying maps' cartography. We extend previous work with crowdsourcing approaches (Fisher 2007; Tammet et al. 2013) by incorporating social network activity as it happens (e.g., Foursquare) to dynamically update and harvest local citizens whereabouts and preferences. "Traveling as a social activity [is] not considered" and the potential of incorporating external content is seldom exploited (Schwinger et al. 2005). Our work's contributions provide an effective sharing of geo-fenced high-level and subjective place information that is harvested from locals' social networks (e.g., Foursquare), by means of a mashup highlighting mobility traces and "heatmaps", at one's pocket reach.

Notably, overlapping dynamic content on a map is not novel. Alves et al.'s (2009) KUSCO mined the web for creating ontologies based on semantic associations to POIs for enriching the description of a place; Crandall et al. (2009) organized photos' location based on visual and temporal features to pinpoint
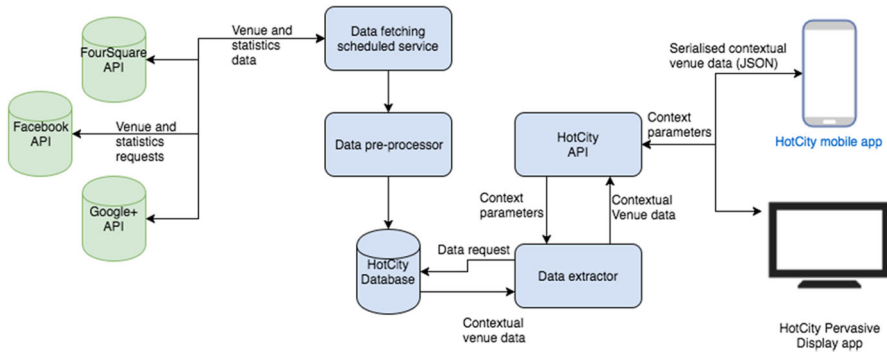
Fig. 1 HotCity service infrastructure architecture

from where did the picture got taken; and Tammet et al.'s (2013) crowdsourcing approach from geo-tagged databases to locate, describe and rate tourism targets in any area of the world. With heatmaps, Hotmap (Fisher 2007) focused on location highlighting techniques for the most frequently viewed locations. Tammet et al.'s work (2013) is largely similar to HotCity in terms of presentation, but differs in functionality: it does not provide a system to the general population to assess in real-time where the "heat" is, utilising snapshots of the available data, without reflecting the city's current state. This is where HotCity makes the biggest contribution: it is a real-time, dynamic and city-wide social context platform available to everyone.

Summarising the previous literature, we can thus conclude that while the harvesting of urban social context has been used in previous studies (i.e. the interactions between people and their environments as recorded via actions in social networks), this has mostly focused on algorithmic recommendation of specific *individual* venues or points of interest to users, and not wider areas where a tourist might visit. In the cases where literature has studied the geospatial visualisation of aggregated tacit social context information (e.g. via heatmaps), it has not concerned itself with how it affects the users' cognitive and decision-making processes during application use. Finally, the use of temporal context and urban dynamics has not been considered in this category of literature. Our work is thus positioned precisely so as to address these gaps in literature.

3 HotCity architecture

Our HotCity prototype was deployed in the city of Oulu, a medium-sized city in Northern Finland. HotCity aggregates information from online sources such as social media (e.g., Foursquare) and citywide sensors (e.g., Wi-Fi location positioning) provided by the panOULU infrastructure and produces citizen-based spatial knowledge (Fig. 1). The data is aggregated on a webserver and is disseminated to a variety of devices, through a purpose-built API (OpenOULU[1]).

---

[1] http://docs.ubioulu.fi/.

In a previous publication (Komninos et al. 2013b) we investigated the use of this data in pervasive displays, but for this paper, our primary goal is to explore users' interaction with HotCity's generated information using a mobile device. On a mobile application we developed, we present a map of downtown Oulu, upon which we superimpose our social data visualizations, allowing users to easily discover novel, popular and interesting areas of Oulu according to user-specified contexts (e.g., time: "Now", "Friday evening"; people: "10 or more").

3.1 Extracting social network information

We rely on three popular social networks as data sources, Foursquare, Facebook and Google Places APIs. For this study, we rely exclusively on data from the Foursquare API, which allowed us to retrieve current information such as real time check-ins and also historically derived information such as total check-ins at a specific location. It's worth mentioning at this point that although Foursquare used to offer a single application for checking in and discovering places, the company has now split the functionality across two applications. The Foursquare mobile app and website are used to discover places (i.e. as a city guide) while the Swarm mobile app allows users to check in. Our research data was collected prior to the split of the service in the two apps. Regardless, the company offers a single unified API for both functions to this date, so our methods are still applicable.

Information from Foursquare checkins is sufficient to represent the social vibe and dynamics of a city, as demonstrated by Kostakos et al. (2010), Noulas et al. (2012) and by Komninos et al. (2013b). Although the other APIs offer valuable data, we exclude them from our implementation for reasons which are discussed in the concluding sections of the paper. For the purposes of our study, we geo-fenced various areas within the city, covering popular commercial and social areas of interest within the city (henceforth: a "listening post"). To identify such posts, we interviewed locals who provided hints on where these locations should be, where they circled all the famous areas around Oulu on certain times of the day on a map. We subsequently collected data from the social networks' APIs at 30-min intervals, per listening post, and retrieved the local businesses data in the vicinity of the listening post. We chose 30-min intervals to *minimize* our margin of error per check-in, as the Foursquare API does not provide the check-in time. Foursquare's check-in timeout policy keeps a user checked into a place for a maximum of 3 h or until he checks into another venue. Finally, our aggregated data consisted of: timestamps, number of check-ins, total amount of check-ins, tags, likes and ratings. As such, the data *does not capture distinct check-ins*, but rather how many people appear to be checked into a venue at any point in time. Each of our listening post contained multiple venues. In addition, if social media users added a new place by "checking-in", our dataset was dynamically extended to include it in our analysis. In previous work (Komninos et al. 2013b) we find that even if the frequency of social interactions (i.e., check-ins) is low, collecting this data over a period of time gives accurate descriptions of urban dynamics within an area of interest.

For scalability and interoperability, our system is capable of disseminating the collected information to a range of ubiquitous devices and services, including

desktop, mobile, wearable and city infrastructure devices through a purpose-built web API that returns JSON formatted data to devices. In this paper, however, we focus on the delivery to native mobile applications, which we describe in Sect. 4.

3.2 Deployment smart city infrastructure

The architecture described in the previous section runs on the PanOULU smart city infrastructure. The city of Oulu is equipped with a citywide ubiquitous computing infrastructure, which includes a free public access Wi-Fi network. A consortium of municipalities, public organizations and ISPs provides the network. A large proportion of the devices that use the network are owned by international visitors that enjoy the open and free wireless Internet access as an alternative to exorbitant international roaming fees of commercial mobile data. At the same time the network has been a valuable resource for numerous projects (e.g., Linden et al. 2012).

3.3 Generating social context heatmap visualisations

Before the description of our mobile application, we discuss here the process of generating heatmap visualisations. To achieve this, on the mobile side, we used the Google Maps' heatmap extension library Mapex.[2] The library requires a list of spatial coordinates (points defined by latitude, longitude and intensity) and once provided with this, it draws a semi-transparent heatmap visualisation as a layer over a simple Google map. As such, we created a relevant call in the HotCity API which returns a list of the venues in a given area (defined as a geographic bounding box via two sets of coordinates) and under temporal context, including the day (e.g. "Friday") and hour for which the information is required (e.g., "15"—meaning 3 pm).

The API call works as follows. It first considers the geographic bounds passed to it and retrieves all venues inside these bounds. Then, for each of these venues, it sums the check-ins captured from Foursquare on that particular day and time of day (used as point intensity), using a subset of our data dating back 60 days from the time of query, and returns a list of venue objects in JSON (latitude, longitude and intensity). We use the 60-day threshold to ensure that seasonality is taken into account (and thus, not displaying for example winter popularity for venues, when the query happens during the summer). In the returned subset, the data is normalised to the nearest integer on to a scale of [10–100], thus giving venues with no check-ins a minimum intensity of 10 and the point with the maximum number of check-ins an intensity of 100. An example for how a single point's intensity is calculated follows in Table 2 below (using a 30-day threshold to shorten the example and assuming the venue with the most check-ins has 60 check-ins).

