

GeoSRS: A Hybrid Social Recommender System for Geolocated Data

Joan Capdevila^{a,c,1}, Marta Arias^{b,c,2}, Argimiro Arratia^{b,c,2,3}

^a*Barcelona Supercomputing Center*

^b*Computer Science Department*

^c*Barcelona Tech / Universitat Politècnica de Catalunya, Spain*

Abstract

We present *GeoSRS*, a hybrid recommender system for a popular location-based social network (LBSN), in which users are able to write short reviews on the places of interest they visit. Using state-of-the-art text mining techniques, our system recommends locations to users using as source the whole set of text reviews in addition to their geographical location. To evaluate our system, we have collected our own datasets by crawling the social network *Foursquare*. To do this efficiently, we propose the use of a parallel version of the *Quadtree* technique, which may be applicable to crawling/exploring other spatially distributed sources. Finally, we study the performance of *GeoSRS* on our collected dataset and conclude that by combining sentiment analysis and text modelling, *GeoSRS* generates more accurate recommendations. The performance of the system improves as more reviews are available, which further motivates the use of large-scale crawling techniques such as the *Quadtree*.

Key words: recommender systems, text mining, quadtree, crawling, social networks, location-based social network

Email addresses: `jc@ac.upc.edu` (Joan Capdevila), `marias@cs.upc.edu` (Marta Arias), `argimiro@cs.upc.edu` (Argimiro Arratia)

¹Supported by Obra Social “la Caixa”

²Supported by MICINN project TIN2011-27479-C04-03 (BASMATI), MINECO project TIN2014-57226-P (APCOM) and Gen. Cat. project SGR2014-890 (MACDA)

³Additional support by MEC project MTM2012-36917-C03-03 (SINGACOM)

1. Introduction

The proliferation of mobile communication and GPS technologies has allowed users to add geographical identification metadata to various social media, such as photographs, text reviews or video, among many others. Location-based social networks (LBSNs) (Zheng, 2012), integrate into a single network user relations (the “social” part) and geo-spatial information (the “location-based” part). By taking into account the physical location of users, LBSNs are bridging the gap between physical world and virtual communities such as *Foursquare*⁴, *Facebook*⁵ or *Twitter*⁶.

The extensive use of these social networking sites has made them invaluable sources of information. However, the sheer volume of data flowing through these sites, even for a single user, has made it increasingly difficult for humans to track all this information. Therefore, most social networking sites implement some sort of Social Recommendation System (SRS) (Guy and Carmel, 2011): for example, Twitter suggests who to follow, Facebook filters and prioritizes posts in users’ walls and Foursquare recommends locations where to go. When the social content is geotagged, it becomes strictly necessary to consider user and item localization in the recommendation paradigm.

Location-based recommendation constitutes a unique application in LBSNs and it substantially differs from traditional recommender systems in the fact that the latter does not take into account the spatial properties of users and items (Mokbel et al., 2011). Moreover, location-based recommendation on top of LBSNs might also benefit from the interaction between the three layers composing a LBSN, namely the user, the location and the content layer (Symeonidis et al., 2014a).

In this paper, we propose a fully integrated system for *information retrieval* of geolocated data and *end-to-end location-based recommender*, suitable for the popular social network *Foursquare*. The reader should note, however, that our methods are applicable to any other social network that contains geolocated time-referenced reviews and hence, in our presentation we abstract from the fact that we are using this particular site.

We believe that fully operative recommender systems on top of LBSNs require end-to-end designs, capable of performing data retrieval from social

⁴<https://foursquare.com/>

⁵<https://facebook.com/>

⁶<https://twitter.com/>

networks, cleaning the noisy and duplicated data, extracting relevant features and, not least, performing recommendation.

Our proposed location-based information system is summarized in Fig. 1. It retrieves the short reviews together with their geolocated venues and reviewers identification as its basis for recommendations. In Foursquare, users are able to *check-in* to places of interest (“venues”), write short reviews (“tips”) for the venues where they checked into, and share this information with users within their social network. For the task of crawling venues, users

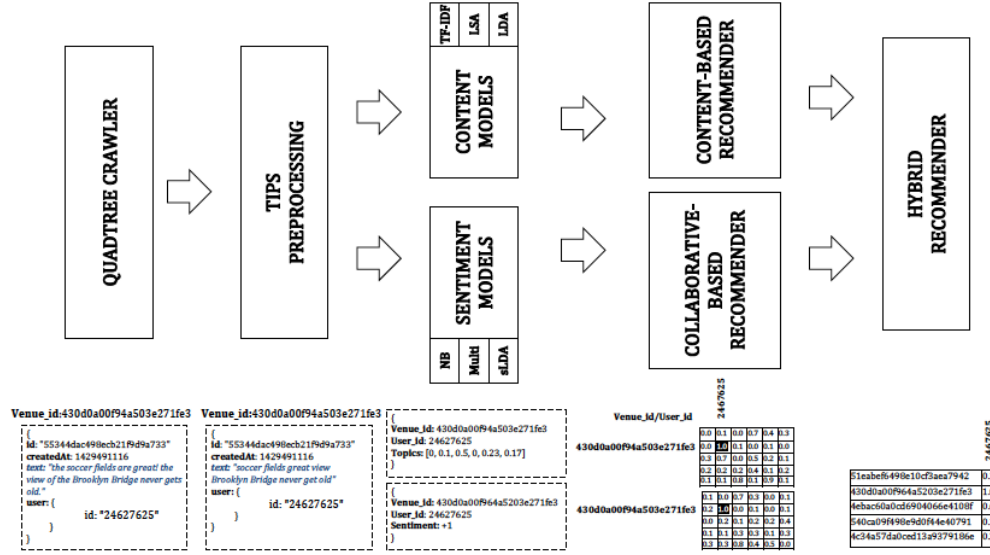


Figure 1: Fully integrated location-based Information System. It includes modules for crawling, preprocessing crawled data, modeling and recommendation. The bottom part of the figure shows examples of data chunks that each module produces.

and tips, we had to use Foursquare’s API, an interface which imposes restrictions to the amount of information one can query and the amount of requests one can make. Therefore, it was imperative to devise a crawling mechanism that would make optimal use of the queries available to us. To achieve this task, we have designed a parallel version of the *Quadtree* algorithm (Samet, 1984), which is very well suited for crawling venues that are spatially distributed, while at the same time gaining considerable throughput. We have found that crawling all venues from large urban areas such as Mexico D.F. or New York in reasonable time was possible by the proposed Quadtree algorithm. We consider the application of the parallel Quadtree

algorithm to this problem an important contribution of our paper and we believe that problems that require sweeping spatial devices (sensors, etc.) could also benefit from it.

To make recommendations our system makes extensive use of user’s reviews (“tips”). In order to extract meaningful information from these free-form reviews, *GeoSRS* relies upon many state-of-the-art techniques for text mining and sentiment analysis, which are evaluated in terms of recommendation accuracy and the ones that outperform are selected to be used in *GeoSRS*. Another relevant contribution from our paper is the increase of accuracy when mixing the review’s sentiment and content into a simple but rather effective weighted hybrid recommender setup (Burke, 2007). Sentiment refers to the global opinion that is reflected in the review (positive, negative or neutral) while content indicates the topics that the review addresses. This enforces the idea that pure review-based choices are not merely based on the opinion reflected on a short review (“The service was too slow”), but also on the content relevant to the user (“This is a kinda working place rather than a coffee shop”).

To evaluate our system, we have collected our own dataset of restaurants and tips from the area of Manhattan in New York City. We have chosen Manhattan due to the high density of venues and the number of active users, to validate both the scalability of the Quadtree crawler and the effectiveness and coverage of the recommender system. Recommender system is evaluated in terms of retrieval accuracy (performance) measures rather than statistical accuracy measures since we do not intend to predict individual venue ratings but relative order among them. Moreover, we propose an evaluation method in which we divide historically the tips dataset in training and testing. Test tips are taken as ground truth to comparatively assess the recommendation. Last but not least, the simple weighted hybrid recommender setup employed in *GeoSRS* is compared against other state-of-the-art configurations such as meta-level and cascade models.

To summarize, this paper proposes for a hybrid recommender system for a location-based social network which is uniquely built upon text reviews. Our main contributions are:

1. Using a parallel version of the *Quadtree* technique as the basic strategy for crawling spatially distributed data.
2. Using sentiment analysis on text reviews to generate the source for collaborative-filtering.

3. Using the aggregated reviews by user and venue respectively to generate the profiling information for content-based recommendation.
4. Using a simple but powerful hybridization technique to improve recommendation performance.
5. Putting it all together into a working information system.
6. Evaluating and comparing *GeoSRS* against other state-of-the-art hybrid systems in terms of IR figures.

The rest of this paper is organized as follows. Section 2 presents an overview of related work in social recommendation on top of LBSNs, sentiment and content analysis systems and general hybrid recommendation techniques. Section 3 proposes a parallel efficient technique to retrieve spatially distributed data sources. Our system *GeoSRS* is described in Section 4. Section 5 brings together the geolocated reviews data set from Foursquare with the working *GeoSRS* system, to assess different possible set-ups using an offline evaluation methodology. Last but not least, Section 6 derives several conclusions from our work and includes directions for future work.

2. Related work

The recommender system proposed in this paper falls within the class of *location-based social recommender systems*, using *sentiment and content analysis of text* combined with collaborative filtering techniques leading to a *hybrid recommender system*. In order to place our system in context of the known literature we briefly review each of these research areas relevant to our work.

Social Recommender Systems (SRS). SRS have arisen as an application of recommender systems to social networks, although they have been accepted lately as a separate discipline in itself (Guy and Carmel, 2011). Nonetheless, researchers have been proposing novel recommender set-ups based on social content for the last five to ten years (Falahi et al., 2012; Guy and Carmel, 2011; Jameson, 2004; Liu and Aberer, 2013).

