The Anatomy of Mobile Location-Based Recommender Systems

Neal Lathia

Abstract

The widespread adoption of smartphones is now putting both the Internet and sensor-rich hardware into the pockets of millions. While recommender systems have become the norm on many web sites, many mobile systems have historically been built as *location-based* services. However, these devices are becoming the ideal interface for *recommender* systems that help users discover, explore, and learn about their physical surroundings. In this chapter, we review the main components of a mobile location-based recommender system: the data that can be used to learn about users and items, the algorithms that have been applied to recommending venues, and the techniques that researchers have used to evaluate the quality of these recommendations, using research that is sourced from a variety of fields. This chapter closes by highlighting a number of opportunities and open challenges related to building future mobile recommender systems.

1 Introduction

The widespread adoption of smartphones–putting both the Internet and sensor-rich hardware into the pockets of millions– is finally bridging the gap between the online and offline worlds. It is now common for mobile phone users to search the web and engage with social media while on the move: the services that were once limited to the desktop computer are now at their fingertips. Furthermore, the vast information repository on the web can now be used to enhance peoples' physical-world experiences. Mobile phones are quickly turning away from being mere portals to the

Neal Lathia
Computer Laboratory, University of Cambridge, Cambridge CB3 0FD, UK, e-mail: neal.lathia@cl.cam.ac.uk

web and towards devices that help users to explore, discover, and interact with their actual surroundings.

One of the key technologies that is enjoying much success in the online world is usage of *recommender systems* to support users' browsing. Online recommender systems take many different shapes: they help users discover movies, music, and e-commerce items of interest, as well as suggesting new friends to connect to in online social networks and providing personalised search results. At the heart of their success is the assumption that a model of users' preferences can be learned by observing their behaviour (expressed as, for example, star-ratings or clicks); huge repositories of data can then be filtered in order to draw out the most interesting results for each person. Mobile phones, instead, have historically been centred on *location*-based services: the underlying paradigm is that the most relevant information for users is about that which is close by. However, the next generation of mobile phones now offer the potential to implement recommender systems to build services that not only leverage users' current location, but also their rich history of preferences and actions. In doing so, a crossroads of multiple lines of research, each with their own rich literature, is being formed. People's usage of mobile systems is of interest to a wide range of fields within Computer Science, ranging from mobile information retrieval [17], sensor research [22], data mining and knowledge discovery [14], human-computer interaction [49], as well as persuasive [26] and ubiquitous computing [47].

This chapter aims to draw together the various lines of enquiry related to location-based personalisation and mobile recommender systems by presenting a structured survey of the key elements of a mobile location-based recommender system. We do so from the point of view of the recommender system itself, beginning with a broad definition of mobile recommender systems. We then cover three features of mobile recommender systems:

1. **Data**. Recording signals of behaviour that reflect users' preferences is the foundation for any mechanism that aims to recommend new places, activities, or friends. In this regard, a growing body of research has delved into collecting data from users about themselves and their surroundings, via participatory sensing, crowd-sourcing, and game-based incentives. Furthermore, a range of research has investigated how to infer users' activities from such data.
2. **Algorithms**. The principal technique behind recommender systems is *collaborative filtering*. While these are readily applied to mobile systems as well, these algorithms have historically taken a "black-box" approach when computing on user ratings: they do not, for example, need to consider the physical distance between places. We will therefore also discuss supervised learning approaches for mobile recommendations and recent research that augments the efficacy of recommendations by taking into account features relating to space (e.g., where people live) as well as preference.
3. **System Evaluation**. The question of evaluating recommender systems is still actively discussed [35]. We complement this research by surveying how mobile recommender systems have been evaluated to date, and how their evaluation differs from more traditional scenarios.

We conclude our survey by discussing emergent themes and set of directions for future research.

1.1 Defining a Mobile Location-Based Recommender System

We begin with a broad definition of the kinds of systems that we describe in this chapter. To date, many *recommender systems* and *location-based systems* have been built and studied as separate entities. Broadly speaking, they can be defined as follows:

- **Recommender Systems** retrieve tailored sets of *items* of interest for each user. A variety of flavours of recommender systems are discussed throughout this handbook: for example, see the chapter on "Data Mining Methods for Recommender Systems" for approaches that recommend based on users' historical preferences (**??**), and "Semantics-Aware Content-Based Recommender Systems" for approaches that recommend based on items' features (**??**).
- **Location-based Systems** or services retrieve information that is tailored to the user's *current location* [72]. Typical applications here include mapping and route-finding services, applications to find nearby services (e.g., restaurants), location-based social networks where friends share their location with one another, traffic notification services, and advertising. The focus is heavily on location, rather than preference.

Historically, the recommender system literature has been characterised by a focus on recommender systems that users interact with using a personal computer, and recommending items that are potentially not 'consumed' immediately after being recommended, such as movies, music, and the contents of e-commerce catalogues. Although these recommendations may often result in real-world interactions (e.g., a movie being sent to your house), they are nevertheless mostly finding content based on what people like. In other words, any spatial relationships between the items (e.g., where a restaurant is relative to another) are not useful when computing recommendations: the focus is on identifying, via a range of machine learning approaches, implicit relationships between items using the feedback or preferences given by the system's users.

The systems that we focus on here, *mobile location-based recommender systems*, take on characteristics of both of the above: they are accessed via mobile devices, use location data (current or otherwise, e.g. historical), involves and leverages users' movement around a physical space and, most importantly, provide *personalised* recommendations that are tailored to users' preferences. To that end, we exclude systems that do not recommend places (or venues; 'items' that are consumed by visiting a specific geographic location), such as when users access their movie recommender (e.g., Netflix account) via a mobile device [34, 52], or seek personalised app-recommendations with their mobile [39]. In that regard, mobile location-based recommender systems may be viewed as a particular kind of *context-aware* recommender system [2], where spatio-temporal data (about *where* and *when* the system is being used) can be used to further personalise results.

In light of the above, what are the tasks that users of mobile location-based recommender systems are seeking to perform?