Each point on the map "radiates" its intensity to a spatial radius (in our case, set empirically to 70 m). Within the intensity radius of each point, the colours of the heatmap change according to its total intensity, hence a single point would radiate for 70 m with a colour from violet to red, according to the intensity of that point.

---

[2] Mapex heatmap library https://github.com/chemalarrea/Mapex.

Table 2 Calculating the intensity of a point (venue) for Fridays 15:00, dataset threshold 30 days, query on submitted on 2013-06-07 15:45

Point ID	Timestamp	Check-ins
4dc0d8a4ae60533f34e50c81	2013-06-07 15:36:37	1
4dc0d8a4ae60533f34e50c81	2013-05-31 15:06:37	2
4dc0d8a4ae60533f34e50c81	2013-05-31 15:36:37	0
4dc0d8a4ae60533f34e50c81	2013-05-24 15:06:37	1
4dc0d8a4ae60533f34e50c81	2013-05-24 15:36:37	5
4dc0d8a4ae60533f34e50c81	2013-05-17 15:06:37	3
4dc0d8a4ae60533f34e50c81	2013-05-17 15:36:37	4
4dc0d8a4ae60533f34e50c81	2013-05-10 15:06:37	3
4dc0d8a4ae60533f34e50c81	2013-05-10 15:36:37	3
	Total check-ins	22
	Normalised intensity	37

The library on the local client then draws the final heatmap, calculating the colour of each pixel according to the summed intensity at that pixel based on the average intensity of overlapping points' radii.

# 4 Sharing social context data with ubiquitous mobile applications

4.1 The pervasive display HotCity app

In Komninos et al. (2013a) we presented the outcomes of a controlled field experiment in Oulu using a pervasive display, part of a city-wide ubiquitous computing infrastructure, which includes a free public access Wi-Fi network and ubiquitous outdoor large interactive touch-screen public displays called "UBI-Hotspots". In this paper we will discuss how the HotCity concept performed in an implementation as a mobile app (Android), using the same experiment set-up as in Komninos et al. (2013a) and we will discuss the comparisons with the pervasive display app as appropriate (the pervasive display interface is shown in Fig. 2c, d). For completeness, in certain places we will also provide additional data about participant behaviour not presented in Komninos et al. (2013a).

4.2 HotCity mobile app in detail

Our application consists of a main screen where users can see a map of downtown of Oulu. Users can interact with this map and gain access to more activities by using buttons that are visible around the screen (Fig. 2a, b).

In particular, in the top left corner we have placed a label that shows the day and the hour for which the heatmap shows data. There are 3 icon buttons on the top right corner of the screen: the clock icon (Fig. 2a, label 1) enables users to change the heatmap depending on the day and the time they choose, the human icon (Fig. 2a, label 2) downloads and shows all the POIs on the map, and the reload icon (Fig. 2a, label 3) refreshes the heatmap data with current information. On the bottom right
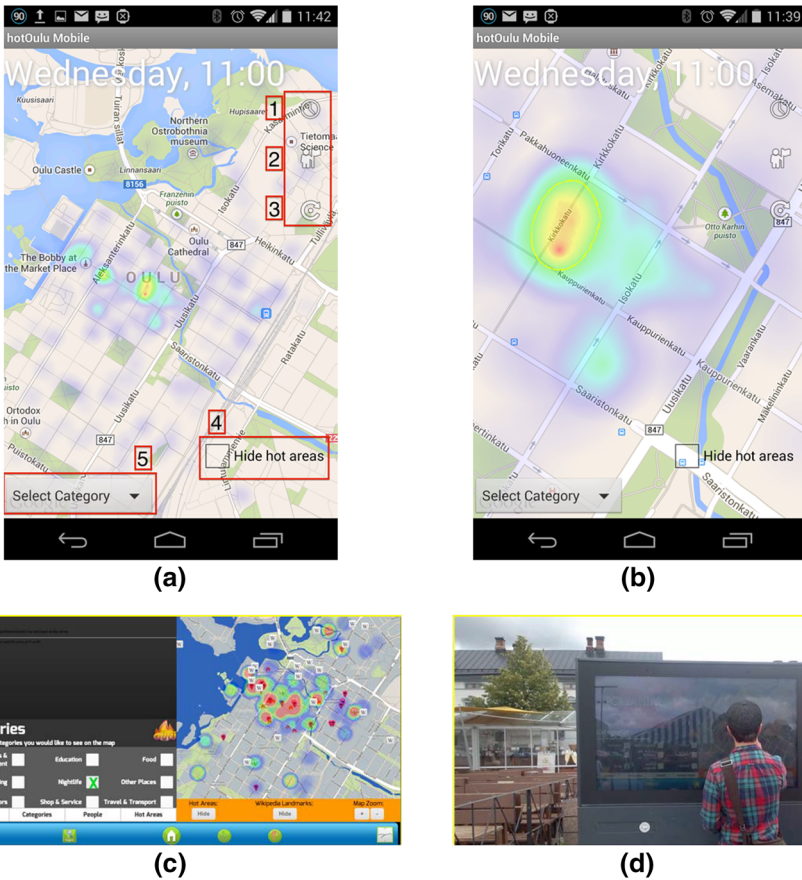
Fig. 2 The HotOulu mobile app main interface (**a**) and detail of the heatmap layer (**b**). Also shown, the pervasive display app interface from Komninos et al. (2013a) (**c**) and display used by a participant (**d**)

corner we have placed a checkbox (Fig. 2a, label 4) that toggles the heatmap overlay. Users can also filter the POIs by category. On the left bottom corner we have placed a select combo box (Fig. 2a, label 5) that shows all the categories that we have data for. With this layout, we developed a service that conveys a range of contextual information to the user in a multi-layered view:

- Layer 0 (2D street map): this layer forms the base on which other layers can be superimposed (i.e., the map of the city itself). It consists of a standard 2D street map, as provided by the Google Maps API. Users can zoom and pan the map using touch gestures (pinch and drag);
- Layer 1 (time-varying heatmap): a layer that shows a heatmap for the specific day and time (Fig. 2b). The map affords users spatial understanding of their surroundings. The heatmap is generated from check-in data throughout the hours of the day and highlights which areas of the city are active or quiet at specific times, this providing a spatial visualization of an area's social "buzz". This

heatmap supports a temporal control, managed by the user: the user can select specific days of the week and time of day (Fig. 3c). The heatmap uses a standard "rainbow" color scheme to indicate the socially active areas. Red shades are used for hot areas while blue for the quiet ones. The heatmap is semi-transparent to allow viewing of the underlying map;

- Layer 2 (POIs): an overlay with categorized points of interest (POIs) within the city, provided by Foursquare's APIs. The POI layer can be filtered by selecting sub-categories using our prototype's interface (Fig. 3a, b).