Classical recommendation paradigms such as content-based recommenders have been enhanced by adding friendship information into the matrix factorization objective function (Liu and Aberer, 2013), configuring the so-called Social Content-based Recommender (SoCo). Collaborative-based filtering has also been redefined by including social information to improve the user

neighborhood identification and deal with data sparsity (Groh and Ehmig, 2007; Ma et al., 2008). According to Groh and Ehmig (2007), social filtering outperforms collaborative-based filtering. Trust-based recommenders refine the notion of social filtering even further by defining several notions of propagation of trust or reputation through the network of users. Examples of this are *FeedbackTrust* (Moghaddam et al., 2009) and *TrustWalker* (Jamali and Ester, 2009) although many others exist (Massa and Avesani, 2007; Andersen et al., 2008; O’Doherty et al., 2012). *FeedbackTrust* improves user-based recommendations by enhancing user’s similarities with trust-based similarities; *TrustWalker* makes item-based recommendations by combining the result of repeated random-walks over the network of connected users.

Others have proposed mechanisms to include interpersonal influences in traditional recommender systems arguing that the interpersonal influence plays a critical role in this scenario (Huang et al., 2010). The term *social regularization* has been coined to refer to the use of a regularization based on social content. In the context of *group recommendation*, which consists in making recommendations to a group of individuals based on their interests, the work of Gartrell et al. (2010) highlights the benefits of incorporating social structure into the recommender.

Several practical SRS can be found in the literature applied to distinct recommendation situations. For example, Diaby et al. (2013) describe an online social network-based job recommender system for recruiters. Tu et al. (2014) present an online dating recommending system that learns the user preferences from a Latent Dirichlet Allocation (LDA) model. Xia et al. (2014) propose a system that takes into account social ties in the area of scientific articles to recommend scholarly papers to users.

Location-based Social Recommender Systems. Recommender systems can be further improved by exploiting geolocated data to take into account the spatial dimension (Sarwt et al., 2013); such systems are generally referred to as *location-based social recommender systems*. Examples in this line of research are (Berjani and Strufe, 2011; Yang et al., 2013; Ye et al., 2010; Zheng et al., 2011); for an introduction to LBSN recommender systems and updates on the state-of-the-art see the tutorial (Zheng, 2012) or the recently published textbook by Symeonidis et al. (2014b). Compared to Yang et al. (2013), we are using the same data source, and due to similarity of contributions we hereby present our contributions with theirs. Yang et al. (2013) proposed a location-based recommendation system that uses ge-

ollocated tips from Foursquare jointly with check-in information to improve the recommendation performance. Our hybrid recommender system *GeoSRS* differs from this work in various aspects. While Yang et al. (2013) use only sentiment information from tips, we also extract the topics structure from them to model user and venue profiles. Additionally, authors based the recommendation scheme on a probabilistic factorization of the user-item matrix that considers social influence, whereas we found out that by using social influence on the collaborative-based branch of our hybrid system performed poorly. Moreover, we evaluate the system by assessing the relative order of near venues compared with the actual attendance rather than calculating the statistical accuracy of the estimated ratings as they do. In our opinion, our evaluation methodology is closer to the real recommendation behavior of a location-based system. Therefore we believe that our work complements their work on other important aspects of location-based social recommender systems that have not been considered before.

Sentiment and text analysis. In this work, we propose the use of some of the already mentioned and more of the state-of-the-art text mining techniques to build a purely review-based recommender system. Several text modelling techniques are assessed under our recommendation scheme, such as Latent Dirichlet Allocation (LDA) (Blei et al., 2003), Latent Semantic Analysis (LSA) (Deerwester et al., 1990), or TF-IDF (Sparck Jones, 1972; Salton and Buckley, 1988). Sentiment analysis is also included in these models by using trained traditional classifiers such as Naive Bayes (McCallum et al., 1998), Multinomial Logistic Regression (Böhning, 1992) or supervised LDA classifiers (Blei and McAuliffe, 2008).

There are many other recommender systems that leverage the information found in free-form text. In fact, previous work show that considering text sources improve standard collaborative filtering techniques. As an example, (Aciar et al., 2007) build a recommender system for consumer products based on product reviews, or (Jakob et al., 2009) improve movie recommendations based on movie reviews. (Reschke et al., 2013) propose a recommendation dialog system built upon narrow questions from reviews, which slightly differs from the recommendation problem definition. In contrast to all these systems, our recommender system bases the whole recommendation on the reviews text data which is modeled in a flexible and general procedure rather than building a text ontology or extracting multiple aspects (Snyder and Barzilay, 2007; Musat et al., 2013).

Social context has been used to improve results in problems other than recommendation. An example that is directly relevant to our work here is the problem of text analysis and in particular of *review quality detection* (Lu et al., 2010). Most previous work in this field treat each review as a stand-alone text document, extracting features from the text and learning a function based on the features (Zhang and Varadarajan, 2006; Kim et al., 2006; Liu et al., 2008). Naturally, taking into account additional information in the form of social relations between authors of reviews should help improve the predictions.

Hybrid social recommenders. A good deal of researchers have been combining multiple recommendation techniques to boost the performance of the so-called hybrid systems (Burke, 2002). Standalone collaborative or content-based recommendation systems suffer from several shortcomings which can be overcome by coupling their individual ratings. In this work, we show that these deficiencies of individual recommenders happen as well in our problem scenario and highlight the benefits of hybridization. Among a broad range of hybridization techniques (Burke, 2007), linear weighted combination seems to be one of the most simple but rather effective mechanisms, and it is the one we implement here. For example, (Mobasher et al., 2004) mix two components, a collaborative and a content-based branch, with a linear weighting scheme to perform movie recommendation. For the sake of comparison, we further develop two other hybrid mechanisms named the *cascade* and the *meta-level* hybrid, which (Burke, 2007) found out to work well when combining two components of differing strength. A cascade hybrid called *Entree* was developed by (Burke, 2002), who combined a knowledge-based and collaborative recommender in a hierarchical manner based on the strength of their recommendation. Finally, a meta-level hybrid that uses content-based recommendation to identify the collaborative neighbors can be found in (Pazzani, 1999).

3. Data retrieval process

3.1. Social networks as sources of open data

Social networks (Facebook, Twitter, Foursquare, LinkedIn, Instagram, etc.) act as consolidated data warehouses unifying distinct users’ social activities into common data schemes, which are often mined by in-house analytic systems or accessible to application developer communities. These platforms

implement data security and privacy policies, accepted beforehand by the user, which rule the accessibility to the users' social data.

We refer to social open data as the content (posts, tweets, tips, publications, photos, etc.) which is visible or public to anyone. In the last decade, open social data became essential in many research fields ranging from recommender systems to urban design and planning (Resch et al., 2012). Among all types of social open data, geolocated data cover a broad range of media types that include geographical coordinates. This is the case of location-based social networks which link some media (tweets, tips, etc.) to the user location gathered through the smart-phone GPS system. It also considers modern urban cities containing thousands of physical sensors spread over large geographical areas, from which their sensed information (e.g. air-pollution, traffic, light-level, etc.) is linked to a geographical coordinate.

Geolocated data, generated either by users acting as social sensors or real physical sensors, can be consulted from third-party data providers (social networks, open data portals, etc.) by means of geospatial queries. Typically, these entities implement control mechanisms to avoid server traffic overload, which hamper the retrieval of the totality of the data at once. Hence, any retrieval process for geolocated data has to carefully take into account these traffic limitations.

3.2. Data retrieval from social networks

Most of social networks give access to their data via Representational State Transfer (REST) Application Programming Interfaces (API). This protocol enables an easy but effective interaction with the social network data contents. However, these web services typically limit the rate of requests per registered application and the response data volumes in order to avoid incoming traffic overload.

For example, the geospatial query to obtain all the venues or places of a given geographical area from the Foursquare social network is constrained to return 50 venues at most. This necessarily implies that this area has to be divided into smaller sub-areas with less than 50 venues each to effectively retrieve all venues. Moreover, the Foursquare platform also limits up to 5000 requests per hour and per registered application, forcing the retrieval process to prioritize optimal requests. Here, we propose and motivate the use of *Quadtree* structures and the Quadtree construction algorithm, to effectively and efficiently retrieve all geolocated content from a given region generated in location-based social networks exhibiting the above-mentioned constraints.

3.3. The Quadtree algorithm

A Quadtree is a data structure based on the principle of recursive decomposition of space. Quadrees are used as hierarchical data representations in the domains of computer vision, image processing, pattern recognition and geographic information systems. The interest in this data structure stems from the fact that it is designed to focus resources on the areas where the information is of greatest density. For a comprehensive survey on quadrees and related hierarchical data structures see (Samet, 1984).

The general Quadtree construction scheme is presented in Algorithm 1, and it works as follows. The algorithm iteratively divides each region or cell into four subregions or subcells when the maximum capacity per cell, denoted N_{max} , is reached. Given two coordinates (the southern-west and northern-east limits) defining the geographical bounding box, the algorithm defines and queues a quadcell object which contains its geographical limits.

Algorithm 1: Quadtree Algorithm.

```

quad ← Quadcell(NElim, SWlim);
queue ← List(quad) ;
while Length(queue) > 0 do
    q ← Pop(queue);           // Obtains quad from the queue
    if CheckQuad(q) then
        aux ← SplitQuad(q);
        Extend(queue,aux);    // Push quad into the queue
    end
end
end
```

For each quadcell in the queue, the algorithm queries the social network REST API, and compares the cardinality of the response against the maximum number of sensors per cell, N_{max} . This task is performed by the *CheckQuad* function described in Algorithm 2.

Algorithm 2: Check Quadcell Function.

```
Function CheckQuad(quad: Quadcell) : boolean
| sensors  $\leftarrow$  APIrequest(quad.NElim,quad.SWlim) ;
| if Length(sensors) <  $N_{max}$  then
| | SaveSensors(sensors);
| | return False;                                // Do not split quad
| else
| | return True;                                // Split quad
| end
end
```

In case the response size exceeds N_{max} , the quadcell is divided into four subregions or children, whose geographical limits are computed from its parents bounds, and these are also added into the pending queue. This task is performed by the *SplitQuad* function described in Algorithm 3. Leaves or quadcells whose API query returned fewer sensors than the maximum allowed, are stored into disk and the quadcells removed from the queue.