1. **Goal-Oriented Search**: Location-based recommender systems often allow users to query for personalised results. Where is the closest restaurants that I would like to have dinner at? Where are nearby shopping areas? These tasks are often associated with a particular intended action (e.g., having dinner, going to a bar), yet with results that can be personalised to each individual.
2. **Location Discovery**: While the above use case captures when users have queries/intents, mobile location-based recommenders can also be used to discover places. What is around me, of interest? What should I see in London? What places are trending nearby, or events happening in my neighbourhood? All these kinds of questions fuel use cases where the user's historical profile can be used to personalise recommended places to see, visit, or attend.
3. **Routing and Transport**. Finally, a number of use cases have appeared in the literature that deal with recommending personalised *routes* to follow. While mainly focused on tourist routes, this use case responds to: how should *I* get from here to there? What route should I walk when I am visiting Barcelona, with my children? And so forth.

Beyond these, there are location-based social matching applications, tailored to find people of interest in particular locations, and behaviour-oriented applications, such as those related to sport and physical activity. While these are potentially amenable to personalisation and recommendation systems, this chapter focuses on those applications that are related to venues and places. The following published surveys review mobile recommender systems more broadly: [29, 42, 68].

2 Data for Mobile Recommender Systems

One of the key differences between mobile- and web-based recommender systems is that the former tend to have access to a broader set of data than the latter. Traditionally, web-based recommender systems' data is described as being either explicit (e.g., a rating or similar value derived from a user's evaluation of an item) or implicit (e.g., a purchase or click; a value derived from the user's behaviour). Mobile systems can also collect these, and more. While recent systems have particularly focused on location and mobility data, mobile systems can collect:

1. **Explicit Data**: Mobile users can, as they do on the web, rate, tag, share, 'like,' or otherwise score an item while on the go. Beyond this, the most prominent explicit action that has emerged across mobile services (e.g., Foursquare, Facebook Places, Google+, Yelp) is the *check-in*: users share their current location with their friends by finding and selecting the venue they are in. In the following, we describe how these relate to preference.

2. **Implicit Data**: As above, users may provide similar implicit data as they do on the web by clicking links, streaming videos, making purchases, or otherwise engaging in an action that is not limited to mobile only. However, the mobile device does provide some differences: there are behaviours (e.g., taking a photo, tracking physical exercise) that are exclusive to mobile devices.
3. **Sensor Data**: Modern smartphones are increasingly sensor-rich. These typically include sensors that can measure location and mobility, co-location, and other facets that describe users' context [47].

Moreover, many mobile systems inherently collect multiple kinds of data at once. Consider, for example a check-in on a location-based social network (e.g., Foursquare). When checking in, users' sensor data is used to identify nearby venues [74]: their check-in action is an explicit signal of presence at that venue; multiple check-ins at the same venue may be considered as an implicit confidence metric of preference for that venue [36], and the timestamp of their check-in implicitly uncovers features of the place where they are [8].

One of the most notable differences between web- and mobile-based recommender system data is that the *item set* that the recommender uses is often dynamic; to follow from the example above, not all possible venues that a user may like to check-in to may be known to the system. A key facet of building a mobile recommender therefore is using the available data in order to learn about *both* items and users. In the following, we describe a number of examples from the literature where data derived from mobile devices is used to build databases that could suitably underpin a recommender system. In particular, we focus on finding and inferring points of interest, learning and modelling mobility data, analysing check-ins, and inferring context and activities from sensors.

2.1 Uncovering Points of Interest and Location Preferences

As mobile devices are used on the go, they become an ideal source of location data. In this section, we describe how this data can be used to learn about both *users* and *items* in a recommender system, and the relations between them.

Location-based recommender systems rely on having a database of Points-of-Interest (POIs) from which to source recommendations. The recent literature has described a number of means of finding and inferring POIs from users' data. A number of systems (e.g., Foursquare) maintain their POI database via crow sourcing; the explicit check-ins that users provide can then be used to uncover venues' spatio-temporal patterns [57]. Others, instead, infer them from implicit data. These include, for example, sourcing POIs by clustering geo-tagged photographs that users upload to services like Flickr [20] [51]. These datasets can be used to automatically extract features of places and events [66], and have also been applied to image search result diversification [40]. Further information about the inferred items can be gathered by intersecting the location data with any available content and tags [41].

While the methods above can be used to populate information about items, geo-tagged photos have also been used to make inferences about users' behaviours. These include identify trips [60], analysing how tourists navigate a city [31], and predicting how people travel [18]. In essence, the seemingly meaningless act of taking a photograph from a geo-aware device can be used to find that (a) the targets of photographs are items that are of interest to users and (b) the users who took those photos are interested in, and have travelled between, those targets that they have captured.

Mobility data has more traditionally been sourced from mobile phones. These include the phones' Global Position System (GPS) [88], sensors, GSM traces [76], as well as the Call Detail Records (CDRs) that are created when devices pair with cellular network communication towers [9]. A full review of the literature analysing these data sources is beyond the scope of this chapter [11] [33]. However, these sources of data uncover a vast range of features about users' behaviours, including how far they tend to travel, their likely mode of travel, and the urban areas they frequent [67]. Clearly, these sources of data encode users' daily routines: the open question, related to recommender systems, is the extent that these also signal users' tastes. Few historical studies shed some light on this issue. Froehlich *et al.* [25] found that mobility patterns correlate with users' preferences: people tend to frequent those places that they like; similarity between users can also be measured from location histories [48]. However, other studies uncover that between 50-70% of users' mobility captures routine behaviours [14]. The fact that such a large proportion of the user data contains places that users will, by definition, be very familiar with challenges the perspective of building recommender systems to facilitate *discovery* of new places. Yet this kind of data has been used to design location-based social activity recommendations [62]. Moreover, GPS traces can be mined for 'interesting' locations [89] in order to recommend locations and activities [86]; further details of the algorithmic approaches appear in the following section.

All of the above data sources share the common trait of requiring processing prior to being used as signals of users' mobility and/or preference. They differ from one another, instead, in how easily and accurately they may be collected. Typically, sources such as GSM and CDRs are only available to mobile operators; GPS and similar on-board location services require a tailor made app-based data collector. While the former kind of data is typically coarse-grained, and GPS can provide much finer-grained samples (both spatially and temporally), fully efficient implementations are dependent on the needs on the underlying application. In particular, continuously querying a phone's GPS sensor will quickly degrade the device's battery: system designers need to trade-off between the sampling accuracy that they seek and the energy efficiency of their application [64]. On the other hand, many applications collect data explicitly from their users, such as via location check-ins. These sorts of systems surface a variety of issues that reflect on data quality, such as the incentives and reasons that users have for contributing at all. Lindqvist *et. al* [49] explored a host of reasons why people participate in these location-based services (often, at the expense of their own privacy). These include: personal tracking, gaming, and social signalling with friends; moreover, they also uncover that there are many places that people proactively chose to *do not* check-in to.