5 Evaluation

5.1 Task and experiment set-up

Our experiment set-up was identical to that in Komninos et al. (2013a). We adopted a scenario-based approach to assess the heatmap's information affordances when visiting a new place. Our scenarios included non-trivial planning tasks. These assumed that the user is in a completely new location for which they are unprepared. They also assume that the user has not made previous plans and has to rely on whatever information can be gathered by our service. Finally, they require the user to not only think about the places which they will visit, but also how far apart they might be, considering they do not know anything about this new place. The scenarios are:

- Task 1: "It's your second day in Oulu and you are walking in the street. Having nothing specific to do, you have enough time to find and visit a landmark (outdoor place)".
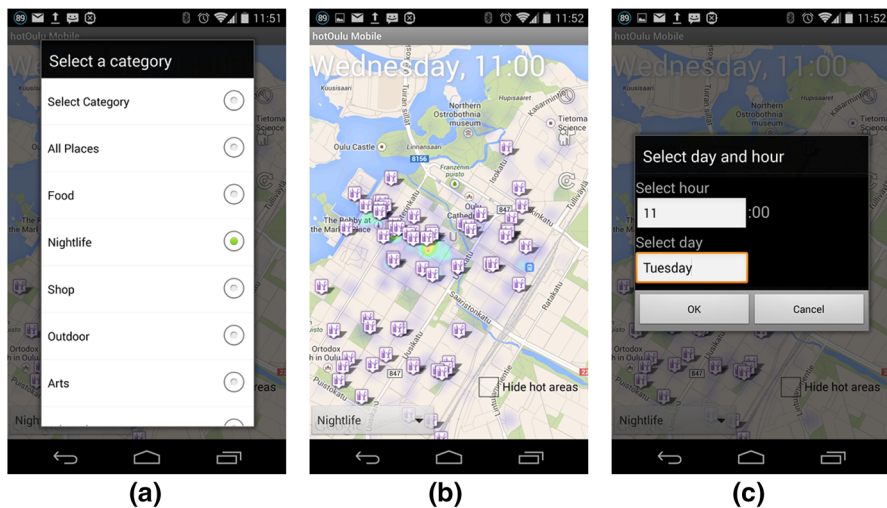
**Fig. 3** Screenshots of our application's control interfaces

- Task 2: "It's 2 PM on Wednesday and you are near Mannerheim Park, so you have to find a restaurant around this area to go" [this park is next to the location of our experiment].
- Task 3: "It's Friday evening (8 PM) and you want to go out for a drink to see the nightlife of Oulu but first you want to find a restaurant to eat".

We conducted a between-subjects experiment design with two distinct groups of users for evaluation. One group was permitted to use Layers 1 (the time-varied heatmap) and 2 (POIs), while the other was limited to just Layer 2 in order to perform the three tasks. For each intervention, we recorded the elapsed time to complete the task, followed by a short user subjective questionnaire to determine complexity, ease of use, learnability, confidence and perceived utility. We further investigated if users' perceptions were affected by their own use of social networks, with a separate questionnaire. All participants were told how the service works by the team that conducted the survey.

5.2 Participants

In order to get a satisfactory and informative result we decided to conduct an experiment with 30 participants (16 male, 14 female), 27 of whom were non-locals (i.e. real tourists from various nationalities). Three researchers approached people of all ages at a fixed outdoor location in the city and invited them to participate. We conducted our experiment in a fixed location because the application and the evaluation process focuses on POI selection using different visual approaches and not on navigating participants around the city. The experiment took place in July, on good weather days, so participants were not hindered by environmental conditions. The participants' age ranged from 14 to 47 years old (m = 27.3 years old, stdev = 7.1). Based on the sample of almost exclusively tourists and their age range (in Valls et al. 2014 it is reported that the average age of tourists who tend to use the Internet and thus, by extent, are likely to use mobile and in-situ digital tourism tools, is approx. 37 years old), we can argue that the field experiment has reasonable ecological and external validity.

Before the task, we asked the participants how familiar they were with various social networks and how often they use them (e.g., checking in a place, tagging etc.), with responses recorded on Likert scales (0 = Don't know this social network, 1–5 = never, rarely, sometimes, often, very often). We found that most of the participants were familiar with Facebook followed by Twitter. Foursquare ranked in fourth place while Instagram is ranked in third (Fig. 4a). Furthermore, we measured the frequency of use of social media for specific actions that relate to spatial context and we found that the actions of checking in, tagging locations and friends to photos rank close to each other with the checking in action first (Fig. 4b), again using Likert scales (0 = Don't know this action, 1–5 = never, rarely, sometimes, often, very often). This was important, as we needed to be sure that the participants knew what a social network is and if they are familiar with the actions that it offers. The results are directly comparable to those in Komninos et al. (2013a) showing that the two subject groups had similar background and attitude to social network use.
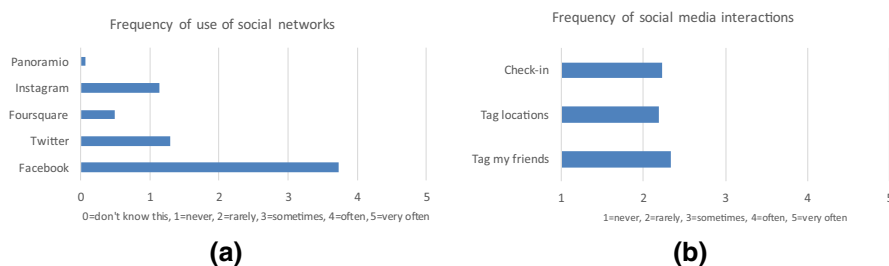
Fig. 4 Participants' engagement with social media (**a**) and frequent actions within these (**b**)

In addition to the above, we asked participants how often they use different sources of information when they visit a new and unfamiliar city (Fig. 5). The responses were measured on Likert scales (1–5 = never, rarely, sometimes, often, very often). For most of the participants the main source of information comes from their family and friends (1st) followed by tourism websites (2nd). Fewer participants are using information from community travel websites (3rd) that ranks close to tourism websites. Finally, traditional information sources like brochures, information offices and guidebooks seem to be used rarely. This can indicate that most people trust and prefer information that comes from people they know and are familiar with and then they use other sources such as the Internet. Again the results are near identical to those in Komninos et al. (2013a), hence validating the similarity of participants' background and attitude towards tourism information search.

5.3 Quantitative results

We recorded the time taken to complete each scenario, and additionally we recorded the interactions of each user with the system: icon clicks and screens viewed. The results of user interaction with the interface elements are shown in (Fig. 6).

The mobile app use data was divided by the heatmap groups and tested for normality. For both groups (heatmap, no-heatmap), data is normally distributed (Shapiro–Wilk, $p > 0.05$). Although we did observe measurable differences in the means of interaction types between the two groups, with the Heatmap group
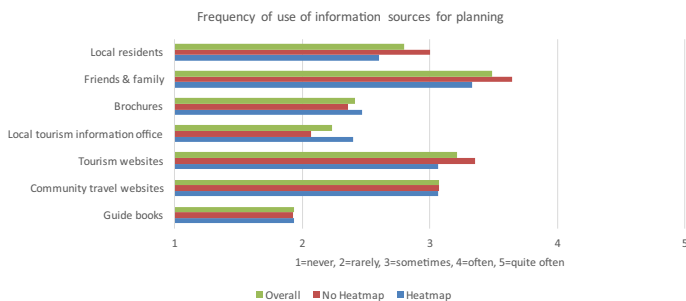
Fig. 5 Participants' habitual trip planning sources

showing slightly higher statistics on both measures (Category clicks $m_{NH} = 6.43$, $m_H = 8.0$ and POI clicks $m_{NH} = 12.0$, $m_H = 15.0$), an independent sample t test showed that the differences are not statistically significant. This result shows that the heatmap did not affect the interaction with the application for the interaction categories that were common across participant groups, which means the heatmap did not encourage participants to investigate more POIs. This result is in line with those in Komninos et al. (2013a), also shown Fig. 6, where we also found that the heatmap did not result in statistically significant differences in category selection and POI selection behaviour. Notably, comparing between the two applications, we observe that users employed the category filters more liberally in the pervasive display case, compared to the mobile in both the heatmap (t = 3.44, p < 0.05, two-tailed independent sample t test) and no-heatmap conditions (t = 6.32, p < 0.01, two-tailed independent sample t test). The comparison of POI clicks was not statistically significant in either condition between the applications. We can only attribute this behaviour to the different layout in the two applications, as in the pervasive display the cognitive cost of switching categories is less for the user, as the map is not occluded, as is in the mobile map (the change only occurs on the left panel of the display, while on the mobile the map is occluded by the category list).