Algorithm 3: Split Quadcell Function.

```
Function SplitQuad(quad: Quadcell) : List
| // Computes the NElim and SWlim for each child
| NElims, SElims, SWlims, NWlims  $\leftarrow$ 
| ChildrenLim(quad.NElim,quad.SWlim) ;
|
| // Creates List of children
| children  $\leftarrow$  List(Quadcell(NElims.NE,NElims.SW),
| Quadcell(SElims.NE,SElims.SW),
| Quadcell(SWlims.NE,SWlims.SW),
| Quadcell(NWlims.NE,NWlims.SW)) ;
| return children;
end
```

In order to interpret the results of the Quadtree algorithm for sensors spread over a bounded region, we simulate the Quadtree execution for different spatial distributions, all containing equal number of sensors (top of Figure 2). In other words, all distributions integrate to N , the number of sensors in the region.

The resulting Quadtree structures (bottom of Figure 2), clearly shows the dependency between the peakiness of the distribution and the number

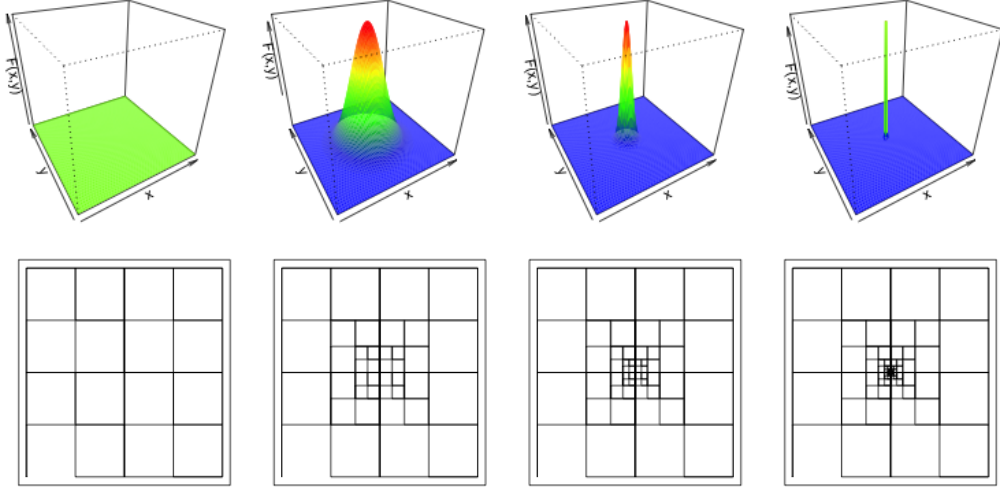


Figure 2: At the top, spatial distribution of geolocated data uniform and normally distributed with standard deviation $\sigma = 100, 10, 0.25$, respectively. Corresponding quadtree structures at the bottom.

of quadcells. Given that each quadcell implies at least one HTTP query to the REST API, peaky distributions of geospatial data result in less efficient quadtree executions compared to flatter and more uniform distributions.

The Quadtree algorithm overcomes the API limitations by first partitioning the geographical space into appropriate subregions and secondly by keeping the HTTP request rate bounded. We have not included the appropriate instructions for the control process over the request rate to keep the description in Algorithm 1 simple. However, the implementation of the Quadtree algorithm should keep control over the time between consecutive requests or the overall number of requests per hour, and sleep the program when necessary.

3.4. *Scaling-up the Quadtree algorithm*

Efficiency and scalability are key with today's volume of data available. Here, we propose a parallel version of the Quadtree algorithm. First, we state the main scalability drawbacks of the proposed algorithm and then revisit the algorithm to propose a new parallel Quadtree algorithm.

We assume that the maximum number of request per hour and per registered application is R_{max} . This limit is imposed by the social network administrator. When the number of sensors, N , is large R_{max} becomes the limiting factor of the algorithm performance.

Our approach uses parallelization of the algorithm into K sub-processes, each using a different registered application key and hence, enabling greater request rates (with K sub-processes the limit becomes KR_{max}). These sub-processes could either run at the same machine or into different machines with distinct public IP addresses, depending on social network directives.

One classical way to tackle this multiprocess problem is by means of the producer-consumer paradigm. The producer is in charge of querying the social network API, generating and queuing the quadcells into a processing queue. Then, the consumer takes each queued quadcell and stores it into the disk. By splitting the two most time consuming sub-processes (querying the API and storing to disk), we experience considerable gains into the parallelization of the Quadtree algorithm.

The producer-consumer scheme for the parallelization of Quadtree, described in Algorithms 4 and 5, works as follows. The producer threads (Algorithm 4) produce quadcells (*prodQuad*) by querying the Social Network API about the quadcells stored in the pending queue (*PendQueue*), checks whether the cardinality of the responses exceed the maximum number of sensors per cell (N_{max}) and splits the quadcells if they exceed. The split quadcells whose number of sensors exceed N_{max} are also push into the pending queue for the next producer. Then, the producer stores the leaves or quadcells that have less sensors than N_{max} into the processing queue (*ProcQueue*), and notifies the condition variable (*Cond*) to release the underlying lock. This notification awakes the consumers threads (Algorithm 5) waiting in the condition variable, which consume quadcells (*consQuad*) from the processing queue. Consuming means pulling the quadcell from the queue and storing the sensors' values into disk. When the consumer thread completes its task, it waits in the condition variable for the next release triggered by the producers.

Algorithm 4: Producer Quadtree Algorithm.

```
Function Producer (PendQueue: Queue, ProcQueue: Queue, Cond:
Condition, API: APIhand ) : void
    while countQuads(PendQueue) > 0 do
        with(Cond);
        q = prodQuad(PendQueue, API);
        putQuad(ProcQueue, q);
        notify(Cond);
    end
    with Cond: ;
    notifyAll(Cond)
end
```

Algorithm 5: Consumer Quadtree Algorithm.

```
Function Consumer (PendQueue: Queue, ProcQueue: Queue,
Cond: Condition ) : void
    while countQuads(PendQueue) > 0 do
        with(Cond) ;
        Wait(Cond) ;
        if countQuads(ProcQueue) > 0 then
            | consQuad(ProcQueue)
        end
    end
end
```

3.5. Case study: the Foursquare platform

The Foursquare platform provides a REST API to interact with its components as well as to access its open social data from registered applications. Querying the platform on the geolocated venues data requires satisfying the response size constraints (50 venues per request) and rate limitations (5000 request per hour and application). The parallel Quadtree algorithm enables an effective and efficient data retrieval process by parallelizing API requests into K_p subprocesses and storing to disk into K_c subprocesses. We found out that using $K_p = 3$ producers and $K_c = 10$ consumers exhibits a proper performance, although the optimization of these parameters is beyond the scope of this paper.

Figure 3 shows the Quadtree structures for three large urban areas: Athens,

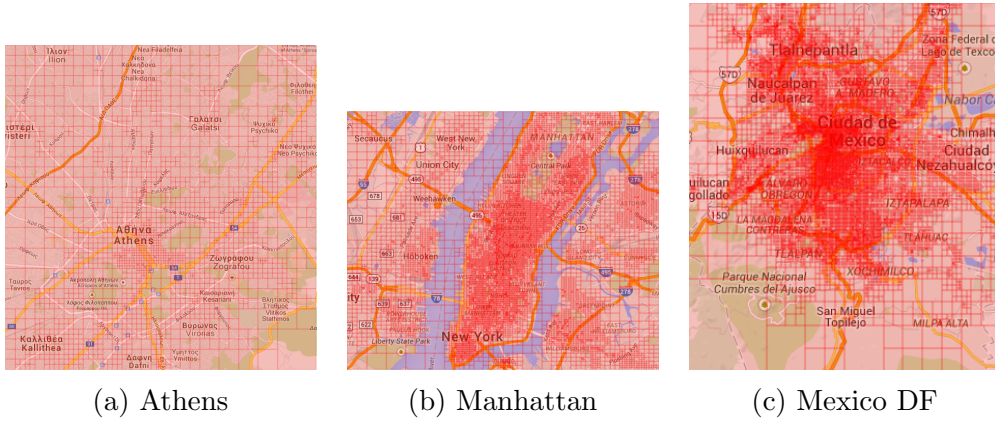


Figure 3: Parallel Quadtree structures

Manhattan and Mexico City. As we could expect, the Quadtree grows deeper in the city downtown and business and commercial subareas, while being shallower in residential areas, parks and city surroundings. Table 1 presents some other features about the Quadtrees built in each of these urban areas.

	<i>Athens</i>	<i>New York</i>	<i>Mexico D.F.</i>
NE lim	38.03, 23.79	40.80,-73.91	19.59, -98.94
SW lim	37.95, 23.69	40.70,-74.07	19.13, -99.36
Venues	48.215	297.924	511.096
Quadcells	2.997	18.682	32.270
Quadcell leaves	2.248	14.011	24.202
Quadtree time (s)	4.432	27.980	32.015

Table 1: Urban Quadtrees with $K_p = 3$ producers and $K_c = 10$ consumers

4. GeoSRS system description

4.1. Overview

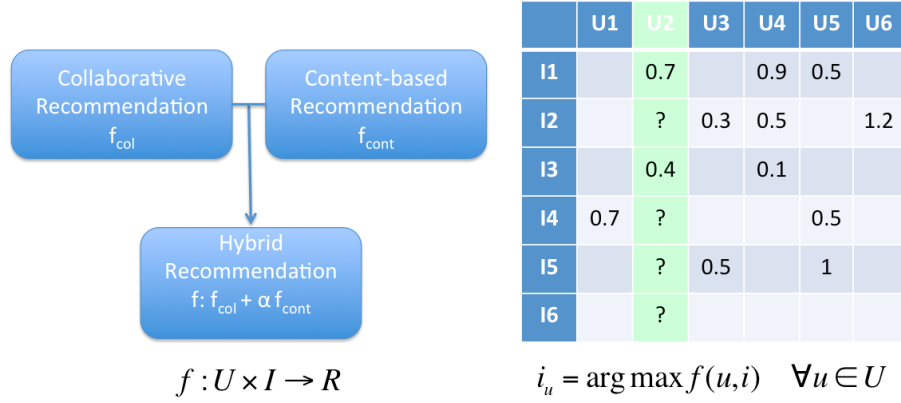


Figure 4: A graphical overview of the hybrid recommender system and its recommendation matrix.

