Social Recommender Systems

Ido Guy

1 Introduction

In the “social web” or “Web 2.0” [79], people play a central role by creating content, annotating it with tags, votes (or ‘likes’), or comments, joining communities, and connecting with friends and followers. *Social media* websites are proliferating and attract millions of users who author content, post messages, share photos with their friends, and engage in many other types of activities. This rapid growth intensifies the phenomenon of *social overload*, where users of social media are exposed to a huge amount of information and participate in vast amounts of interactions. Social overload makes it harder on the one hand for social media users to choose which sites to engage in and for how long and on the other hand makes it more challenging for social media websites to attract users and retain them.

Social Recommender Systems (SRS) are recommender systems that target the social media domain. They aim at coping with the social overload challenge by presenting the most relevant and attractive data to the user, typically by applying personalization techniques. The “marriage” between recommender systems (RS) and social media has many potential benefits for both sides. On the one hand, social media introduces many new types of data and meta-data, such as tags and explicit online relationships, which can be used in a unique manner by RS to enhance their effectiveness. On the other hand, recommender systems are crucial for social media websites to enhance the adoption and engagement by their users and thus play an important role in the overall success of social media. It should be noted that traditional RS, such as user-based collaborative filtering, are social in their nature since they mimic the natural process where we seek advice or suggestions from

I. Guy (✉)
Ben-Gurion University of the Negev, Beer-Sheva, Israel
e-mail: idoguy@acm.org

F. Ricci et al. (eds.), *Recommender Systems Handbook*,
https://doi.org/10.1007/978-1-0716-2197-4_22

other people [88]. Yet, in this chapter we focus on those recommender systems that are aimed for the social media domain, which we term *social recommender systems* [43].

This chapter focuses on two key areas of SRS, social media content recommendation and people recommendation. We dedicate a section to each of these areas, reviewing the different sub-domains, their unique characteristics, the applied methods, case studies in the enterprise, and open challenges. SRS consist of more areas, such as recommendation of tags and groups (communities), however, these are left beyond the scope of this chapter. The remainder of the chapter is organized as follows: the next two sections discuss in detail content and people recommendation. The following section discusses key aspects characterizing SRS as raised throughout its preceding two sections. The chapter concludes by reviewing emerging SRS domains and open challenges.

2 Content Recommendation

Social media introduced many new types of content that can be created and shared by any user in a way that has never been possible before. Users became the center of every social media website and in many cases were the ones creating the actual content of the site: textual content as in Wikipedia and WordPress; photos as in Flickr and Facebook; and video as in YouTube. Users also have a key role in providing feedback and annotating existing content on social media websites. Comments allow users to add their own opinion; votes and ratings allow them to 'like' (or dislike) favourite posts; and tags allow them to annotate the content with keywords that reflect their own viewpoint. These new types of feedback forms allow RS to implicitly infer user preferences and content popularity by analyzing the crowd's feedback.

In the social media era, articulated relationships have become available through social network sites (SNSs) [11] and changed the world of content recommendation. While in the past such relationships could only be partially extracted by surveys and interviews, and later by mining communication patterns from phone logs or email that are highly sensitive privacy-wise, the availability of relationships in social networks allows tapping into one's network of familiar people (Facebook, LinkedIn) or people of interest (Twitter) in a simpler way without infringing privacy. The use of the friend list instead of or alongside the list of similar people as in traditional CF has been broadly proven to be productive for enhancing content recommendations. Sinha and Swearingen [97] were among the first to compare friend-based recommendation with traditional methods and showed their effectiveness for movie and book recommendation. Golbeck [36] showed that friends can be a trusted source for movie recommendation. Groh and Ehmig [38] compared collaborative filtering with friend-based "social filtering" and showed the advantage of the latter for club recommendation within a German SNS. Costa and Ortale [20] use signed social relations (trust and distrust) to produce personalized recommendations using a model-based approach. Yang et al. [110] provide a survey of collaborative filtering based SRS, classifying them into matrix factorization-based approaches and neighbourhood-based approaches. Eirinaki et al. [26] survey large-scale recommender systems that take advantage of the characteristics of the underlying social network. They focus on the variety and volatility of social bonds and tackle the problems of size and speed of change in social graphs. Overall, recommendation based on friends enhance recommendations' accuracy; allow the user to better judge the recommendations since s/he is familiar with the respective people; spare the need for explicit feedback from the user in order to calculate similarity; and help cope with the cold-start problem for new users.