In Komninos et al. (2013a) we did not report on the quality of the POIs selected by participants, but since this data is available we will present it here and compare it with that from mobile app users. With regard to the actual POIs that were selected by our participants as their response to the task at hand, we note that the heatmap visualization on the mobile seems to drive participants towards selecting a smaller variety of distinct points ($n_{NH} = 43$, $n_H = 38$). The same observation applies to the pervasive displays ($n_{NH} = 31$, $n_H = 27$). To examine the quality of those POIs in terms of popularity, we multiplied the total check-ins at each point by the number of times that POI was selected as a final choice, then divided by the sum of these
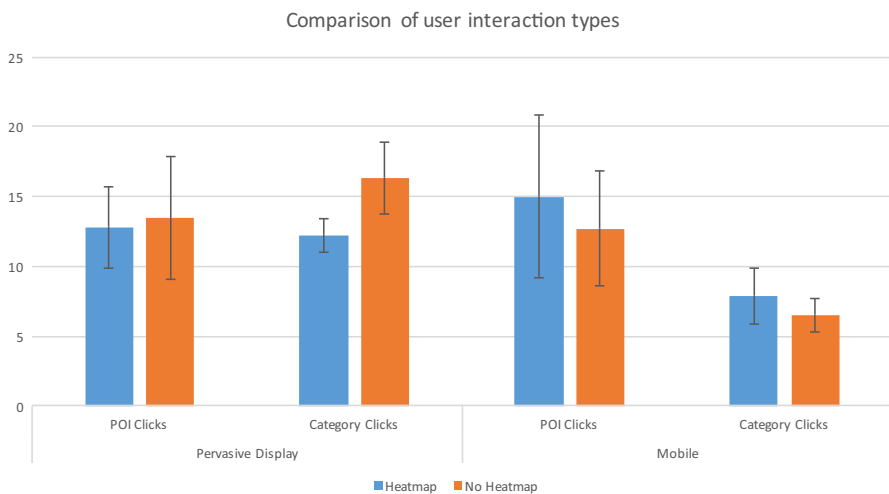

Fig. 6 Participants' observed interaction types

multiplications by the total number of participants who provided a final choice. This metric provides an estimate of average POI popularity per group member. Further, by looking at the popularity of the places selected in all tasks, as indicated by total venue checkins, we observe that when using the mobile app, the heatmap group was able to select on average more popular POIs than the no-heatmap group in tasks 2 and 3, showing that the heatmap visualization can result in better user experience for more complex information seeking tasks (Fig. 7). On the other hand, while using the pervasive display, the groups not using the heatmap selected POIs with greater popularity, showing that in that case the heatmap visualisation did not improve their ability to find popular places.

In Table 3, we show some measures of the areas that included participants' selections for each task. In both the mobile app and pervasive display, the heatmap resulted in exploration of larger areas across all tasks (considering the sum of the explored areas in m^2, with covered areas calculated from the convex hulls defined by the outermost selected POIs). We further plotted the POIs selected by the participants against a heatmap depicting the overall popularity of city areas in Oulu, which we call the "Safety Zone" (i.e. the objectively popular part of the city, shown by the red area of the heatmap). Care must be taken here not to confuse this with the heatmap information shown to participants, as this heatmap represents the total popularity of venues based on check-ins, irrespective of temporal context.

It should be borne in mind that the location of our experiment was in the midst of this most popular part of the city, hence it is natural to expect that participants would seek places that are not too far away from it, regardless of condition. However, our observations show a different behaviour between the heatmap and no-heatmap conditions. We observe that with the heatmap enabled, in both the mobile and pervasive display apps, participants shifted their attention to select more POIs that cluster near the bounds of the most popular part of the city (65.57 and 72.73% respectively), compared to the no-heatmap condition (45.31 and 36.21% respectively). This behavioural shift towards exploring the bounds of the safety zone came
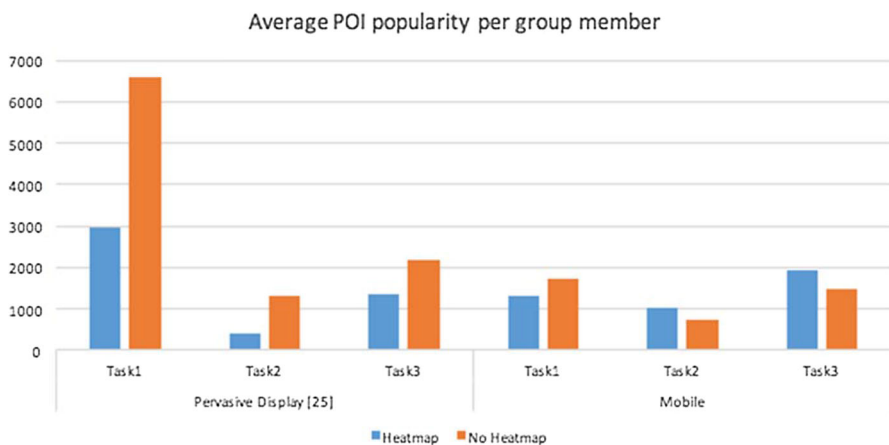
Fig. 7 Popularity of selected POIs for each task

Table 3 Participant-selected POI spatial characteristics (area coverage)

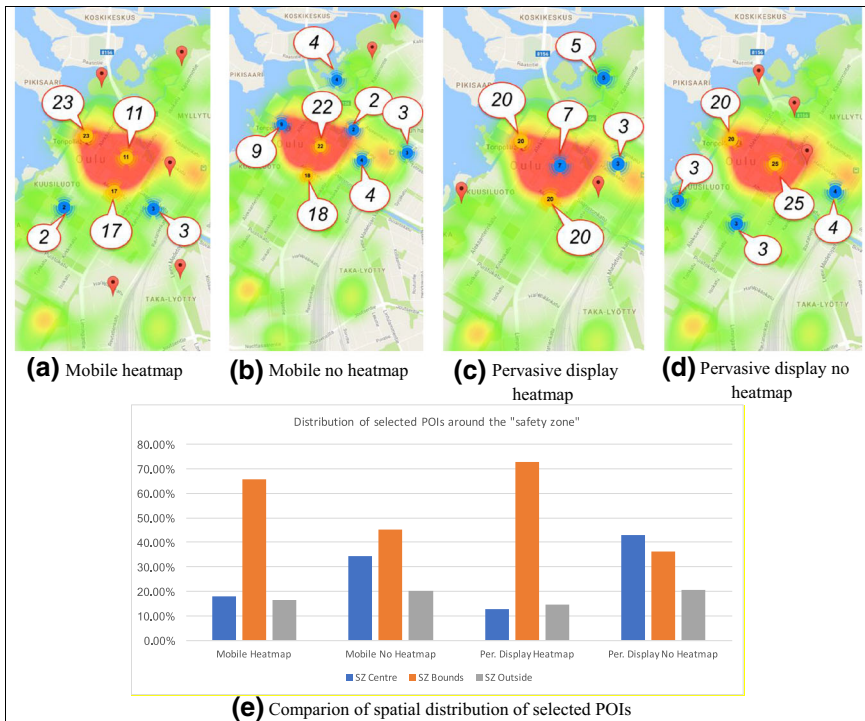|	Covered area	Task 1 (m^2)	Task 2 (m^2)	Task 3 (m^2)	Overall
Mobile	No heatmap	506,766.16	63,725.89	332,506.51	902,998.56
	Heatmap	1,149,845.65	10,515.16	548,231.46	1,708,592.27
Pervasive Displays	No heatmap	390,867.94	2,681.61	179,894.95	573,444.49
	Heatmap	301,055.26	6,177.62	380,142.44	687,375.32

mostly by at the expense of exploring the centre of the safety zone, as can be seen by the nominal percentage differentials ($\Delta_M = -16.34\%$, $\Delta_{PD} = -30.38\%$) and to a lesser extent from exploring fewer points outside the bounds of the safety zone ($\Delta_M = -3.92\%$, $\Delta_{PD} = -6.14\%$). This is exemplified by Fig. 8 below. In this figure, the yellow and blue markers show "clusters" of two or more POIs that are spatially very close together and the number in the marker shows the cluster size, i.e. number of spatially proximal points (we use this visualisation to prevent the problem of marker occlusion).

Combined with the larger area covered using the heatmap condition (Table 3), it can be concluded that the heatmap caused a difference in the information seeking behaviour and allowed for more "adventurous" exploration, however confining it still within the "safety" of being mostly conducted in the most popular part of town.