GeoSRS uses a rating-based scheme to perform item recommendation. This means that given a utility function $f : U \times I \rightarrow R$, a rate or measure of importance is estimated from the set of all users, U , and the set of all items, I . The recommendation problem can be seen as the following optimization problem,

$$i_u = \arg \max_{i \in I} f(u, i) \quad \forall u \in U \quad (1)$$

in which, for each user u , the goal is to determine the item, i , that maximizes the utility function among all available items. The item, i_u , that satisfies this maximization is recommended to the user u . However, the utility function is usually not defined for the whole space $U \times I$, and consequently, the recommender needs to estimate the missing ratings to perform an accurate recommendation, as shown in the recommendation matrix from (Figure 4).

GeoSRS is a hybrid social recommender system that combines a collaborative and a content-based subsystems to define the utility function which estimates the ratings for the unseen items. In particular, the utility function is defined from the user-item affinity as well as the community opinion by combining them into a hybrid system, given by the following linear combination,

$$f = f_{col} + \alpha f_{cont}$$

where α is a positive real, and while the content-based branch, f_{cont} , mainly uses the inferred features of the item and the profile of the user, the collaborative branch, f_{col} , takes advantage of the “wisdom of the crowd”.

4.2. Content-based branch

The content-based recommendation branch brings into the system the user preferences and item features with the purpose of making recommendations according to the user profile. Following the introduced nomenclature, the utility function for the content-based branch predicts the rate for both seen and unseen items. The rate, R , is predicted from the dot product of the users profile vectors, w_u , and items features vectors, w_i . Thus, the utility function can be redefined as,

$$f_{cont} : w_u \times w_i \rightarrow R \quad (2)$$

Due to the fact that items and users spaces are large, mechanism to reduce dimensionality or filtering items/users are often used.

4.2.1. Similarity Measurement

Similarity measurement among pairs of users and items can be either calculated using heuristic models or statistical models learned from the underlying data (Adomavicius and Tuzhilin, 2005). Our proposed approach for the content-based branch uses a cosine-based heuristic model in order to compute the similarity rate between user profiles and items features. Hence, the utility function formulation can be expressed as:

$$f_{cont} = \cos(w_u, w_i) = \frac{\vec{w}_u \vec{w}_i}{||\vec{w}_u|| ||\vec{w}_i||} \quad (3)$$

where w_u and w_i are the profile vectors for the users and items defined from the reviews content models described below.

4.2.2. Review content models

Text reviews have been preprocessed using standard techniques, in particular we have removed stop-words and punctuation marks, stripped white space, and converted all text to lower case.

In this work we consider several state-of-the-art techniques for review modeling. The aim of modeling reviews is to build descriptive profiles for both users and items. To this end, we aggregate all reviews by a user, and

consider the aggregated reviews as a single text document to be modeled, w_u . Analogously, we aggregate all reviews of a given item in order to build a document that once modeled will represent the item, w_i . In what follows, we list and briefly explain the modeling techniques we have used to model aggregated text reviews into K components for both profile vectors.

Term Frequency - Inverse Document Frequency. TF-IDF maps the contents of a review into a set of K keywords depending on the measurement of importance through the calculation of a numeric statistic (Sparck Jones, 1972; Salton and Buckley, 1988). The numeric statistic is calculated for each word from the review and then, the K more relevant words are kept.

Latent Semantic Analysis. LSA can be used to map the contents of a review into K latent concepts which turn out to be words that are close in meaning (Deerwester et al., 1990). LSA is performed applying single value decomposition of a term-document matrix. This matrix contains the terms or words in the rows and the documents or reviews in the columns. The real number in the intersection of rows and columns indicates the occurrence of terms in the document. The K largest eigenvalues and their corresponding eigenvectors lead to a rank K approximation of the term-document matrix.

Latent Dirichlet Allocation. LDA is a generative probabilistic model in which documents or reviews are represented as random mixtures over latent topics, and each topic is characterized by a distribution over words (Blei et al., 2003). LDA is used to model reviews contents into topics that occurs in a review. Intuitively, a review can talk about the *service* of a *Japanese* restaurant, and another review refers to the *food* of a *local* restaurant. Both reviews lead to different topics if the space of topics is large enough to differentiate among them.

4.3. Collaborative branch

The collaborative branch of the review-based recommender system aims to gather the opinion of the crowd from the posted reviews. Recommendation is done based on the opinion that the close neighborhood has about a specific item. Thus, we define the utility function f for the seen items $r_{u',i}$ to be the sentiment indicator that comes from the sentiment analysis of a text review that a user wrote about an item. For the unseen items, the heuristic proposed is an averaging over the rates of those similar users (“neighbors”) who have seen the item. Formally,

$$f_{col} = r_{u,i} = \frac{1}{N} \sum_{u' \in N_u} r_{u',i} \quad (4)$$

where N_u represent the N users in the neighborhood of u . Determining N and the neighborhood is key to achieve a good rate estimate for unseen items.

4.3.1. Neighborhood Identification

The neighborhood identification plays a key role to achieve an accurate prediction of the recommender rate. Recommender systems have traditionally used the similarity among users from the matrix of user-item ratings (Diaby et al., 2013). Similarity usually takes into account pairs of ratings that both users have rated. Similarity measurement has been usually calculated by means of a heuristic mechanism such as the Pearson correlation coefficient or cosine-based similarity (Adomavicius and Tuzhilin, 2005).

We propose to use the cosine-based similarity for the review-based recommender collaborative branch, as the heuristic measurement of similarity of two pairs of user vectors \vec{u}_x, \vec{u}_y . This is

$$sim(u_x, u_y) = cos(u_x, u_y) = \frac{\vec{u}_x \vec{u}_y}{||\vec{u}_x|| ||\vec{u}_y||} \quad (5)$$

The user vectors are defined from sentiment ratings over all items. With this approach, pairs of users who have both rated positively or negatively pairs of items, share similar opinions for the unseen items. The cosine-based similarity becomes a useful metric to identify them.

Lately, social recommender systems have introduced novel mechanism to identify the neighbor for collaborative filtering. Groh and Ehmig (2007) propose to use social friendships or relationship to generate the neighborhood for a collaborative recommender. According to them, this approach could outperform the traditional collaborative filtering.

4.3.2. Review sentiment models

Text reviews have also been preprocessed using the aforementioned techniques before applying the following supervised sentiment models. These models have been trained with tagged reviews data and then used to predict the sentiment for the out-of-sample reviews. In fact, we have used three different data sets to build, test and validate the sentiment classifiers:

- A training data set. It is composed of 1500 text reviews. These reviews are tagged using AFINN (Nielsen, 2011), a list of english words rated for valence with an integer between minus five (negative) and plus five (positive). The rating is discretized in three values -1, 0 and 1 via a modified version of sign function, indicating the negative, neutral and positive reviews respectively.
- A testing data set. It is composed of 500 text reviews which also tagged using the AFINN approach.
- A validation data set. It is composed of 200 text reviews that we have manually tagged as positive, neutral or negative reviews.

Bernoulli Naive Bayes sentiment classifier. A Naive Bayes classifier considers a probabilistic model with naive independence assumptions between the features to predict the probability distribution of a sample over the set of classes. We aim to simply model the review sentiment by using a Naive Bayes classifier which considers a bag-of-words model for the review text. This means that we consider words as the model features without taking their position in the review into account and assuming independence among words of a given class (positive, negative or neutral). This can be achieved by using the Bernoulli Naive Bayes model, which defines features as independent binary inputs describing whether or not a word occurs in a given review (see McCallum et al. (1998) for a clarification on the different types of Naive Bayes classifiers).

Multinomial Logistic Regression sentiment classifier. Since we aim to estimate three possible outcomes for each review, positive, negative and neutral reviews, a multi-class classifier is needed. The Multinomial Logistic regression generalizes the logistic regression to a multi-class setting (Kim et al., 2002). This classifier estimates the probabilities of the possible outcomes based on a set of features. It differs from the Naive Bayes classifier in that there is no need for statistical independence on the set of features or words used in the multinomial logit regression. A drawback, compared with Naive Bayes, is the fact that estimating the regression coefficients from a Multinomial classifier is much complex and generally requires an iterative process.

Supervised Latent Dirichlet Allocation. The SLDA classifier for sentiment analysis extends the LDA topic model with response variables for each document (Blei and McAuliffe, 2008). The response variable for a sentiment

analysis classifier corresponds to the sentiment class: positive, negative or neutral. Documents and their responses are jointly modeled to find the latent topics that best predict the responses for future unlabeled documents.

4.4. Hybrid set-up

Using a hybrid linear combination of the collaborative and content-based branches, we pretend to minimize the drawbacks of each individual approach. The hybrid model can be represented as follows, with $\alpha > 0$,

$$f : U \times I \rightarrow R \implies f : f_{col} + \alpha f_{cont} \quad (6)$$

On the one hand, collaborative approaches highly depend on the availability of a critical mass of users who have rated enough items to effectively predict unseen items. One way to overcome *the sparsity of the user-item rating matrix* is to jointly use profiling data to identify the close neighborhood. Our approach minimize the sparsity problem by combining the rates from the content-based branch when the collaborative rates are not accurate enough.

On the other hand, content-based approaches basically fail to recommend items whose features have not been rated yet by the user. In other words, the content-based recommender is *overspecialized* to the features from the seen items since user profiling arises from the items already seen by the user. Item diversity is added into our system to overcome this overspecialization. By combining the collaborative sentiment into the proposed system, unseen item features are also recommended enhancing the overall accuracy.

There also exist drawbacks which are common to both approaches, such as the cold start problem (Lika et al., 2014). For example, the *new user problem* happens when a user has not reviewed any items or very few. In this situation, the recommender is not capable of performing a recommendation and this lowers the system coverage. The *new item problem* is also a drawback difficult to minimize in both set-ups. It occurs when an item is new into the system and it has not yet been reviewed by any user. Since no user feedback exist for these items, the system is not capable of performing an accurate recommendation, also impacting the system coverage.

These two drawbacks are mainly due to the lack of data and hence their solution necessarily undergoes through the process of effectively crawling more data from users and items.