2.2 Behavioural Inferences from Smartphone Sensors

While the previous section mainly dealt with what systems can learn about users and items from mobility data (including smartphone sensors, like GPS), there is a growing field of research that focuses on further behavioural aspects that can be inferred from sensors [43]. To date, these have not been widely applied to recommender systems. We include this brief review here since (a) these inferences provide insights into users, and is thus relevant to user modelling using smartphones, and (b) to highlight future opportunities that may emerge from applying these inferences to personalised systems.

1. **Activity Recognition**. Data from smartphone sensors (e.g., the accelerometer) has been used to monitor and detect users' current activities. These include whether the user is walking, sitting, driving, or talking [16]. Moreover, smartphone sensors have been used to detect users' contexts, including whether they are in an environment where music is being played [50].
2. **Transportation Modes**. Combinations of accelerometer and GPS data have been used to infer how users are moving between places, detecting transportation modes such as bicycles, cars, buses, or subways [77]. These kinds of inferences have, more broadly, been used to monitor users' 'green' behaviours [26], indicating how inferences from sensor states can be used to profile users' behaviours.
3. **Sociability**. Smartphone sensors have also been used to detect users' social networks and interactions [24]. A mixture of Bluetooth, accelerometers, and microphone sensors has been applied to detect users' colocations and interactions [65], both to quantify those users who are more sociable and provide feedback to users. Other work takes similar data into the domain of recommendation, by recommending online contacts based on physically sensed colocations [61].

There are a number of challenges related to the above, which include collecting data *efficiently*, without overly draining devices' batteries, designing *accurate* inference algorithms in order to infer the higher-level behaviours that are relevant to the user-modelling task at hand. However, these methods promise to deliver highly granular data about users: where they go, whom they interact with, their activities and routines, and more: just as locations reflect preference, future mobile recommender systems may use sensor inferences to augment user profiles.

3 Computing Recommendations in Mobile Applications

In this section, we describe approaches that have been proposed in order to compute recommendations for mobile users. In particular, we focus on how the problem of generating recommendations related to venues (e.g., restaurants, shops) has been formulated into well-defined machine learning problems, that can then be tackled by learning from the kinds of data described in the previous section.

We begin by briefly reviewing the equivalent in traditional recommender systems. In general, a recommender system will have a set of *items* and *users*; any given user may have rated a fraction of the items (or performed equivalent actions, if the system deals with implicit data: we use the term ‘rating’ to generically mean a preference value). The task of the system is to recommend, to each user, those items that he/she will be interested in—perhaps with a number of constraints. To do so, the system computes personalised *predictions* for those items that a user has not rated. These predictions can then be used to *rank* items according to estimated preference; the user is presented with a list of items ordered according to how interested the system has forecasted that user will be in them. Broadly, therefore, the two main approaches to web recommendation focus on rating prediction and item ranking [21] (**??**). In mobile recommender systems, some of these principles continue to apply: in the following, we review variants of them that have been tailored to particular mobile recommendation scenarios.

These variants have emerged for two reasons: first, those tasks that users seek to accomplish in mobile settings often differ from what people do on the web, and it is questionable as to whether the problem of information overload is applicable at all. Further, there are a variety of challenges related to applying machine learning to tasks related to mobile scenarios. These include limitations that are a result of the data itself (e.g., inferring preference from mobility, differentiating between positive and negative experiences from implicit datasets), as well as our current understanding of the limits to the predictability of any data that can be collected [33, 14]. Finally, there are also differences in the users themselves, who may be locals or tourists and may be interested in geographical regions of varying size.

3.1 Overview of Recommendation Formulations

In this section, we examine how the problem of recommending places to mobile users has been defined as a formal prediction problem. In particular, we consider four variants of the broad problem: (1) recommending venues of particular categories, (2) recommending the *next* place that a user may like to visit, (3) recommending *new* places that users have yet to visit, and (4) recommending *routes* that users may like to take as they navigate a particular space. While each of these can generally be considered as place-focused recommendation problems, they each capture differences in users’ needs from a mobile recommender.

1. **Categorical Recommendation**. The setting that is likely to mirror 'traditional' single-category online recommendation is that of recommending venues of a particular type (e.g., restaurants). Systems described in the literature focus on shops [78, 82] restaurants [69, 81], and cultural/tourist travel [5, 83]. Much like movie recommendation (**??**), the items in this setting tend to all be the same—the task is therefore to rank them appropriately, perhaps with the only added constraint of being within a particular radius of the user's current location. In [69], items are described with n-dimensional vectors that include further attributes of each venue (e.g., average cost); this way of representing the data allows for the system which takes a conversational approach to recommending the best restaurant.
2. **Predicting the Next Place**. Let us assume that a user having dinner at a restaurant; she now would like a recommendation for bars or clubs to go to once she has finished eating. This is an example of seeking a recommendation for the *next* place to visit: the relevant inputs to this query are (a) the user and her preferences/location history, (b) the current location of the user, and (c) the current time of day. Formally, let us represent a user's location history as a time series of venues that end at the current venue V_n at time t_n:

$$P_u = \big((V_0,t_0),(V_1,t_1),...,(V_{n-1},t_{n-1}),(V_n,t_n)\big) \tag{1}$$

   Given a set of candidate venues L, the prediction task is thus to predict which venue V_{n+1} the user should visit. More broadly, the goal is to *rank* venues such that the venue V_{n+1} that the user would like to visit next is placed as highly as possible within the recommendation list [56]. To do so, a ranking score $\hat{r}_{u,v}$ is computed for every venue v in $(L \setminus \{V_n\})$ (i.e., all the venues except the one where the user currently is) using features from all users' location histories (as described below).