Finally, in terms of time taken to complete the tasks (seconds), again data was divided by the heatmap groups and tested for normality. The data in this case is normally distributed for all tasks in the Heatmap group, and for Task 2 in the No Heatmap group. The application recorded every action along with a unique ID used to identify each participant and a timestamp. The timestamp provided information about the time taken to complete each task and action. Moreover, for each task the three researchers noted the time the participant started the task and the time that made the final selection in order to confirm the duration of each task. For Task 2 an independent sample t test shows a very small difference in the time that participants took to complete the task ($m_{NH} = 53.6$ s, $m_H = 53.07$ s) which is not statistically significant. For Task 1 ($m_{NH} = 121.0$ s, $m_H = 89.93$ s) and Task 3 ($m_{NH} = 64.67$ s, $m_H = 104.67$ s), a Mann–Whitney test shows that the difference is only statistically significant for Task 3 ($p < 0.05$), which the Heatmap group took longer to complete. In comparison, the task time completions were far longer in Komninos et al. (2013a) and in that experiment we didn't discover any statistically significant differences attributable to the use of the heatmap in the pervasive display. The results are summarized in Fig. 9.

5.4 Qualitative results

After each task, every participant completed a NASA TLX questionnaire, in order to obtain their subjective opinion on the retrieval process. The findings are summarized in Fig. 10 and Table 4 shows the relevant comparison statistical test results. For all subscales, we found the distribution of responses to be non-normal,

(a) Mobile heatmap **(b)** Mobile no heatmap **(c)** Pervasive display heatmap **(d)** Pervasive display no heatmap

(e) Comparion of spatial distribution of selected POIs

Fig. 8 Spatial distribution of selected POIs using the mobile app Heatmap (**a**) and No Heatmap (**b**), and the pervasive display Heatmap (**c**), No Heatmap (Komninos et al. 2013a) (**d**) and comparison (**e**)

hence Mann–Whitney U tests show that the null hypotheses can be rejected for Mental Demand, Perceived Performance and Effort (p < 0.01), in contrast with the

Fig. 9 Duration of tasks
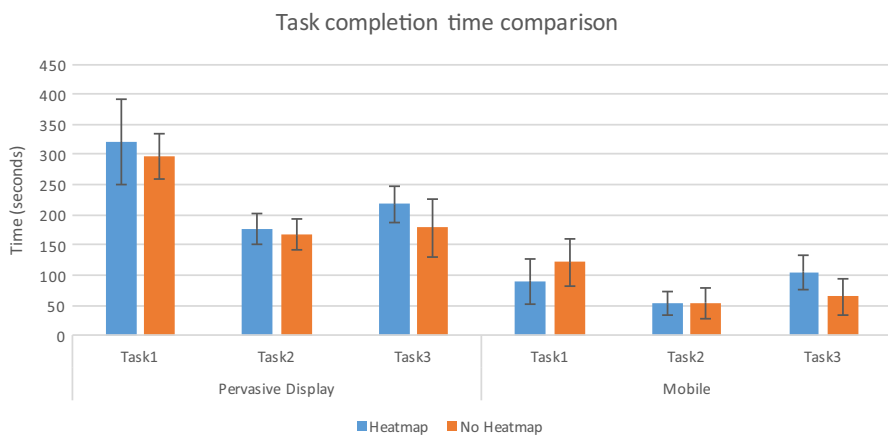
Comparison of NASA-TLX ratings
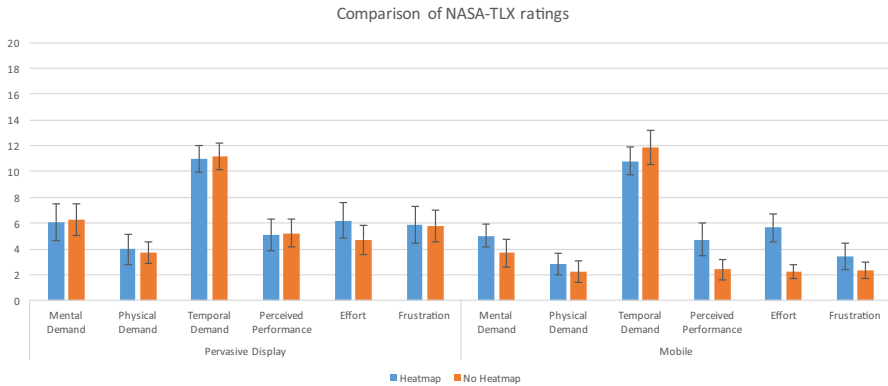
Fig. 10 NASA-TLX subjective evaluation scores

Table 4 NASA-TLX comparison statistics for mobile users (Mann–Whitney U tests)

|	Mean		St. Dev.		Z	p
	Heatmap	No heatmap	Heatmap	No heatmap		
Mental Demand	**5.000**	**3.667**	2.989	3.730	− 3.340	**0.001**
Physical demand	2.822	2.267	2.908	2.824	− 1.190	0.234
Temporal demand	10.822	11.822	3.548	4.543	− 0.773	0.439
Perceived performance	**4.733**	**2.422**	4.364	2.671	− 3.677	**0.000**
Effort	**5.644**	**2.200**	3.713	1.869	− 4.926	**0.000**
Frustration	3.400	2.333	3.543	2.140	− 1.682	0.093

Results with statistical significance ($p < 0.05$) are shown in bold

results in Komninos et al. (2013a) where we found no statistically significant differences in any of the questionnaire subscales.

These results show that our users' perceptions converge towards greater confidence in their choices when using a heatmap on the mobile device, since they appear to believe that they were more successful in identifying a good POI candidate with the help of the heatmap. This, however, comes at a cost of higher mental demand and required effort. Despite the statistically significant differences, it should be observed that on a positive note, Mental Demand, Effort and Frustration are quite low overall for both groups (for NASA-TLX, lower is better on all subscales except Performance). On the other hand, a negative aspect is that perceived Performance is quite low, even for the Heatmap group, while for both groups, Temporal Demand is relatively high. As concerns perceived performance, its low nominal value can be attributed to the fact that in our experiment, participants were not actually asked to visit and spend time at their selected venue, to verify whether they indeed believed they made a good choice. As such, it is understandable that participants might have reserved reasonable doubt about their performance, hence the result should be interpreted as better performance with the
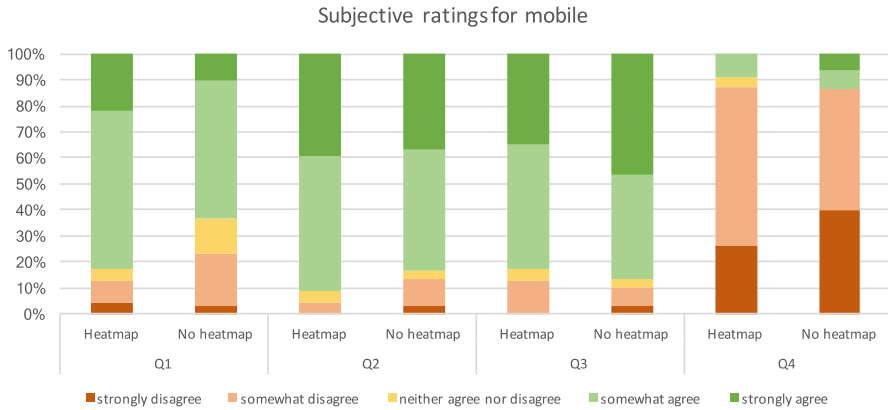
Fig. 11 Subjective evaluation scores (mobile)
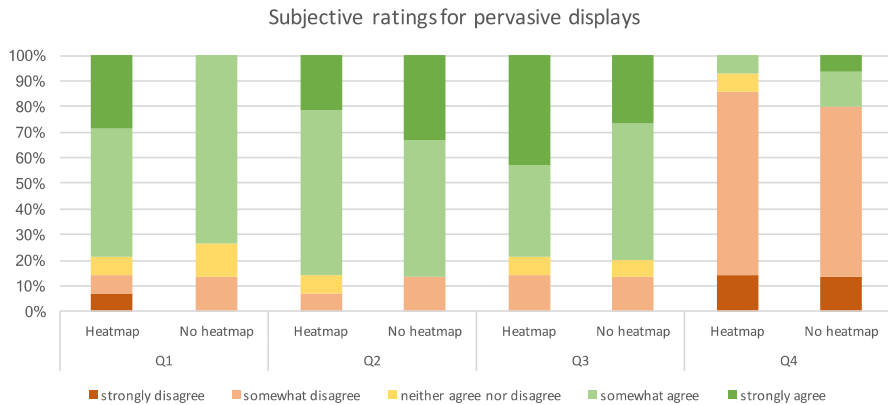
Fig. 12 Subjective evaluation scores (pervasive display Komninos et al. 2013a)

heatmap in identifying the best possible choice from those presented to them in the app, within a reasonable timeframe and in the context of the scenario given to them.