5. System evaluation

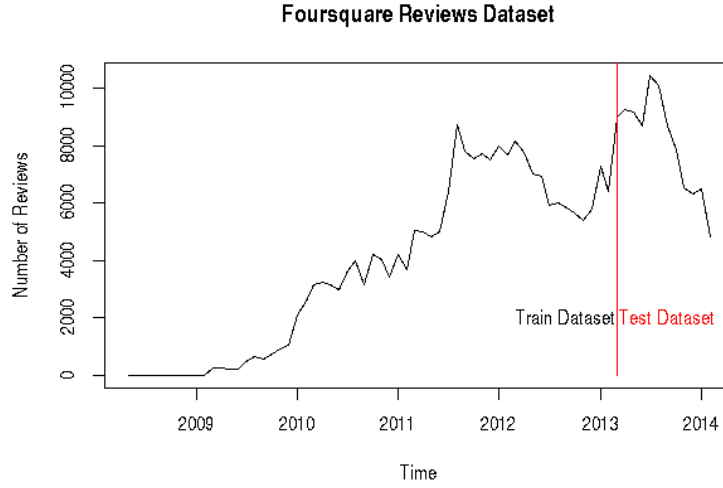
The evaluation methodology will measure how effectively *GeoSRS* recommends an unseen item to a given user. It will then make use of the social network to assess the performance accuracy of the system by comparing the actual user feedback about an item against the potential recommendation rate for that item. Hence, we assume that there exists some degree of causality between the fact of purchasing/experiencing an item and the subsequent action of reviewing the product/experience.

This approach also provides a simple but powerful mechanism to measure the recommendation coverage, or simply, the proportion of recommendations that the system is able to output. The inability to make a recommendation can happen due to the facts that either the user provides feedback for the first time –*new user problem*– or the item has not yet been reviewed –*new item problem*. For example, a user who checks into a Foursquare venue like a restaurant to have some food, would probably provide its own feedback in form of a review at the social network. The interpretation of this feedback (positive, negative or neutral) is used in our evaluation approach to be compared against the recommendation rate from *GeoSRS* at the reviewing time and within the neighborhood of that restaurant.

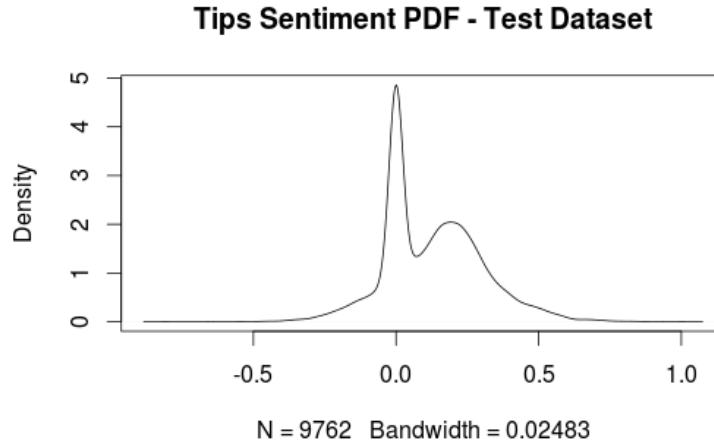
Next, we describe the Foursquare data set that is later used to assess *GeoSRS* in terms of recommendation coverage and performance accuracy.

5.1. The Foursquare Restaurant Tips Data set

The Foursquare restaurant tips data set consists of 309,640 short reviews or tips from the Manhattan region. The whole data set is split into training (70% oldest tips) and test set (30% newest tips).



(a) Number of Foursquare restaurant reviews over time.



(b) Tips Sentiment Probability Density Function

Figure 5: Description of the Foursquare dataset used in our evaluation.

In Figure 5a, we show the number of reviews as a function of time, from the first review in 2008 until the last crawling execution in February 2014. Furthermore, this graphic shows a positive tendency in the number of reviews created since 2008, what indicates us that Foursquare has been growing since

then, despite some seasonal patterns. As it has been said, the whole data set is split into two subsets, which are also indicated in the plot.

Figure 5b shows the distribution of sentiment polarity of tips in the test dataset. While negative values in the pdf express negative sentiment, positive sentiment is contained in the interval $(0, 1]$. On the other hand, neutral user feedback is represented by a sentiment value of zero. For evaluation purposes, we consider neutral feedback as a positive experience since most of them are statements such as “*Try the hot chocolate!*”; we make the assumption that the user implicitly rates the experience as positive.

We also note that the test data set is unbalanced in terms of positive and negative ratings, reflecting the reality of Foursquare tips data from early 2013 till 2014. This might not be extremely relevant when assessing the performance accuracy of this data set, but it is relevant when comparing different recommender set-ups and the generalization capabilities.

5.2. Recommendation Coverage

Recommendation coverage measures the domain of items over which the system can perform recommendations (Ge et al., 2010). Typically, the term coverage has been associated with (1) the percentage of the items for which the system is able to generate a recommendation and (2) the percentage of the available items which are effectively recommended. Here, we adopt the former definition since we gather the latter in the concept of performance accuracy defined later. Thus, formally, *Coverage* is defined as the percentage of available items (I) for which the recommender system can output a prediction (I_P) i.e. $Cov = (|I_P| / |I|) \cdot 100$.

In *GeoSRS* and within the proposed evaluation scheme, I_P only depends on the presence or absence of historical data (tips) in the training data set for an item and user from the test data set. Given that the content-based branch outputs a rate always that a given item and user have more than one tip, the overall recommendation coverage does not depend on the collaborative branch.

In Figure 6, we plot the recommendation coverage as a function of time in order to show how the coverage softly increases when the number of tips starts to be large (cf. Figure 5a) and the proportion between new and old items/users stabilizes.

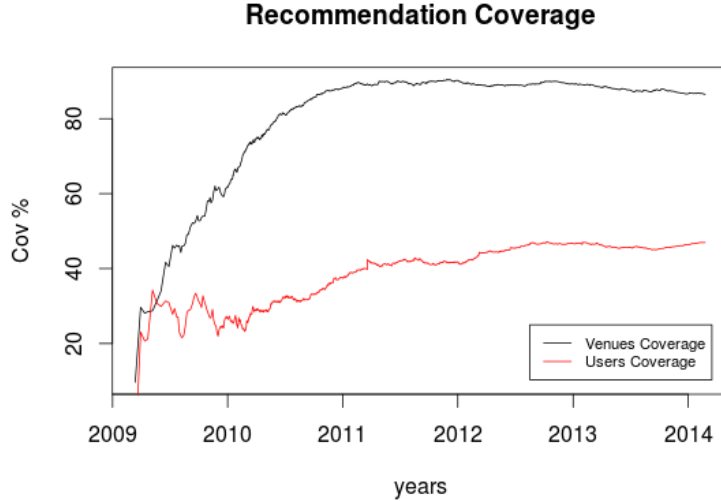


Figure 6: Recommendation coverage for venues and users.

From this picture, we can also see that while there is a good coverage value on restaurants venues (85% in 2014), coverage on users is quite low (45% in 2014). This low recommendation coverage on users is mainly due to two reasons. On the one hand, there are new users who signed up to Foursquare during the test data set time frame (2013-2014); on the other hand, some users might have tipped just once Manhattan restaurants, but they might have tipped restaurants out of Manhattan. This poor recommendation coverage on users could be mitigated by adapting the crawling system to also crawl tips from users, instead of just crawling tips from venues.

5.3. Performance Accuracy

Performance accuracy measures the goodness of the rate prediction against the actual rate (Shani and Gunawardana, 2011; Hernández del Olmo and Gaudioso, 2008). Typically, it is formulated as the number of successful cases over all the cases in the test data set,

$$accuracy = \frac{\# \text{ good cases}}{\# \text{ overall cases}} \quad (7)$$

As a difference with most of the evaluation methodologies in recommender systems, our approach assumes that a successful recommendation is the one

that leads the user to a positive experience. The positiveness of the experience is measured through the user feedback in the review, while the recommendation is generated by the engine.

Since the overall number of cases can be decomposed into positive and negative cases, the accuracy can be also expressed in terms of the average error rate, $\bar{\epsilon}$, as,

$$accuracy = 1 - \frac{\# \text{ negative cases}}{\# \text{ overall cases}} = 1 - \bar{\epsilon} \quad (8)$$

where ϵ_k is the recommendation error rate for the k th experiment, defined as follows:

$$\epsilon_k = \begin{cases} \frac{N_{pos_k}}{N_k} & \text{Sentiment is **positive**} \\ 1 - \frac{N_{pos_k}}{N_k} & \text{Sentiment is **negative**} \end{cases} \quad (9)$$

where N_k is the number of items recommended in the k experiment and N_{pos_k} , the position in a ordered list that occupies the item that the user has experienced.

Notice that if a review is positive, then ϵ_k is low if the item is ranked at the top of the recommendation list, while if a review is negative, ϵ_k is low if the item is ranked at the bottom of the list.

5.3.1. Collaborative-based model results

The proposed collaborative models gather the opinion of the crowd by averaging the feedback from users whose taste is similar to the recommended user, a.k.a. neighborhood. As introduced earlier, the N neighbors are identified through the k -nearest neighbor algorithm that uses the cosine similarity of the historical user experiences.

First, we plot the Area Under the Curves (AUC) of the recommendation error cumulative density functions (cdf) which let us interpret the goodness of the models with the available Foursquare data sets.

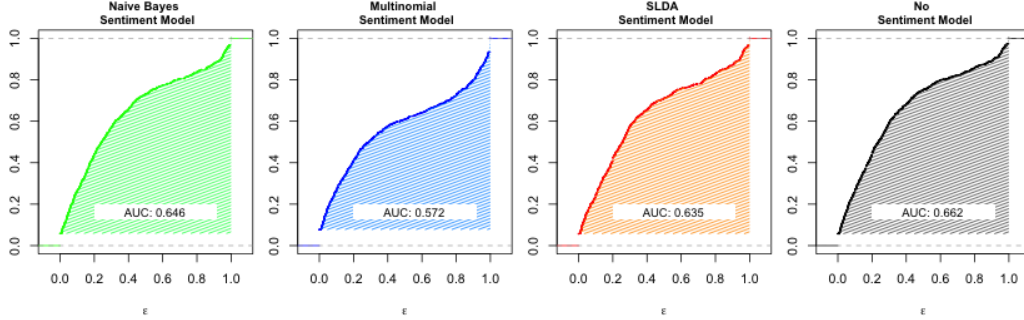


Figure 7: AUCs from the recommendation error cdf for different sentiment models.