$$\hat{r}_{u,v} = P(v = V_{n+1}|u,V_n) \tag{2}$$

   This problem has been tackled in the literature using both Foursquare check-in data [56] and GPS and WiFi log data (although not from the perspective of recommendation) [71, 53]. Successfully predicting next places with these datasets, however, highlights one of the open challenges of this method when applied to a recommendation scenario: part of the success may be attributed to the habitual or otherwise routine mobility that is captured in the data [33]; in essence, predicting that a user will go from home to work and back again seems to have little value from the perspective of a recommender system. To tackle this shortcoming, researchers have narrowed the scope of what venues in L are candidate for recommendation: the following section focuses on one subset of these.
3. **Predicting New Venues**. Since recommender systems are often described as tools to facilitate discovery, another problem that mobile systems may tackle is that of predicting the *previously unvisited* venues that a user may like to go to. This problem has been formally defined by Noulas et. al [55] as follows. Given the set of venues U that a user u has historically visited over a period of

time $(t-\Delta,t)$, the aim is to predict those venues in $(L \setminus U)$ that the user may like to visit in time $(t,t+\Delta)$. The choice of time period Δ is a parameter that determines the extent that this approach may be amenable to venue *rediscovery*, e.g. predicting venues that the user has not been to yesterday, last week, last month, or ever at all.

Like above, this approach has its own shortcomings. Most notably, this approach can only predict and recommend novel locations that the system already knows about (i.e., that are available in the training data set), which is a particular instance of the cold-start problem.

4. **Recommending Routes to Follow**. Research that has focused on recommending to tourists often deals with personalised routes that this kind of user may follow as they explore a new area [15]; tourists' digital footprints can be directly uncovered from the photographs they take while on tour [30, 32], which can then be used to construct personalised tours [12, 1]. The idea here is somewhat akin to recommending a playlist of tracks [6], albeit with added geographic constraints: the formal task is to compute a sorted list of places to visit that optimises against both the user's preference, the time to travel between stop points, and any other contextual factors.

   In essence, this setting may be viewed as an instance of the 'next place' problem, where the task is to recommend the next N places based on a number of constraints. For example, the system in [73] considers the time since a user has visited a place of a particular category, in order to diversify results. Similarly, the system in [13] also considers venue opening/closing times, the routes between places, and the 'best' times to visit particular venues. Finally, the system in [4] also considers the kinds of tourists groups that are requesting recommendations–including, for example, whether the group includes children.

We note that a number of other variant prediction problems that are relevant to mobile recommendations exist; the above are a selection that have appeared in the recent literature. These include, for example, discovering new events [75].

3.2 Algorithmic Approaches to Venue Recommendation

In this section, we review some of the algorithmic approaches that have been adopted for mobile recommendation. Many of these leverage the principles underlying collaborative filtering [87, 28]: a full review of collaborative filtering is beyond the scope of this chapter (**??**). Broadly speaking, when users can be represented as vectors of the venues that they have a preference for, and venues ('items') can be represented as vectors of the users who have a preference for them, the entire family of collaborative filtering approaches can be applied.

A particular characteristic of location-based recommender systems is that recommendation results may need to be pre- and/or post-filtered in order to localise the results to a particular geographic area [2]. These approaches are, more generally, typically applied in *context-aware* recommender systems: in the location-based domain, this may, for example, entail pre-filtering by only training on those ratings that match the current target one and/or post-filtering by removing some of the ranked items (e.g., "only show me recommendations within a 5 kilometre radius").

We begin by describing baseline approaches that may be suitable to compare any recommendation algorithm against. The include:

- **Popularity**. Although non-personalised, popularity is a strong baseline to consider when recommending venues. Popularity may be defined in a number of ways: geographically, by absolute number of visitors, by visitors' frequency of visits, or by category. While this approach does not personalise results, it captures the fact that popular venues are–by definition–places that many people will like to go to. A personalised variant of popularity could, for example, rank places based on a user's historical patterns (e.g., ranking coffee shops highly if a user tends to visit this category often).
- **Proximity**. Since the 'items' in venue recommender systems have an inherent geographical layout, another baseline to compare against is that of simply recommending venues by geographical distance from the user's current position [62]. This baseline does not consider preference, historical mobility, or any other contextual factors–yet captures users' tendency to travel over short distances [54].

An approach that has emerged in the recent literature [56, 55] revolves around extracting features from the data, creating binary-labelled datasets, and applying supervised learning in order to learn the likelihood that a user visit a particular venue. There are three kinds of features that can be extracted from mobility preference data. These include:

- **Place Features**. Beyond any categorical/attribute data that is available for venues, mobility data can be used to infer aspects of places that are, more broadly, related to the behaviours of those people who attend them. These include (a) the overall popularity of the venue, (b) the popularity of that venue at a particular time of day, or day of week, (c) the popularity of the venue within its particular category or geographic space. Popularity can be defined both in terms of absolute visits or the unique number of users who have visited a place.
- **User Features**. As above, beyond any attribute data available for a user, the mobility data can expose a number of features about preference for venues. These include (a) the frequency or proportion of times that the user has historically visited a place, (b) the user's prior likelihood of visiting a place of a particular category, and (c) the distance of a venue from the geographic centroid of a users' historical mobility. If a social network is also available, similar features can be extracted for a user's friends for each venue–capturing the importance of friend's mobility in determining how users navigate places [23].
- **Structural Features**. Finally, mobile data about users and places also inherently encodes a number of structural properties that are a result of both places and users combined. These include geographic features: the distance between places, and the *rank* distance between neighbouring venues. A sizeable amount

of data about users' mobility allows for features relating to transition probabilities: what is the likelihood that a person go from venue A to venue B, or from category A to category B? The benefit of these features is that they are not solely based on geography; they uncover features that relate places without needing to know about the spatial layout of the items.

Using any of the available features described above, each visit by a user to a venue can be turned into a positively labelled *instance* that can be used to train any supervised learning approach (e.g. linear regression and decision trees [56]). However, doing so using only positively labelled instances will lead to poor results, as the training data is highly skewed. To overcome this, researchers have augmented their training data by randomly selecting unvisited venues to construct negatively labelled instances. This approach effectively reduces the problem of ranking into one of training a regression model with binary data [19]; learning on extracted features seeks to determine what aspects of venues attract users to them–to then be able to compute ranking scores for other venues that can be provided as recommendations.