Further from the NASA-TLX questionnaire, we asked participants to provide their subjective feedback on a 5-point Likert scale (1 = Strongly Disagree, 5 = Strongly Agree) on a further four questions, adapted from the System Usability Scale, shown below.

Q1. I feel confident that I made good choices with help from this system
Q2. I would use this system frequently on a visit to an unfamiliar city
Q3. Most people would learn to use this system very quickly
Q4. The map display was needlessly complex.

Figures 11 and 12 shows the distribution of responses for each question. In terms of confidence in choices, we see more positive responses when using the heatmap in

both the mobile and pervasive display devices. Participants seem equally positively disposed to using our system in both the mobile and pervasive display versions, regardless of including the heatmap. Similarly we do not observe significant differences in their perception of how easy to learn the system would be, indicating that the heatmap element appears to be intuitive. The level of negative responses in Q4 indicates that participants also did not feel that the heatmap overly complicates the map on the mobile device, although they felt the map on the pervasive display was slightly more burdened by this element.

Lastly, we collected free-form feedback from participants, in the shape of comments at the end of the experiment. Participants without the heatmap visualization gave overall positive feedback on the ease of use and utility of the system (8 comments). Amongst them we also found many suggestions for improvements (12 in total) and the most desired feature was more information about venues such as opening hours, venues local comments and some sort of rating (participants suggested a star system or a 1–5 scale). Two participants also commented on the utility of the application as being best suited to tourists or young people. Heatmap participants enjoyed the idea of visualizing social context (12 comments in total), as indicated by their enthusiastic feedback:

- *Very nice app! young people tag themselves in cool places (because they wanna tell they were there) and then non locals people can find those places.*
- *Very fun app if you want to avoid popular places/be exactly there.*
- *It's quite interesting to see that people really check in to those places, in that way heatmap info is interesting.*
- *Good application because there are places where people have really visited. It really helps. I would use this application when visiting in a new city with pleasure.*
- *It was good that with this app you know where people go.*
- *Good idea! If a place is popular is also a good one. The app doesn't require much effort to use.*
- *It is good thing that this application shows the most popular places where people have checked in. This is good when traveling abroad. Usually there are good food where are the most popular places.*

Participants also provided some comments on improving the app, which again related to being able to see more information such as opening hours and menus (6 comments). Two participants also felt that the application could be a little faster (in terms of downloading and presenting data) and a further four commented that the categories were "too broad" and needed further refinement.

6 Free-running deployment of the pervasive display HotCity app

As a final step in our paper, we present some findings from the free-running deployment of the HotCity app in the pervasive displays in Oulu. After the end of the experiment in Komninos et al. (2013a), we left the application running
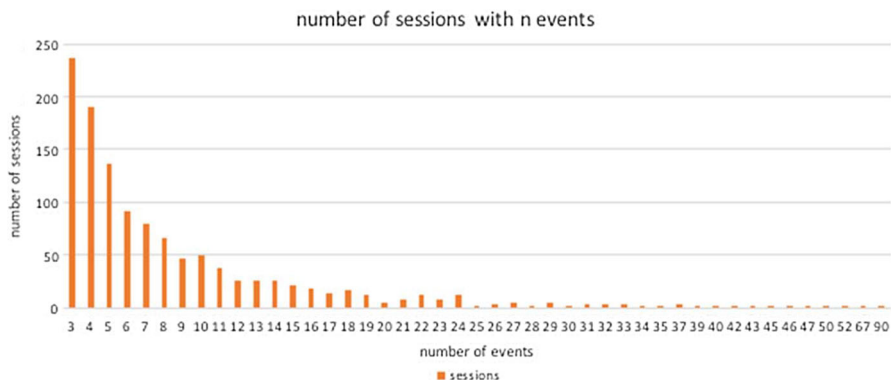
Fig. 13 Distribution of sessions based on number of interaction events during free use

Table 5 Interaction events in sessions with 5–24 events

Event	Occurrences	Percentage
Switch to home tab	1030	17.66
Switch to heatmap tab	746	12.79
Heatmap updates	583	9.99
Heatmap toggle	249	4.27
Switch to categories tab	738	12.65
Select a POI category	1155	19.80
Switch to people tab	523	8.97
People updates	366	6.27
Wikipedia POI clicks on map	80	1.37
Wikipedia POI toggle on map	254	4.35
POI clicks on map	109	1.87
Total events	5833	100.00

on the Ubi-Hotspots of the city, allowing any user to freely access it for their own use. Although by no means a controlled experiment, and noting that we are unable to obtain any user characteristics from this mode of use, we present here some interesting statistics of use in the 3 years that the application has been available.

In total, the application has been used 1188 times. From this usage, the majority of sessions has very few interactions (clicks on the various interface elements), namely 1–3 interaction events. For our analysis of use thus we consider only those sessions which have between 5–24 events, as these make up approximately 60% of all usage cases, i.e. 713 sessions (Fig. 13).

One immediately interesting observation is that only 6% (42) sessions actually involved the selection of any POIs for viewing. The remaining sessions involved only the manipulation of heatmap (selecting a time and day to update the heatmap) or seeing venues with more than a given number of persons checked in currently. In Table 5 below, we outline the distribution of interaction events, where it is clearly shown that
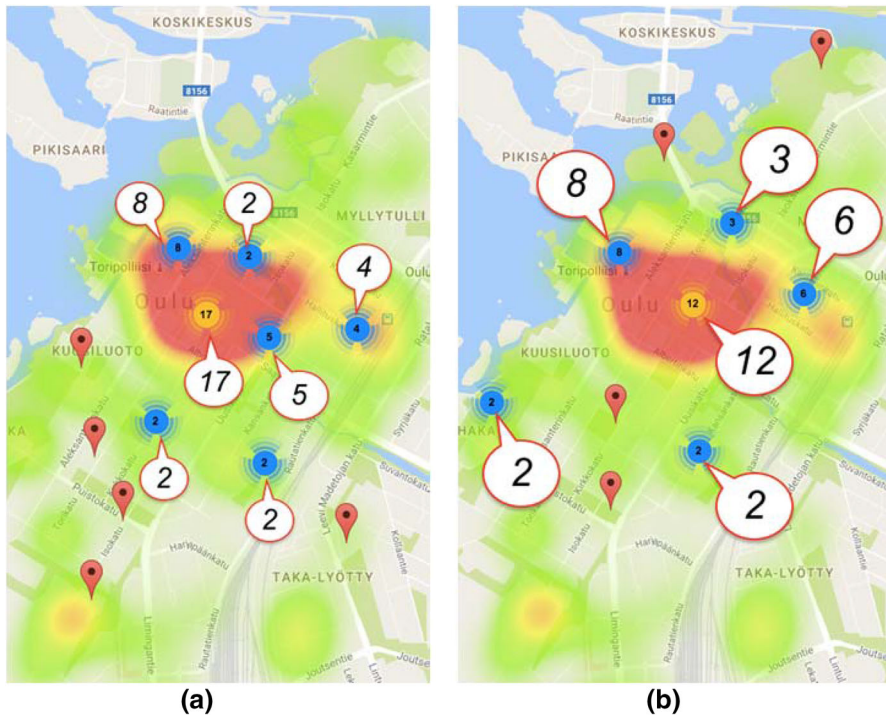
Fig. 14 POI selections during free use with the heatmap displayed (**a**) and hidden (**b**)

users were mostly interested in exploring the heatmap, altering the POI categories visible on the map and seeing POIs where people are currently checked in.

In those sessions where POIs were selected for viewing, 42% of the POIs were clicked while the heatmap was disabled, while 58% were clicked with the heatmap shown. We subsequently visualised the selected POIs over a popularity heatmap of the city, to see whether the presence of the heatmap seemed to affect the spatial distribution of the users' selection. The popularity heatmap is derived by the overall number of checkins in all venues, regardless of time of day. The results are shown in Fig. 14.