As Figure 7 suggests Naive Bayes outperforms the Multinomial and SLDA models. Even more relevant is the fact that using no sentiment performs better than using a sentiment model. The collaborative model without sentiment considers all tips as positive instead of using the sentiment model which extracts the opinion from the reviewer.

However, the fact of using an unbalanced data set with more positive than negative experiences, benefits those models that are biased towards positive ratings. As shown in the table below, Naive Bayes, and obviously the no-sentiment model, have bigger proportions of positive ratings in the training data sets (because the trained sentiment classifier was generated in this way) and they perform better in test data set which has greater number of positive experiences.

	<i>SLDA</i>	<i>Multinomial</i>	<i>Naive Bayes</i>	<i>No Sentiment</i>
# positivie ratings	191.809	142.758	205.585	212.947
# negative ratings	21.138	70.189	7.362	0

Table 2: Number of positive and negative predicted ratings

In order to mitigate the unbalancing effects described above, we propose to subsample the testing data set in two balanced data sets and then compute the performance accuracy. In order to estimate properly the cumulative density function and the AUC statistic, we repeat the sub-sampling process hundreds of times and average the results. The results of this assessments are shown in Figure 8 for each of the sentiment models.

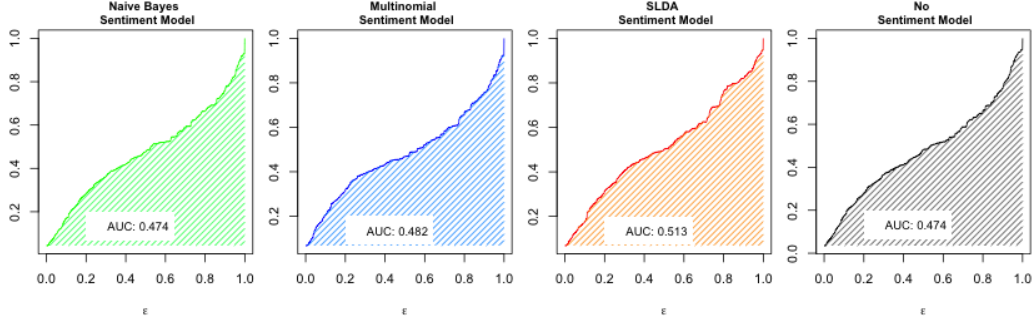


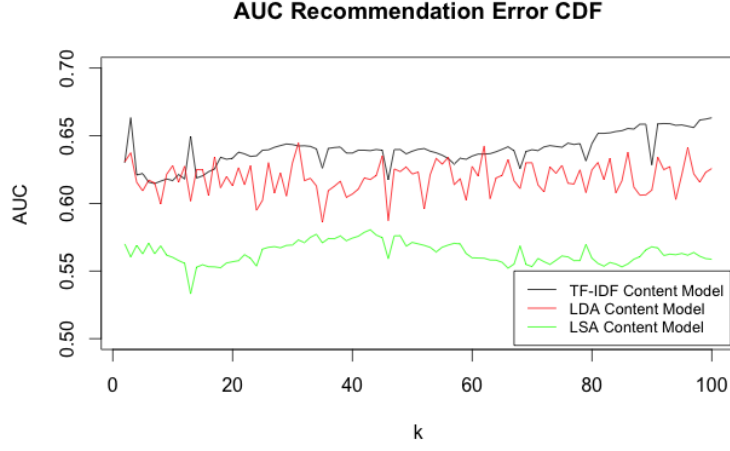
Figure 8: AUCs from the recommendation error cdf for different sentiment models in a balanced data set.

The results of this analysis show that the sentiment models behaves better than the collaborative model without sentiment, except for Naive Bayes which behaves similarly. Furthermore, we realize that the SLDA models, which takes into account the word polysemy, performs also better than others when used in a collaborative recommender.

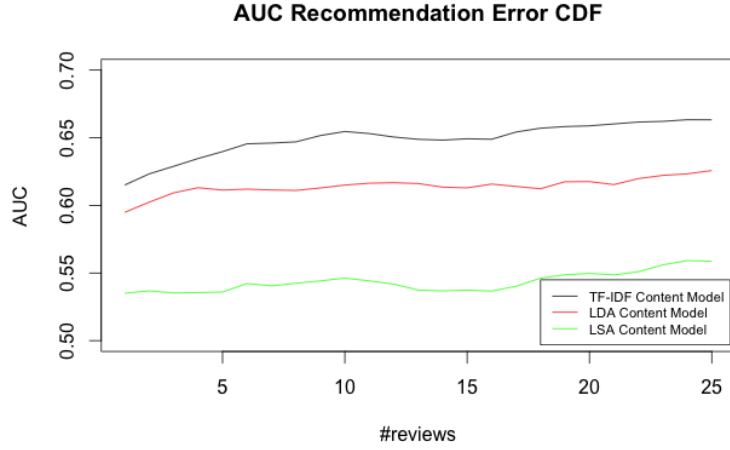
5.3.2. Content-based model results

Content-based models cope with the user tastes, which are represented by features extracted from the tips data. We have presented earlier three state-of-the-art text models to represent the user tastes into data features. Note that the content model gathers the a priori knowledge about a restaurant. In other words, it takes into account the similarity between a user and a venue based on their preferences and traits, respectively. Consequently, we assess the content model against the fact that user has gone or not to the recommended venue, more than evaluating the experience on it, since we understand that the goal of the content-based model is to match similar profiles more than taking into account the global opinion about a restaurant.

In what follows, we compare the performance accuracy of these three models by examining the AUC of the recommendation error cdf. In this section, the recommendation error is defined as if all tips in the test data set were positive experiences.



(a) AUC of the cdc as a function of the number of features.



(b) AUC of the cdf as a function of the number of reviews

Figure 9: Performance of the content-based branch.

Figure 9a shows the performance accuracy as a function of the feature space size. Note that, in TF-IDF, features are keywords; in LSA, semantic concepts; and in LDA, document topics. Although the complexity that entails LSA and LDA models, the simplicity of TF-IDF content model tends to generate more accurate recommendations for feature space size ranging from

2 to 100 features, since it outputs a greater AUC and hence less recommendation error.

Despite the noisiness of these curves, clearly related to noisy text reviews, there is a tendency that states that the greater the number of features, the more AUC and the better performance accuracy is. However, the feature space size directly impacts the computation time and hence the velocity to make a recommendation. From now on, we use $k = 100$ for all three models, which guarantees a fair trade off between performance accuracy and computation cost. The different cdf and their AUC for each of the content models with $k = 100$ features can be seen in Figure 10.

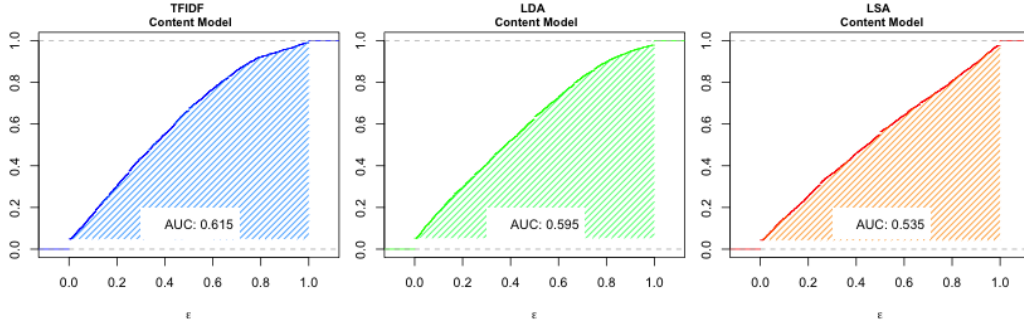


Figure 10: AUC from the cdf of the recommendation error

One of the aspects that could affect the performance accuracy of the content-based models is the number of reviews or tips per venue and user. As shown in Figure 9b, the performance accuracy increases with the number of reviews per venue and user. The graphic plots the AUC for observations of users and venues in the test data set with more than a given number of reviews. Consequently, it says that the more knowledge we have about the venue and the user, the better we can estimate its recommendation. Note also that TF-IDF outperforms the LSA and LDA for any number of tips.

In general, it can be interpreted that not only the number of the reviews influence the performance accuracy and coverage of a review-based recommendation system, but other features about the data such as the quality of the reviews, the heterogeneity of topics reviewed, the influence of the reviewer, amongst others.

As a closing remark for this section we point out that we believe our selection of metrics for capturing the quality of our rankings are adequate for the task and data at hand. The reader should note that we are not

predicting ratings, and so more common measures such as MAE or RMSE are not appropriate in our context. Please refer to (Adomavicius and Tuzhilin, 2005; Lü et al., 2012) for more in depth discussion on the topic of classification accuracy metrics for ranking systems.

5.3.3. *GeoSRS: our hybrid recommender system*

Finally, we show results of *GeoSRS*, obtained by linearly combining both content-based and collaborative approaches, as discussed in Section 4.1. Our aim in this section is twofold. Firstly, we want to optimize the coefficient α from Eq. (6) that minimizes the recommendation error. In other words, we seek to find α values that maximizes the AUC of the recommendation error cumulative density function. Hence, we simulate several recommendations for different α value and computed its associated AUC to find the one that maximizes it. Secondly, we aim to show that the hybrid approach consistently outperforms each individual component.

To globally assess the recommender, we consider again the evaluation model with the testing experiences classified as positive or negative. The SLDA sentiment model is chosen since it scored higher in the balanced evaluation test. Furthermore, TF-IDF is chosen as the content-based model to be included in this hybrid approach.