A second approach that has been recently applied to recommending venues is by using *random walks* (often, with restart [80]), which is well-known in the context of web search [58]. This approach is suitable for a dataset that can be represented as a graph; broadly, the algorithm begins at node i and moves across the graph's nodes with particular transition probabilities: eventually, the steady-state of this walk is reached, and defines the probability of being at any particular node j, or, put another way, the relevance of j with respect to i. In the case of venue recommendation, the graph we have at hand contains nodes of users and venues. Links between users and venues define the preference (e.g., historical visits) of a user to a venue, and weights on those links are the transition probabilities. If we have a social network, there are also links between users; this approach has been used both the recommend places [55] as well as recommend links to be added to the social network [7]. While powerful, this method suffers from the perspective of scalability: for example, in [55], a separate random walk was computed for each user.

4 Evaluating Mobile Recommendations

A critical step of all recommender system research is applying a methodology to evaluate the quality of the recommendations [35]. Mobile recommender systems are no exception; in fact, many of the techniques that have been applied to evaluate recommendation quality can be similarly applied to this domain. For a full review of recommender system evaluation, please refer to the relevant chapter in this handbook (**??**). Broadly speaking, just like in web settings, mobile recommendation evaluations can be conducted using *quantitative* and *qualitative* methods.

Quantitative methods mirror precisely what is traditionally done with web data: data sets are split into appropriate training and test sets, and the predictive power of learning algorithms is measured, after they have been given the training set, on the hidden test set. However, while many web experiments focus on prediction accuracy, since mobile data is often unary (i.e., a check-in) or implicit (e.g., from location traces), then *ranking* metrics are more often appropriate. For example, in [62] the percentile-ranking metric is used to evaluate the quality of recommended events, a metric that was previously used with implicit data [36]. In this case, a successful recommendation would highly rank those events that users subsequently attended. It therefore defines $gone_{u,j}$ as a binary flag that reflects whether user u attended event j, and $rank_{u,j}$ as the normalised rank of the event j in u's recommendations. The percentile-rank is defined as:

$$\overline{rank} = \frac{\sum_{u,j} gone_{u,j} \times rank_{u,j}}{\sum_{u,j} gone_{u,j}} \tag{3}$$

A number of recent studies have also provided *qualitative* evaluations of their systems [13]. Much like their web-based equivalents, these studies entail building a system, recruiting participants, and evaluating the recommendations using surveys, interviews, or similar methods. While offering similar benefits, such as a finer grained understanding of user experience, they do also tend to suffer from similar drawbacks; for example, they often face cold-start settings and are relatively small-scale. For example, Tintarev *et. al* [79] evaluated a mobile tourist recommender by having recruited participants complete a questionnaire. This questionnaire was used to generate personalised points of interest, and the system was then evaluated by examining the number of venues visited, as well as their popularity and novelty. Similarly, the Magitti system [10] was evaluated in the field, allowing the researchers to understand concepts such as omissions, distance to recommended places, and the transparency/explainability of the recommendations.

All of the studies above indicate that evaluating a mobile recommender system begins by evaluating the recommender system as it would be evaluated on the web. However, limiting studies to these evaluations alone will not expose the complex mesh of values that users seek in a successful recommendation, including aspects that are also applicable to the web (novelty, diversity, explainability) as well as aspects that are unique to mobile settings (distance, time of day, geographic representativity, venue opening hours, etc.).

5 Conclusions and Future Directions

In this chapter, we have reviewed the basic components of mobile location-based recommender systems: the tasks that these systems seek to support, the (explicit, implicit, or sensor) data that can be used to build them, how these kinds of recommendations are defined as formal prediction problems, the algorithms that have been applied to them in the recent literature, and how these systems are both quantitatively and qualitatively evaluated. In doing so, a number of themes have emerged; a number of open challenges remain as we look forward to the future research in this domain. We close this chapter by describing a number of these challenges:

1. **Context**. Mobile recommender systems are, arguably, even more tied to users' current context than their web-based equivalents: those venues that people seek to discover will be highly dependent on (beyond their preferences) where they are, the time of day, who they are with, and perhaps even how they feel. While the concept of context is emergent in the recommender system literature [2, 3], fitting it appropriately into mobile recommender systems requires revisiting how context can be defined, collected, and applied to this domain.
2. **Hierarchical Item Sets**. In traditional recommender systems, 'items' are well-defined entities (books, e-commerce items) that often do not overlap, and may be 'dynamic' in terms, for example, their stockroom availability [37]. In mobile recommender systems, 'items' are dynamic in that they may be venues that are open, closed, or permanently moved; they may be events that have varying temporal qualities (e.g., a theatre production that lasts for one month vs. a rock concert that only happens on one night); or indeed they may have varying geographic spans (such as a venue vs. a neighbourhood [85]). In essence, the items in mobile recommenders are strongly structured and relate to one another both hierarchically and spatio-temporally. One problem that emerges here is that historical mobility-preference data detracts from these system's ability to recommend upcoming events of interest, that will have no associated data. Future work can explore how these dynamics may be learned or detected, and, perhaps more importantly, how to appropriately structure a recommender system that balances between distance and preference: should such a system recommend that the user travel to somewhere distant in exchange for a high preference match, or recommend somewhere nearby that does not fully fit their profile?
3. **Privacy**. All of the potential that mobile recommender systems uncover seems to conflict with users' privacy: the data that we have described above includes instances of both users' selective exposure of their location as well as passive location tracking. Future systems may consider including obfuscation mechanisms that re-introduce certain levels of privacy into the collected data [63]: more work is required to understand how this would impact users' recommendations, and how to overcome any shortcomings.
4. **Proactivity and Interruptions**. As smartphones accompany their owners throughout their daily life, and are often within arms reach of their owners [22], mobile location recommender systems can also proactively send notifications to their users about places of potential interest that are around them [27]. The challenge with this feature is understanding the balance between pushing relevant information to users and not overly burdening them with a constant stream of interruptions. Recent work [59] has analysed interruptions within the context of mobile experience-sampling: future work could focus on whether a system could similarly learn about how to appropriately interrupt users to deliver recommendations.
5. **Different Users and Items**. This chapter has focused on recommending places to people. Future mobile systems need not limit themselves to this paradigm. For example, recent work has used mobility patterns to recommend public transport fares [44] and personalise service status updates [45]. Similarly, recent

work has recommended passenger pick-up locations to cab drivers (and vice versa) [84], recommended where to place new retail stores in a city [38], and recommended places to groups of people [70]; the definition of what constitutes a ‘user’ and an ‘item’ is open to many further interpretations.