We note from this figure that in the case of the heatmap being displayed, 63% of viewed POIs are clustered in the two most popular areas of the city, while in the case of the heatmap being hidden, only 49% are clustered in the central and most popular area of the city. We note also that participants' choices are more geographically restricted, covering an area of 1,201,158.77 m^2 with the heatmap (perimeter = 4390.381 m), while without the heatmap the covered area is 1,316,066.43 m^2 (perimeter = 4975.198 m) This provides some evidence that the heatmap can have an effect on guiding users' choices, though given the nature of the data, we report this conclusion with some caution.

7 Discussion and conclusion

Visitations in a city are an experience that include both the tourist's presence at individual venues (e.g. a museum or a restaurant) and also entire areas (e.g. the market district or historic neighbourhoods). Even though recommendations for specific venues have been the focus of much study via the mining of social network data, the recommendation of areas is a subject in literature that remains unexamined. Our comprehensive review of literature on tourism-oriented system infrastructures including mobile and/or access to urban mobility patterns highlighted various shortcomings of current research, notably in the depiction of spatial context, the consideration of temporal context, the usability and cognitive impact of area context visualisations on tourists, and the duration and limited scope of current studies. These shortcomings form the main motivation of our work and frame the positioning of our contribution to the field.

Based on these, we hypothesized that providing a highlight as a background layer, instead of a foreground layer, informs the user of area importance without obfuscating their understanding of the underlying maps' cartography. More specifically, we attempted to investigate the effect of a heatmap visualization could serve as such a layer for conveying urban social context. The main goal of our paper was to extend the findings from our previous study using pervasive displays, into the use case of smartphones, and to compare between the two. Our results show that users behaved similarly with both a traditional map-and-POI and a heatmap augmented system. We did not find many significant differences in both quantitative and qualitative data on the users' experience with *information seeking tasks* using the two systems. This is a positive outcome, since the potential additional complexity introduced by the heatmap layer did not negatively affect our participants' behavior when using maps.

More importantly, we also found evidence that a heatmap layer that is used as a "recommender" system for areas, compared to highlighting individual points, provides distinct advantages in the provision of social context. Although heatmap users selected on average the same amount of candidate POIs for the tasks, their final choices were more concentrated spatially, which shows that a heatmap can indeed act as an attractor to specific regions of the urban environment. Furthermore, on average, heatmap users tended to select venues of better quality, as indicated by the Foursquare venue popularity of the chosen places. Feedback from our heatmap group participants shows that their perception of urban space is altered through the use of the heatmap, as it affords an understanding of social use under the additional context of time. The heatmap can be seen to support exploration without dictating specific venues as destinations. Users can discover and be attracted by areas, rather than follow goal-driven instructions, which may lead to disappointment (e.g. venues that are fully booked, not busy enough, isolated from other venues) and leave users without suitable alternatives. Further, visitors using the heatmap can get a better feel for the true "vibe" of a city as it evolves throughout the day and visit contextually relevant areas, making their experience more worthwhile. This demonstrates that using such approach is a great tool to market and highlight special city areas on

demand. To some extent, our prototype enhanced the visitor's experience using the widely available mobile phone.

Perhaps though, the most revealing finding from this data is that it provides ample evidence of our hypothesis that users really don't seem interested in individual POIs. Instead users were more interested in exploring the dynamics of areas at different times in the day and further to obtain an idea of the categories of venues in these areas. This is a radical shift from the design of most current travel and location-based mobile and web services, which place a strong focus on the individual selection of single POIs (e.g. TripAdvisor, FourSquare, Google Maps). We feel strongly that designers of such services should take this point into account and provide alternative cartographic visualisations that will support the users' desire to be informed but also to allow for serendipitous discovery and exploration of urban areas, instead of attempting to guide users to particular POIs.

These advantages seem to come at a price. Our users felt greater confidence in their choices when using a heatmap, albeit at a cost of higher mental demand and required effort to achieve these performances. Fortunately, these costs are not substantial as the differences between the traditional and heatmap visualisations are small, and considerably below the mid-point in our measurement scales. Feedback from our participants was also very encouraging and shows that our idea is perceived as helpful and worth pursuing further.

As with all studies, the internal and external validity of the experiments are a concern which merits discussion. In Kjeldskov et al. (2005) it is shown that the usability evaluation of mobile guides benefits differently from a range of different evaluation approaches, from discount usability evaluation, to lab tests and field trials. In this case, we opted for a semi-controlled field experiment, which enhances the ecological validity of our findings, due to the small, but nonetheless reasonable, number of participants, who were by large majority actual tourists, as well as due to the situatedness of the experiment. In this experiment, the increased ecological validity impacts the internal validity of our findings—participants might have been distracted by the environment, or felt pressed for time to complete the experiment and move on to their other activity. Hence we would like to repeat some of this work in a laboratory setting, examining the usability of heatmaps with think-aloud techniques, a process which would allow us to bring out the mental models that describe the users' interaction with the system. Such models would allow us to better design the information visualisation options and repeat field studies. Although we demonstrated the usefulness of a heatmap-type visualization of social context, in the future, we aim to take our concept further by exploring social context sharing through different visualizations. Based on participant feedback and motivated by the issue of cognitive load, we believe a better approach might be found in a visualization that affords a better mapping between real-world definitions of space and social context. As an example, we have started looking at automatically produced choropleth type visualizations (i.e. a map split in geographic segments, which are shaded or patterned in proportion to a statistical variable being displayed on the map. Such maps could that segment urban layouts based on road data from OpenStreetMap, in an effort to better encompass commonly used concepts such as "city blocks" and "neighbourhoods". Figure 15 shows output from preliminary

Fig. 15 The choropleth segmentation algorithm (**a**) paired with social media data from Foursquare (**b**)

work that we have done on a choropleth-generating urban segmentation algorithm, which is paired with social media data to create an environmental social context visualization.

The external validity of our findings could be improved by including a larger population sample with more participants spanning a variety of ages and backgrounds. For our current study, we had to assign realistic, but still artificial tasks to our participants. A future study that involves free use of the application for the duration of a visit would be desirable, although such a study would be hard to orchestrate with real tourists. To do this successfully, we would need tourists to explicitly state their information seeking needs every time the app is used and track their movement during the period of usage, raising significant concerns for privacy and potentially safety of the visitors. However, the analysis of free-use data obtained over a three-year period provides further supporting evidence to the findings of our controlled experiments.

As mentioned previously, our work relies solely on data from the Foursquare API, though our platform does collect data from both Facebook and Google APIs. The Facebook API provides social information such as "Likes" and "Tags" (people tagging posts or images with venues and locations). Using Google Places API we collect information about the user ratings of an establishment or location. However the use of data from multiple sources poses significant inherent challenges. Firstly, physical venues are represented non-uniformly across the various APIs, through variations in naming (e.g. a place called "The Arc" can be represented as "Arc", "Arc Café", "Café Arc") and spatial coordinates, which may differ by several meters (in some cases, up to several hundred). This is understandable, many of these venue representations are generated by users themselves who rely on inaccurate positioning methods or even cases where a place has moved premises. These changes are not homogeneously reflected across all APIs. Another challenge is a scarcity of understanding of social media data correlations (e.g. is a venue with many "Likes" on Facebook likely to attract many "Check-ins" on Foursquare or a high rating on Google Places?) and an understanding of which values that are pertinent to a place are most indicative of popularity. These challenges go beyond the scope of this paper, but definitely open up an interesting and important area of further research that merits following-up.