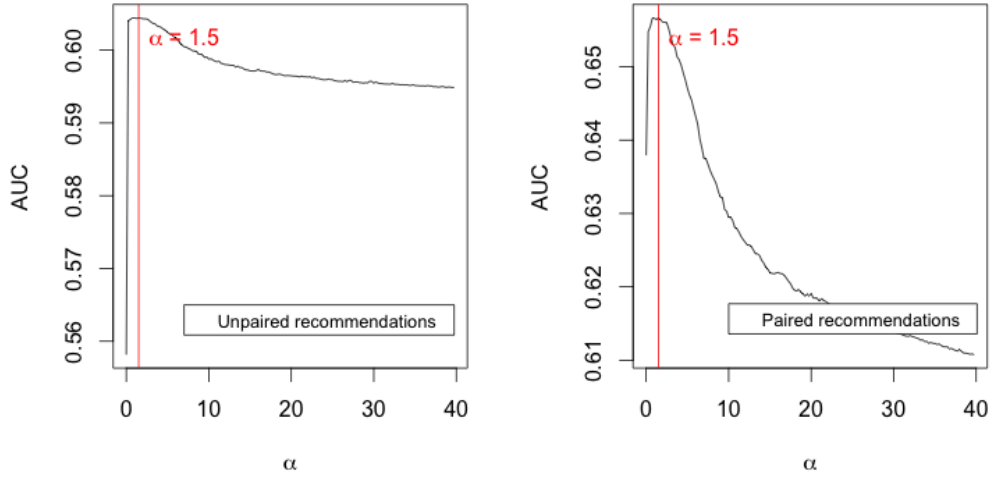


Figure 11: Hybrid Linear Combiner Optimization

Figure 11 plots two curves, both representing the AUC of *GeoSRS* for

different values of α ranging from 0 to 20, see Eq.(6). Notice that the collaborative branch does not always output predictions due to lack of close neighbors to the given user. The left-most plot shows the performance of *GeoSRS* for all reviews in the test data set, whereas the right-most plot shows the performance of *GeoSRS* reviews that can be covered (make a recommendation) by both system branches (“paired” recommendations). As the graphic shows, an α equal to 1.5 maximizes AUCs from both scenarios. Note also that $\alpha = 0$ corresponds to the collaborative recommender, while $\alpha \rightarrow \infty$ corresponds to the content-based recommender.

Next, we plot the empirical cdf for both situations with $\alpha = 1.5$. Figure 12a shows the performance accuracy for paired recommendations. That the hybrid system has greater AUC can also be clearly seen in Table 3.

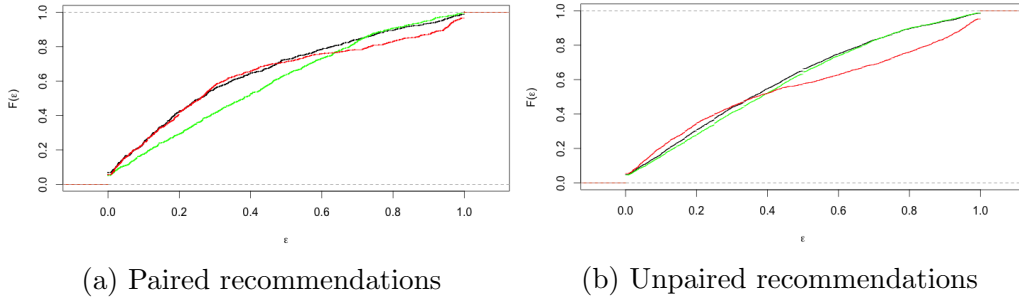


Figure 12: Empirical cumulative density functions for Content-based (green), Collaborative (red) and Hybrid (black) set-ups

Figure 12b plots the empirical cumulative density function for the real Foursquare data set dominated by unpaired recommendations, what this means is that recommendations do not necessarily have a collaborative prediction for each content-based. Particularly, we observe that in our situation for every 5 content-based recommendations, there is only one collaborative recommendation. This has a huge impact into the overall system evaluation because unpaired recommendations weights more than paired recommendations and the content-based branch seems to perform more accurately. As shown in Figure 12b, the Hybrid curve is now closer to the content-based curve than to the collaborative which is much lower because it cannot output an appropriate recommendation for those users without neighborhood.

By exposing the results split into two separate data sets (paired and unpaired recommendations), the fact of using a linear model always guarantees

greater performance. The hybrid recommender performance benefits from either when the collaborative branch can work out a rating or when it cannot. Nonetheless, the biggest contributions happens when the collaborative branch can recommend.

It is relevant to say that restricting the recommender just to those users with a collaborative prediction (paired recommendation) has a huge impact to the overall recommendation coverage since the collaborative branch has low coverage due to the high sparsity in commonly rated venues. Table 3 shows the relationship between the performance accuracy in terms of AUC for different system configurations and the recommendation coverage in each data set situation mentioned above.

	Hybrid	Content-based	Collaborative	Coverage
# Paired rec.	0.6566	0.5984	0.6351	8.34%
# Unpaired rec.	0.6044	0.5922	0.5514	39.81%

Table 3: AUC and coverage for different set-ups of the review-based recommender

For the sake of comparison, we implement two state-of-the-art hybridization techniques that have been shown to work specially well when employing two components of different strengths (for example, collaborative and content-based) named meta-level and cascade (Burke, 2007).

Regarding the meta-level setup, we build a collaborative through content approach similar to (Pazzani, 1999), which uses the content-based recommendation matrix to identify the close neighborhood of each user under evaluation. The unseen rates are estimated by means of averaging their neighbor’s sentiment rates, computed beforehand through the sentiment analysis model. On the contrary, the cascade configuration builds a strictly hierarchical hybrid which first applies the stronger branch, in our scenario this is the collaborative, and then it uses the content-based branch for those items which the collaborative could not decide well. A very similar system was proposed by (Burke, 2002) for restaurant recommendation, but their hybrid system relied on a knowledge-based in place of a content-based recommender branch.

Fig. 13 shows the empirical density curves of the error for weighted, meta-level and cascade hybrids. As it is shown, the simple weighted hybrid system used in *GeoSRS* outperforms the other two enabling a simple but rather powerful hybridization technique.

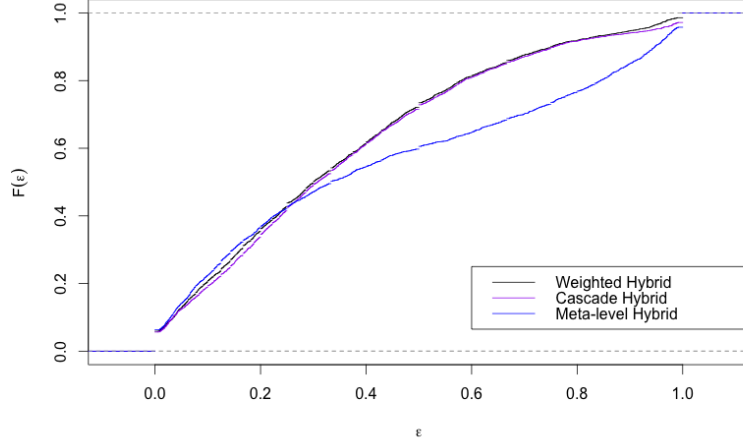


Figure 13: Comparison against state-of-the-art hybridization techniques

6. Conclusions and future work

GeoSRS enables social network users to mitigate the information overload problem by digging into text reviews data and recommending personalized items according to their preferences based on their past reviews. *GeoSRS* supposes a novel attempt to purely integrate text reviews data into a working recommending system.

This paper assesses different *GeoSRS* set-ups which use several state-of-the-art text mining techniques. According to the proposed offline evaluation approach, the results justify combining text review content and sentiment into an hybrid recommending engine. This is particularly relevant in balanced and paired data sets. Concretely, our results show that for our data set the best configuration is given by TF-IDF for the content-based branch and SLDA for the collaborative branch with a coefficient of $\alpha = 1.5$. Moreover, the evaluation outcomes point out the importance of the data retrieval process and the quality of the data set, by showing that the larger the number of reviews, the greater the performance accuracy. Not least, we have also shown the benefits of using a simple hybridization technique like the weighted linear combination compared to employing more complex techniques such as cascade or meta-level.

Within this evaluation scheme and *GeoSRS*, we show the goodness of

recommendation coverage is tied to the fact of having historical reviews for all users and all items. We observe that the growth of the social network activity over the years directly improves this evaluation figure, but we also state that the more retrieved data, the better the coverage. We speculate that the quality of our recommendations should improve as the number of reviews increases over time.

In order to directly impact the performance accuracy and recommendation coverage of the *GeoSRS*, this article also motivates the use of *Quadtree* algorithm to efficiently retrieve geolocated data, such as reviews from Foursquare. A parallel version of *Quadtree* is detailed in this article with the aim to effectively crawl data from large urban areas. As a result, *GeoSRS* together with the *Quadtree* crawling process become key components of a recommending system capable of working in large data sets of reviews for large urban areas.

Finally, we can conclude that *GeoSRS* results, as in many other recommender systems, highly depend on the availability and quality of the data. We have shown that by balancing data sets or making paired recommendations in the hybrid set-up, the outcome of the system could differ. Because of this, we deeply encourage *GeoSRS* to undergo through online evaluation tests, such as randomized experiments or A/B testing, in order to mitigate the skewness on the Foursquare test data sets and get more reliable results.

Acknowledgements. The authors gratefully acknowledge the extensive comments of the anonymous reviewers that lead to an improvement of the original manuscript.

References

- Aciar, S., Zhang, D., Simoff, S., Debenham, J., 2007. Informed recommender: Basing recommendations on consumer product reviews. *IEEE Intelligent Systems* 22 (3), 39–47.
- Adomavicius, G., Tuzhilin, A., Jun. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Trans. on Knowl. and Data Eng.* 17 (6), 734–749.
- Andersen, R., Borgs, C., Chayes, J., Feige, U., Flaxman, A., Kalai, A., Mirrokni, V., Tennenholtz, M., 2008. Trust-based recommendation systems:

- an axiomatic approach. In: Proc. of the 17th Int. Conf. on World Wide Web (Beijing). ACM, pp. 199–208.
- Berjani, B., Strufe, T., 2011. A recommendation system for spots in location-based online social networks. In: Proc. of the 4th Workshop on Social Network Systems (Salzburg). ACM, p. 4.
- Blei, D., McAuliffe, J., 2008. Supervised topic models. In: Platt, J., Koller, D., Singer, Y., Roweis, S. (Eds.), *Advances in Neural Information Processing Systems* (Vancouver). Vol. 20. MIT Press, pp. 121–128.
- Blei, D. M., Ng, A. Y., Jordan, M. I., Mar. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.* 3, 993–1022.
- Böhning, D., 1992. Multinomial logistic regression algorithm. *Annals of the Institute of Statistical Mathematics* 44 (1), 197–200.
- Burke, R., 2002. Hybrid recommender systems: Survey and experiments. *User modeling and user-adapted interaction* 12 (4), 331–370.
- Burke, R., 2007. Hybrid web recommender systems. In: *The adaptive web*. Springer, pp. 377–408.
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., Harshman, R., 1990. Indexing by latent semantic analysis. *J. of the American Soc. for Information Sci.* 41 (6), 391–407.
- Diaby, M., Viennet, E., Launay, T., 2013. Toward the next generation of recruitment tools: An online social network-based job recommender system. In: Proc. of the 2013 IEEE/ACM Int. Conf. on Advances in Social Networks Analysis and Mining (Niagara Falls). ASONAM '13. pp. 821–828.
- Falahi, K., Mavridis, N., Atif, Y., 2012. Social networks and recommender systems: A world of current and future synergies. In: Abraham, A., Hasanien, A.-E. (Eds.), *Computational Social Networks*. pp. 445–465.
- Gartrell, M., Xing, X., Lv, Q., Beach, A., Han, R., Mishra, S., Seada, K., 2010. Enhancing group recommendation by incorporating social relationship interactions. In: Proc. of the 16th ACM Int. Conf. on Supporting Group Work (Sanibel Island). GROUP '10. pp. 97–106.