The list above constitutes a brief set of ideas about future directions for mobile recommender systems. As smartphones’ ability to collect valuable data increases, these devices are beginning to draw the interest of researchers beyond the computer sciences [46]; the future work in this domain has the potential of having far-reaching implications across both research and practical applications.

References

1. Abowd, G., Atkeson, C., Hong, J., Long, S., Kooper, R., Pinkerton, M.: Cyberguide: a Mobile Context-Aware Tour Guide. Wireless Network **3** (2007)
2. Adomavicius, G., Tuzhilin, A.: Context-Aware Recommender Systems. In: ACM Recommender Systems, pp. 335–336. Lausanne, Switzerland (2008)
3. Adomavicius, G., Tuzhilin, A.: Context-aware recommender systems. In: Recommender systems handbook, pp. 217–253. Springer (2011)
4. Ardissono, L., Goy, A., Petrone, G., Segnan, M., Torasso, P.: Intrigue: Personalized Recommendation of Tourist Attractions for Desktop and Handset Devices. Applied Artificial Intelligence **17** (2003)
5. Ardissono, L., Kuflik, T., Petrelli, D.: Personalization in Cultural Heritage: The Road Travelled and the One Ahead. User Modelling and User-Adapted Interaction **22**(1), 73–99 (2012)
6. Baccigalupo, C., Plaza, E.: Case-Based Sequential Ordering of Songs for Playlist Recommendation. Lecture Notes in Computer Science **4106**, 286 – 300 (2006)
7. Backstrom, L., Leskovec, J.: Supervised Random Walks: Predicting and Recommending Links in Social Networks. In: ACM WSDM. Hong Kong, China (2011)
8. Bawa-Cavia, A.: Sensing the Urban: Using Location-Based Social Network Data in Urban Analysis. In: Workshop on Pervasive Urban Applications. San Francisco, USA (2011)
9. Becker, R., Caceres, R., Hanson, K., Isaacman, S., Loh, J., Martonosi, M., Rowland, J., Urbanek, S., Varshavsky, A., Volisky, C.: Human Mobility Characterization from Cellular Network Data. Communications of the ACM **56**(1), 74–82 (2013)
10. Bellotti, V., Begole, B., Chi, E., et al.: Activity-Based Serendipitous Recommendations with the Magitti Mobile Leisure Guide. In: ACM CHI. Florence, Italy (2008)
11. Blondel, V. (ed.): 3rd Conference on the Analysis of Mobile Phone Datasets. Boston, USA (2013)
12. C.H. Tai D.N. Yang, L.L., Chen, M.S.: Recommending Personalized Scenic Itinerary with Geo-Tagged Photos. In: ICME. Hannover, Germany (2008)
13. Cheverst, K., Davies, N., Mitchell, K., Friday, A., Efstratiou, C.: Developing a Context-Aware Electronic Tourist Guide: Some Issues and Experiences. In: ACM CHI. The Hague, The Netherlands (2000)
14. Cho, E., Myers, S., Leskovec, J.: Friendship and Mobility: User Movement in Location-Based Social Networks. In: ACM KDD. San Diego, USA (2011)
15. Choudhury, M., Feldman, M., Amer-Yahia, S., Golbandi, N., Lempel, R., Yu, C.: Automatic Construction of Travel Itineraries using Social Breadcrumbs. In: ACM Hypertext. Ontario, Canada (2010)
16. Choudhury, T., Borriello, G., Consolvo, S., Haehnel, D., Harrison, B., Hemingway, B., Hightower, J., Klasnja, P., Koscher, K., LaMarca, A., LeGrand, L., Lester, J., Rahimi, A., Rea, A.,

Wyatt, D.: The Mobile Sensing Platform: An Embedded Activity Recognition System. IEEE Pervasive Computing **7**(2), 32–41 (2008)
17. Church, K., Smyth, B.: Understanding the Intent Behind Mobile Information Needs. In: International Conference on Intelligent User Interfaces, pp. 247–256. Sanibel Island, FL, USA (2009)
18. Clements, M., Serdyukov, P., deVries, A.P., M.J.T.Reinders: Using Flickr Geotags to Predict User Travel Behaviour. In: ACM SIGIR. Geneva, Switzerland (2010)
19. Cohen, W., Schapire, R., Singer, Y.: Learning to Order Things. Journal of Artificial Intelligence Research **10**(1), 243–270 (1999)
20. Crandall, D., Backstrom, L., Huttenlocher, D., Kleinberg, J.: Mapping the World's Photos. In: WWW. Madrid, Spain (2009)
21. Deshpande, M., Karypis, G.: Item-Based Top-N Recommendation Algorithms. ACM Transactions on Information Systems **22**(1), 143–177 (2004)
22. Dey, A., Wac, K., Ferreira, D., Tassini, K., Hong, J., Ramos, J.: Getting Closer: An Empirical Investigation of the Proximity of Users to their Smartphones. In: ACM Ubicomp. Beijing, China (2011)
23. Domenico, M.D., Lima, A., Musolesi, M.: Interdependence and Predictability of Human Mobility and Social Interactions. In: Nokia Mobile Data Challenge Workshop. Newcastle, United Kingdom (2012)
24. Eagle, N., Pentland, A.: Reality Mining: Sensing Complex Social Systems. Personal and Ubiquitous Computing **10**(4), 255–268 (2006)
25. Froehlich, J., Chen, M., Smith, I., Potter, F.: Voting With Your Feet: An Investigative Study of the Relationship Between Place Visit Behavior and Preference. In: ACM Ubicomp (2006)
26. Froehlich, J., Dillahunt, T., Klasnja, P., Mankoff, J., Consolvo, S., Harrison, B., Landay, J.: UbiGreen: Investigating a Mobile Tool for Tracking and Supporting Green Transportation Habits. In: ACM CHI. Boston, USA (2009)
27. Gallego-Vico, D., Woerndl, W., Bader, R.: A Study on Proactive Delivery of Restaurant Recommendations for Android Smartphones. In: ACM RecSys Workshop on Personalization in Mobile Applications. Chicago, USA (2011)
28. Gao, H., Tang, J., Hu, X., Liu, H.: Exploring Temporal Effects for Location Recommendation on Location-Based Social Networks. In: ACM Recommender Systems. Hong Kong, China (2013)
29. Gavalas, D., Bellavista, P., Cao, J., Issarny, V.: Mobile Applications: Status and Trends. Journal of Systems and Software **84**(11), 1823–1826 (2011)
30. Girardin, F., Blat, J., Calabrese, F., Fiore, F.D., Ratti, C.: Digital Footprinting: Uncovering Tourists with User-Generated Content. IEEE Pervasive Computing **7**(4), 36–43 (2008)
31. Girardin, F., Calabrese, F., Fiore, F.D., Ratti, C., Blat, J.: Digital Footprinting: Uncovering Tourists with User-Generated Content. IEEE Pervasive Computing **7**(4), 36–43 (2008)
32. Girardin, F., Fiore, F.D., Ratti, C., Blat, J.: Leveraging Explicitly Disclosed Location Information to Understand Tourist Dynamics: A Case Study. Journal of Location-Based Services **2**(1), 41–54 (2008)
33. Gonzalez, M., Hidalgo, C., Barabasi, A.L.: Understanding Individual Human Mobility Patterns. Nature **453**(5) (2008)
34. van der Heijden, H., Kotsis, G., Kronsteiner, R.: Mobile Recommendation Systems for Decision Making on the Go. In: IEEE ICMB (2005)
35. Herlocker, J., Konstan, J., Terveen, L., Riedl, J.: Evaluating Collaborative Filtering Recommender Systems. ACM Transactions on Information Systems **22**, 5–53 (2004)
36. Hu, Y., Koren, Y., Volinsky, C.: Collaborative Filtering for Implicit Feedback Datasets. In: IEEE ICDM, pp. 263–272. Pisa, Italy (2008)
37. Jambor, T., Wang, J.: Optimizing Multiple Objectives in Collaborative Filtering. In: ACM Recommender Systems, pp. 55–62. Barcelona, Spain (2010)
38. Karamshuk, D., Noulas, A., Scellato, S., Nicosia, V., Mascolo, C.: Geo-Spotting: Mining Online Location-Based Services for Optimal Retail Store Placement. In: ACM KDD. Chicago, USA (2013)