References

Alghamdi H, Zhu S, El Saddik A (2016) E-tourism: mobile dynamic trip planner. In: 2016 IEEE international symposium on multimedia (ISM). IEEE, pp 185–188

Alves AO, Pereira FC, Biderman A, Ratti C (2009) Place enrichment by mining the web. In: European conference on ambient intelligence. Springer, Berlin, pp 66–77. https://doi.org/10.1007/978-3-642-05408-2_9

Bao J, Zheng Y, Mokbel MF (2012) Location-based and preference-aware recommendation using sparse geo-social networking data. In: Proceedings of the 20th international conference on advances in geographic information systems. ACM, pp 199–208

Bellotti V, Price B, Rasmussen P, Roberts M, Schiano DJ, Walendowski A, Begole B, Chi EH, Ducheneaut N, Fang J, Isaacs E, King T, Newman MW, Partridge K (2008) Activity-based serendipitous recommendations with the Magitti mobile leisure guide. In: Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems—CHI '08, pp 1157–1166

Bergé LP, Serrano M, Perelman G, Dubois E (2014) Exploring smartphone-based interaction with overview+ detail interfaces on 3D public displays. In: Proceedings of the 16th international conference on human–computer interaction with mobile devices and services. ACM, pp 125–134

Braunhofer M, Ricci F (2017) Selective contextual information acquisition in travel recommender systems. Inf Technol Tour 17(1):5–29

Braunhofer M, Elahi M, Ricci F (2014) Techniques for cold-starting context-aware mobile recommender systems for tourism. Intell Artif 8(2):129–143

Calabrese F, di Lorenzo G (2011) Estimating origin—destination flows using mobile phone location data. Cell 10(2011):36–44

Calabrese F, Reades J, Ratti C (2010) Eigenplaces: segmenting space through digital signatures. Pervasive Comput 9(2010):78–84

Carter S, Marlow J, Cooper M (2017) No app needed: enabling mobile phone communication with a tourist kiosk using cameras and screens. In: Proceedings of the 2017 ACM international joint conference on pervasive and ubiquitous computing and proceedings of the 2017 ACM international symposium on wearable computers. ACM, pp 221–224

Cheverst K, Davies N, Mitchell K, Friday A, Efstratiou C (2000) Developing a context-aware electronic tourist guide: some issues and experiences. In: Proceedings of the SIGCHI conference on human factors in computing systems—CHI '00, pp 17–24

Churchill EF, Nelson L, Denoue L, Girgensohn A (2003) The plasma poster network: posting multimedia content in public places. Interfaces 1(2003):599–606

Colomo-Palacios R, García-Peñalvo FJ, Stantchev V, Misra S (2017) Towards a social and context-aware mobile recommendation system for tourism. Pervasive Mobile Comput 38:505–515

Crandall DJ, Backstrom L, Huttenlocher D, Kleinberg J (2009) Mapping the world's photos. In: Proceedings of the 18th international conference on world wide web WWW, p 761

Fisher D (2007) Hotmap: looking at geographic attention. IEEE Trans Graph 13(6):1184–1191

Girardin F, Calabrese F, Fiore FD, Ratti C, Blat J (2011) Digital footprinting: uncovering tourists with user-accessed digital footprinting: uncovering tourists with user-generated content. Pervasive Comput 7(2011):36–43

Grubert J, Pahud M, Grasset R, Schmalstieg D, Seichter H (2015) The utility of magic lens interfaces on handheld devices for touristic map navigation. Pervasive Mobile Comput 18:88–103

Herzog D (2017) Recommending a sequence of points of interest to a group of users in a mobile context. In: Proceedings of the eleventh ACM conference on recommender systems. ACM, pp 402–406

Hinrichs U, Carpendale S, Valkanova N, Kuikkaniemi K, Jacucci G, Vande Moere A (2013) Interactive public displays. IEEE Comput Graph Appl 33(2):25–27

Jiang B, Yin J, Zhao S (2008) Characterizing human mobility patterns in a large street network. Phys Rev E 80:17

Kim M, Kotz D (2005) Modeling users' mobility among WiFi access points. In: WiTMeMo 05 papers presented at the 2005 workshop on wireless traffic measurements and modeling, pp 19–24

Kjeldskov J, Graham C, Pedell S, Vetere F, Howard S, Balbo S, Davies J (2005) Evaluating the usability of a mobile guide: the influence of location, participants and resources. Behav Inf Technol 24:51–65

Koch M, Kötteritzsch A, Fietkau J (2017) Information radiators: using large screens and small devices to support awareness in urban space. In: Proceedings of the international conference on web intelligence. ACM, pp 1080–1084

Kostakos V, O'Neill E, Penn A, Roussos G, Papadongonas D (2010) Brief encounters: sensing, modeling and visualizing urban mobility and copresence networks. ACM Trans Comput Hum Interact 17(2010):1–38

Linden T, Heikkinen T, Kostakos V, Ferreira D, Ojala T (2012) Towards multi-application public interactive displays. PerDis 2012(2012):1–9

Lindqvist J, Cranshaw J, Wiese J, Hong J, Zimmerman J (2011) I'm the mayor of my house: examining why people use foursquare-a social-driven location sharing application. In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM, pp 2409–2418

Majid A, Chen L, Chen G, Mirza HT, Hussain I, Woodward J (2013) A context-aware personalized travel recommendation system based on geotagged social media data mining. Int J Geogr Inf Sci 27(4):662–684

McCarthy JF, Costa TJ, Liongosari ES (2001) UniCast, OutCast and GroupCast: three steps toward ubiquitous, peripheral displays. Proc UbiComp 2001:332–345

Meehan K, Lunney T, Curran K, McCaughey A (2013) Context-aware intelligent recommendation system for tourism. In: 2013 IEEE international conference on pervasive computing and communications workshops (PERCOM workshops). IEEE, pp 328–331

Noulas A, Scellato S, Lambiotte R, Pontil M, Mascolo C (2012) A tale of many cities: universal patterns in human urban mobility. PLoS One 7(5):e37027

O'Hare GM, O'Grady MJ (2003) Gulliver's Genie: a multi-agent system for ubiquitous and intelligent content delivery. Comput Commun 26(11):1177–1187

Poslad S, Laamanen H, Malaka R, Nick A, Buckle P, Zipl A (2001) CRUMPET: creation of user-friendly mobile services personalised for tourism, pp 28–32.

Quercia D, di Lorenzo G, Calabrese F, Ratti C (2011) Mobile phones and outdoor advertising: measurable advertising. IEEE Pervasive Comput 10(2011):28–36

Rattenbury T, Good N, Naaman M (2007) Towards automatic extraction of event and place semantics from flickr tags. In: Proceedings of the 30th annual international ACM SIGIR conference on research and development in information retrieval SIGIR 07, p 103

Savage NS, Baranski M, Chavez NE, Höllerer T (2011) I'm feeling LoCo: a location based context aware recommendation system. In: Proceedings of the 8th international symposium on location-based services (LBS'11)

Schwinger W, Grün C, Pröll B, Retschitzegger W, Schauerhuber A (2005) Context-awareness in mobile tourism guides—a comprehensive survey. Rapport technique. Johannes Kepler University Linz, Linz

Souffriau W, Vansteenwegen P, Vertommen J, Berghe GV, Oudheusden DV (2008) A personalized tourist trip design algorithm for mobile tourist guides. Appl Artif Intell 22(10):964–985. https://doi.org/10.1080/08839510802379626

Tammet T, Luberg A, Priit J (2013) Sightsmap: crowd-sourced popularity of world places. Inf Commun Technol Tour 2013:314–325

Umanets A, Ferreira A, Leite N (2014) GuideMe—a tourist guide with a recommender system and social interaction. Proc Technol 17:407–414

Valls JF, Sureda J, Valls-Tuñon G (2014) Attractiveness analysis of European tourist cities. J Travel Tour Mark 31(2):178–194

van Setten M, Pokraev S, Koolwaaij J (2004) Context-aware recommendations in the mobile tourist application COMPASS. Adaptive hypermedia and adaptive web-based systems. Springer, Berlin

Komninos A, Besharat J, Ferreira D, Garofalakis J (2013a) HotCity: enhancing ubiquitous maps with social context heatmaps. In Proceedings of the 12th International Conference on Mobile and Ubiquitous Multimedia. ACM, Lulea, p 52. https://doi.org/10.1145/2541831.2543694

Komninos A, Stefanis V, Plessas A, Besharat J (2013b) Capturing urban dynamics with scarce check-in data. IEEE Pervasive Comput 12(4):20–28

Yang D, Zhang D, Yu Z, Yu Z (2013) Fine-grained preference-aware location search leveraging crowdsourced digital footprints from LBSNs. In: Proceedings of the 2013 ACM international joint conference on pervasive and ubiquitous computing. ACM, pp 479–488