- Ge, M., Delgado-Battenfeld, C., Jannach, D., 2010. Beyond accuracy: Evaluating recommender systems by coverage and serendipity. In: Proc. of the Fourth ACM Conf. on Recommender Systems (Foster City). RecSys '10. pp. 257–260.
- Groh, G., Ehming, C., 2007. Recommendations in taste related domains: Collaborative filtering vs. social filtering. In: Proc. of the 2007 Int. ACM Conf. on Supporting Group Work (Sanibel Island). GROUP '07. pp. 127–136.
- Guy, I., Carmel, D., 2011. Social recommender systems. In: Proc. of the 20th Int. Conf. Companion on World Wide Web (Hyderabad). WWW '11. pp. 283–284.
- Hernández del Olmo, F., Gaudioso, E., 2008. Evaluation of recommender systems: A new approach. Expert Syst. Appl. 35 (3), 790–804.
- Huang, J., Cheng, X.-Q., Guo, J., Shen, H.-W., Yang, K., 2010. Social recommendation with interpersonal influence. In: Proc. of the 2010 Conf. on ECAI 2010: 19th European Conf. on Artificial Intelligence. pp. 601–606.
- Jakob, N., Weber, S. H., Müller, M. C., Gurevych, I., 2009. Beyond the stars: Exploiting free-text user reviews to improve the accuracy of movie recommendations. In: Proc. of the 1st Int.l CIKM Workshop on Topic-sentiment Analysis for Mass Opinion (Hong Kong). TSA '09. pp. 57–64.
- Jamali, M., Ester, M., 2009. Trustwalker: a random walk model for combining trust-based and item-based recommendation. In: Proc. of the 15th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining (Paris). ACM, pp. 397–406.
- Jameson, A., 2004. More than the sum of its members: Challenges for group recommender systems. In: Proc. of the Working Conference on Advanced Visual Interfaces (Gallipoli). AVI '04. pp. 48–54.
- Kim, S.-B., Rim, H.-C., Yook, D., Lim, H.-S., 2002. Effective methods for improving Naive Bayes text classifiers. In: PRICAI 2002: Trends in Artificial Intelligence. Springer, pp. 414–423.
- Kim, S.-M., Pantel, P., Chklovski, T., Pennacchiotti, M., 2006. Automatically assessing review helpfulness. In: Proc. of the 2006 Conf. on Empirical Methods in Natural Language Processing. EMNLP '06. pp. 423–430.

- Lika, B., Kolomvatsos, K., Hadjiefthymiades, S., 2014. Facing the cold start problem in recommender systems. *Expert Systems with Applications* 41 (4, Part 2), 2065 – 2073.
- Liu, X., Aberer, K., 2013. SoCo: A social network aided context-aware recommender system. In: *Proc. of the 22nd Int. Conf. on World Wide Web. WWW '13*. pp. 781–802.
- Liu, Y., Huang, X., An, A., Yu, X., 2008. Modeling and predicting the helpfulness of online reviews. In: *Eighth IEEE Int. Conf. on Data Mining, ICDM '08*. pp. 443–452.
- Lü, L., Medo, M., Yeung, C. H., Zhang, Y.-C., Zhang, Z.-K., Zhou, T., 2012. Recommender systems. *Physics Reports* 519 (1), 1–49.
- Lu, Y., Tsaparas, P., Ntoulas, A., Polanyi, L., 2010. Exploiting social context for review quality prediction. In: *Proc. of the 19th Int. Conf. on World Wide Web. WWW '10*. pp. 691–700.
- Ma, H., Yang, H., Lyu, M. R., King, I., 2008. SoRec: Social recommendation using probabilistic matrix factorization. In: *Proc. of the 17th ACM Conf. on Information and Knowledge Management. CIKM '08*. pp. 931–940.
- Massa, P., Avesani, P., 2007. Trust-aware recommender systems. In: *Proc. of the 2007 ACM Conf. on Recommender systems. ACM*, pp. 17–24.
- McCallum, A., Nigam, K., et al., 1998. A comparison of event models for Naive Bayes text classification. In: *AAAI-98 Workshop on Learning for Text Categorization. Vol. 752*. pp. 41–48.
- Mobasher, B., Jin, X., Zhou, Y., 2004. Semantically enhanced collaborative filtering on the web. In: *Web Mining: From Web to Semantic Web*. Springer, pp. 57–76.
- Moghaddam, S., Jamali, M., Ester, M., Habibi, J., 2009. Feedbacktrust: Using feedback effects in trust-based recommendation systems. In: *Proc. of the 3rd ACM Conf. on Recommender Systems. RecSys '09*. pp. 269–272.
- Mokbel, M. F., Bao, J., Eldawy, A., Levandoski, J. J., Sarwat, M., 2011. Personalization, socialization, and recommendations in location-based services 2.0. In: *PersDB 2011 Workshop, Seattle, Washington, USA*.

- Musat, C.-C., Liang, Y., Faltings, B., 2013. Recommendation using textual opinions. In: Proc. of the 23rd Int. Joint Conf. on Artificial Intelligence. AAAI Press, pp. 2684–2690.
- Nielsen, F. Å., 2011. A new ANEW: Evaluation of a word list for sentiment analysis in microblogs. In: Proc. of the ESWC2011 Workshop on 'Making Sense of Microposts': Big things come in small packages. pp. 93–98.
- O'Doherty, D., Jouili, S., Van Roy, P., et al., 2012. Trust-based recommendation: An empirical analysis. In: Sixth ACM Workshop on Social Network Mining and Analysis (SNA-KDD 2012), Aug. 12, 2012.
- Pazzani, M. J., 1999. A framework for collaborative, content-based and demographic filtering. *Artificial Intelligence Review* 13 (5-6), 393–408.
- Resch, B., Zipf, A., Beinat, E., Breuss-Schneeweis, P., Boher, M., 2012. Towards the live city: paving the way to real-time urbanism. *International Journal on Advances in Intelligent Systems* 5 (3-4), 470–482.
- Reschke, K., Vogel, A., Jurafsky, D., 2013. Generating recommendation dialogs by extracting information from user reviews. In: *ACL* (2). pp. 499–504.
- Salton, G., Buckley, C., 1988. Term-weighting approaches in automatic text retrieval. *Information Processing & Management* 24 (5), 513–523.
- Samet, H., Jun. 1984. The quadtree and related hierarchical data structures. *ACM Comput. Survey* 16 (2), 187–260.
- Sarwt, M., Levandoski, J. J., Eldawy, A., Mokbel, M. F., 2013. Lars*: An efficient and scalable location-aware recommender system. *IEEE Transactions on Knowledge and Data Engineering* 99.
- Shani, G., Gunawardana, A., 2011. Evaluating recommendation systems. *Recommender Systems Handbook*, 257–297.
- Snyder, B., Barzilay, R., 2007. Multiple aspect ranking using the good grief algorithm. In: Proc. of the Human Language Tech. Conf. of the North American Chap. of the Assoc. of Comput. Linguistics (HLT-NAACL). pp. 300–307.

- Sparck Jones, K., 1972. A statistical interpretation of term specificity and its application in retrieval. *J. of Documentation* 28, 11–21.
- Symeonidis, P., Ntempos, D., Manolopoulos, Y., 2014a. Location-based social networks. In: *Recommender Systems for Location-based Social Networks*. Springer, pp. 35–48.
- Symeonidis, P., Ntempos, D., Manolopoulos, Y., 2014b. *Recommender Systems for Location-based Social Networks*. Springer.
- Tu, K., Ribeiro, B., Jensen, D., Towsley, D., Liu, B., Jiang, H., Wang, X., 2014. Online dating recommendations: matching markets and learning preferences. In: *Proc. of the Companion Pub. of the 23rd Int. Conf. on World Wide Web Companion*. pp. 787–792.
- Xia, F., Asabere, N. Y., Liu, H., Deonauth, N., Li, F., 2014. Folksonomy based socially-aware recommendation of scholarly papers for conference participants. In: *Proc. of the 23rd Int. Conf. on World Wide Web Companion*. pp. 781–786.
- Yang, D., Zhang, D., Yu, Z., Wang, Z., 2013. A sentiment-enhanced personalized location recommendation system. In: *Proc. of the 24th ACM Conference on Hypertext and Social Media*. ACM, pp. 119–128.
- Ye, M., Yin, P., Lee, W.-C., 2010. Location recommendation for location-based social networks. In: *Proc. of the 18th SIGSPATIAL Int. Conf. on Advances in Geographic Information Systems*. ACM, pp. 458–461.
- Zhang, Z., Varadarajan, B., 2006. Utility scoring of product reviews. In: *Proc. of the 15th ACM Int. Conf. on Information and Knowledge Management*. CIKM '06. pp. 51–57.
- Zheng, Y., 2012. Tutorial on location-based social networks. In: *Proc. of the 21st Int. Conf. on World Wide Web (WWW)*. ACM.
- Zheng, Y., Zhang, L., Ma, Z., Xie, X., Ma, W.-Y., 2011. Recommending friends and locations based on individual location history. *ACM Trans. Web* 5 (1), 5:1–5:44.