39. Karatzoglou, A., Baltrunas, L., Church, K., Bohmer, M.: Climbing the App Wall: Mobile App Discovery through Context-Aware Recommendations. In: ACM CIKM. Maui, Hawaii (2012)
40. Kennedy, L., Naaman, M.: Generating Diverse and Representative Image Search Results for Landmarks. In: WWW. Madrid, Spain (2008)
41. Kennedy, L., Naaman, M., Ahern, S., Nair, R., Rattenbury, T.: How Flickr Helps us Make Sense of the World: Context and Content in Community-Contributed Media Collections. In: ACM MM. Augsburg, Germany (2007)
42. Kenteris, M., Gavalas, D., Economou, D.: Electronic Mobile Guides: A Survey. Personal and Ubiquitous Computing **15**(1), 97–111 (2011)
43. Lane, N., Miluzzo, E., Lu, H., Peebles, D., Choudhury, T., Campbell, A.: A Survey of Mobile Phone Sensing. IEEE Communications Magazine (2010)
44. Lathia, N., Capra, L.: Mining Mobility Data to Minimise Travellers' Spending on Public Transport. In: ACM KDD. San Diego, California (2011)
45. Lathia, N., Froehlich, J., Capra, L.: Mining Public Transport Usage for Personalised Intelligent Transport Systems. In: IEEE ICDM. Sydney, Australia (2010)
46. Lathia, N., Pejovic, V., Rachuri, K., Musolesi, M., Rentfrow, P.: Smartphones for Large-Scale Behaviour Change Interventions. IEEE Pervasive Computing, Special Issue on Understanding and Changing Behaviour **12**(3) (2013)
47. Lathia, N., Rachuri, K., Mascolo, C., Rentfrow, P.: Contextual Dissonance: Design Bias in Sensor-Based Experience Sampling Methods. In: ACM Ubicomp. Zurich, Switzerland (2013)
48. Li, Q., Zheng, Y., Xie, X., Chen, Y., Liu, W., Ma, W.: Mining User Similarity Based on Location History. In: Intl. Conf. on Advances in Geographic Information Systems. Santa Ana, USA (2008)
49. Lindqvist, J., Cranshaw, J., Wiese, J., Hong, J., Zimmerman, J.: I'm the Mayor of My House: Examining Why People Use Foursquare - a Social-Driven Location Sharing Application. In: ACM CHI. Vancouver, Canada (2011)
50. Lu, H., Pan, W., Lane, N., Choudhury, T., Campbell, A.: SoundSense: Scalable Sound Sensing for People-Centric Applications on Mobile Phones. In: ACM MobiSys. Krakow, Poland (2009)
51. Marlow, C., Naaman, M., Boyd, D., Davis, M.: Position Paper, Tagging, Taxonomy, Flickr, Article, ToRead. In: Collaborative Web Tagging Workshop (WWW) (2006)
52. Miller, B., Konstan, J., Riedl, J.: PocketLens: Toward a Personal Recommender System. In: ACM TOIS (2005)
53. Monreale, A., Pinelli, F., Trasarti, R., Giannotti, F.: WhereNext: A Location Predictor on Trajectory Pattern Mining. In: ACM SIGKDD, pp. 637–646. Paris, France (2009)
54. Noulas, A., Scellato, S., Lambiotte, R., Pontil, M., Mascolo, C.: A Tale of Many Cities: Universal Patterns in Human Urban Mobility. PLoS ONE **7**(5) (2012)
55. Noulas, A., Scellato, S., Lathia, N., Mascolo, C.: A Random Walk Around the City: New Venue Recommendation in Location-Based Social Networks. In: IEEE International Conference on Social Computing. Amsterdam, The Netherlands (2012)
56. Noulas, A., Scellato, S., Lathia, N., Mascolo, C.: Mining User Mobility Features for Next Place Prediction in Location-based Services. In: IEEE Internationcal Conference on Data Mining. Brussels, Belgium (2012)
57. Noulas, A., Scellato, S., Mascolo, C., Pontil, M.: An Empirical Study of Geographic User Activity Patterns in Foursquare. In: AAAI ICWSM. Barcelona, Spain (2011)
58. Page, L., Brin, S., Motwani, R., Winograd, T.: The PageRank Citation Ranking: Bringin Order to the Web. In: Technical Report Stanford InfoLab. Stanford, USA (1999)
59. Pejovic, V., Musolesi, M.: InterruptMe: Designing Intelligent Prompting Mechanisms in Pervasive Applications. In: ACM Ubicomp. Seattle, USA (2014)
60. Popescu, A., Grefenstette, G.: Deducing Trip Related Information from Flickr. In: WWW. Madrid, Spain (2009)
61. Quercia, D., Capra, L.: FriendSensing: Recommending Friends Using Mobile Phones. In: ACM RecSys. New York, USA (2009)

62. Quercia, D., Lathia, N., Calabrese, F., Lorenzo, G.D., Crowcroft, J.: Recommending Social Events from Mobile Phone Location Data. In: IEEE ICDM. Sydney, Australia (2010)
63. Quercia, D., Leontiadis, I., McNamara, L., Mascolo, C., Crowcroft, J.: SpotME If You Can: Randomized Responses for Location Obfuscation on Mobile Phones. In: ICDCS. Minneapolis, USA (2011)
64. Rachuri, K., Mascolo, C., Musolesi, M.: Energy-Accuracy Trade-offs of Sensor Sampling in Smart Phone based Sensing Systems. In: Mobile Context Awareness: Capabilities, Challenges and Applications Workshop. Springer, Copenhagen, Denmark (2010)
65. Rachuri, K., Mascolo, C., Musolesi, M., Rentfrow, P.: SociableSense: Exploring the Trade-offs of Adaptive Sampling and Computation Offloading for Social Sensing. In: ACM MobiCom. Las Vegas, USA (2011)
66. Rattenbury, T., Good, N., Naaman, M.: Toward Automatic Extraction of Event and Place Semantics from Flickr Tags. In: ACM SIGIR, pp. 103–110 (2007)
67. Ratti, C., Pulselli, R., Williams, S., Frenchman, D.: Mobile Landscapes: Using Location Data from Cell Phones for Urban Analysis. Environment and Planning B **33**(5), 727–748 (2006)
68. Ricci, F.: Mobile Recommender Systems. Journal of IT & Tourism **12**(3), 205–231 (2011)
69. Ricci, F., Nguyen, Q.N.: Critique-Based Mobile Recommender Systems. OGAI Journal (2005)
70. Salamo, M., McCarthy, K., Smyth, B.: Generating Recommendations for Consensus Negotiation in Group Personalization Services. Personal and Ubiquitous Computing **16**(5), 597–610 (2012)
71. Scellato, S., Musolesi, M., Mascolo, C., Latora, V., Campbell, A.: NextPlace: A Spatio-Temporal Prediction Framework for Pervasive Systems. In: Ninth International Conference on Pervasive Computing. San Francisco, USA (2011)
72. Schiller, J., Voisard, A. (eds.): Location-Based Services. Morgan Kaufman Publishers (2004)
73. van Setten, M., Pokraev, S., Koolwaaij, J.: Context- Aware Recommendations in the Mobile Tourist Application COMPASS. In: Adaptive Hypermedia and Adaptive Web-Based Systems. Eindhoven, The Netherlands (2004)
74. Shaw, B., Shea, J., Sinha, S., Hogue, A.: Learning to Rank for Spatiotemporal Search. In: ACM WSDM. Rome, Italy (2013)
75. Sklar, M., Shaw, B., Hogue, A.: Recommending Interesting Events in Real Time with Foursquare Checkins. In: ACM Recommender Systems. Dublin, Ireland (2012)
76. Sohn, T., Varshavky, A., LaMarca, A., Chen, M., Choudhury, T., Smith, I., Consolvo, S., Hightower, J., Grisworld, W., de Lara, E.: Mobility Detection Using Everyday GSM Traces. In: ACM Ubicomp. Orange County, USA (2006)
77. Stenneth, L., Wolfson, O., Yu, P., Xu, B.: Transportation Mode Detection using Mobile Phones and GIS Information. In: ACM SIGSPATIAL. Chicago, USA (2011)
78. Takeuchi, Y., Sugimoto, M.: CityVoyager: an Outdoor Recommendation System Based on User Location History. Ubiquitous Intelligence and Computing (2006)
79. Tintarev, N., Amatriain, X., Flores, A.: Off the Beaten Track: A Mobile Field Study Exploring the Long Tail of Tourist Recommendations. In: MobileHCI. Lisbon, Portugal (2010)
80. Tong, H., Faloutsos, C., Pan, J.: Fast Random Walk with Restart and Its Applications. In: IEEE International Conference on Data Mining. Hong Kong, China (2006)
81. Tung, H., Soo, V.: A Personalized Restaurant Recommender Agent for Mobile E-Service. In: Proceedings of IEEE International Conference on e-Technology, e-Commerce, and e-Services, pp. 259–262. Washington DC, USA (2004)
82. Yang, W., Cheng, H., Dia, J.: A Location-Aware Recommender System for Mobile Shopping Environments. Expert Systems with Applications (2008)
83. Yoon, H., Zheng, Y., Xie, X., Woo, W.: Social Itinerary Recommendation from User-Generated Digital Trails. Personal and Ubiquitous Computing **16**(5), 469–484 (2012)
84. Yuan, N., Zheng, Y., Zhang, L., Xie, X.: T-Finder: A Recommender System for Fidning Passengers and Vacant Taxis. IEEE Transactions on Knowledge and Data Engineering **25**(10) (2013)

85. Zhang, A., Noulas, A., Scellato, S., Mascolo, C.: Hoodsquare: Modeling and Recommending Neighborhoods in Location-Based Social Networks. In: SocialCom. Washington DC, USA (2013)
86. Zheng, V., Zheng, Y., Xie, X., Yang, Q.: Collaborative Location and Activity Recommendations with GPS History Data. In: WWW. Raleigh, North Carolina (2010)
87. Zheng, V.W., Cao, B., Zheng, Y., Xie, X., Yang, Q.: Collaborative filtering meets mobile recommendation: A user-centered approach. In: AAAI (2010)
88. Zheng, Y., Li, Q., Chen, Y., Xie, X., Ma, W.: Understanding Mobility Based on GPS Data. In: ACM Ubicomp. Seoul, Korea (2008)
89. Zheng, Y., Zhang, L., Xie, X., Ma, W.: Mining Interesting Locations and Travel Sequences From GPS Trajectories. In: WWW. Madrid, Spain (2008)