In recent years, studies that harness deep learning approaches to content SRS have emerged. Sun et al. [98] suggested an attentive recurrent network-based approach for temporal social recommendation. They model users' complex dynamic and general static preferences over time by fusing social influence among users with two attention networks. Taneja and Arora [101] prioritize contextual dimensions for user modeling using neural networks and tensor factorization. Tahmasebi et al. [100] suggest a hybrid social recommender system utilizing a deep autoencoder network. The proposed approach uses collaborative and content-based filtering, as well as social influence. Fan et al. [27] propose a Bi-LSTM with attention mechanism to extract "deep" social sequences, which consider information from not only direct neighbors but also distant neighbors. Their approach demonstrates good performance over the Ciao and Epinions datasets.

The remainder of this section reviews key domains of social media content recommendation, such as blogs, microblogs, news, and multimedia. We then briefly discuss group recommendation, which is especially relevant for recommendation of social media content. Following, a case study of social media recommendation within the enterprise is presented in detail. The section concludes with a summary of key points.

2.1 Key Domains

Blogs Blogs are one of the classic social media applications and a natural ground for recommendation techniques. They typically consist of inherent hierarchy that SRS need to take into account. At the top of this hierarchy is the blog itself, which may be owned by an individual user or a community, and is often focused on a topic or domain. The blog includes different blog posts (or blog entries) that include one article by a single author. The author (and sometimes other users) can usually annotate the post with appropriate tags, which also serve for dissemination to relevant populations. The post's readers can add comments and can often also vote for (or 'like') the post [50]; other authors can use a trackback to reference the post from their own post. In one of the early studies of blog recommendation, Arguello et al. [5] explored personalized recommendation of whole blogs (as opposed to blog posts) using the TrecBlog06 dataset [72]. Given a query that represented the user's topical interests, two document models were explored: the first included a single large document that was based on concatenation of all the blog's posts and the second was based on smaller documents, each representing a single post, while aggregation was made at ranking time. Evaluation indicated that both models performed equally well and that hybridization of both further improved the results.

Multimedia Multimedia recommendation is challenging due to the lower amounts of textual data and the extremely large size of the content. One of the most popular social media websites, YouTube, includes an advanced recommender system that drives a large portion of the user traffic and helps direct users to more relevant videos. Davidson et al. [23] stated that the goals of the YouTube recommendations are to be recent and fresh, diverse, and relevant to the user's recent actions. They also stated that users should understand why a video was recommended to them, thus incorporating explanations in the YouTube RS. As described in their paper, YouTube recommendations are based on the user's personal activity on the site and are expanded by a variant of collaborative filtering (CF) over the co-visitation graph. Ranking is done based on a variety of signals for relevance and diversity. A later study examined the impact of recommendation on excessive use of online video streaming services and found it to be high [56].

Community Question and Answering Social or community question-and-answering (SQA or CQA) websites, such as StackOverflow, Quora, and Yahoo Answers, allow users to ask various types of questions and receive (and vote for) answers from the crowd. Questions may a wide variety of domains and seeks for information, conversation, or both [7, 45, 51]. As such, they also serve as a fertile ground for different types of recommender systems for both question askers and answerers. The challenge here is twofold: on the one hand, recommend to askers similar previously-asked questions to avoid redundant burden on answerers and spread of similar information in many question pages; on the other hand, recommend answerers with questions they may want to answer and increase overall answer engagement on the website. As one example, Szpektor et al. [99] experimented with recommendation of questions to potential answerers on the Yahoo Answers website. They discovered that topic relevance was not a good enough basis for recommendation. Diversity and freshness also played a key role: on the one hand, a novel and somewhat different question was more likely to arouse answerer's attention and on the other hand it was extremely important for answerers to receive questions that are very fresh, typically only a few minutes old.

Jobs LinkedIn is one of the most successful SNSs and as the world's largest professional network it has many unique recommendation challenges, such as of companies and of professional groups. Another specifically interesting example is the recommendation of job opportunities. Such recommendation can have a tremendous influence on people's lives as it can ultimately lead to a career change. Recommendation needs to take into account many aspects, such as location alternatives, candidate's experience, and timing. Wang et al. [106] shed some light on the job recommendation task at LinkedIn and particularly focus on the timing of recommendation. Their statistical model considered the tenure between two successive decisions to estimate the likelihood of a user's decision to make a job transition at a given point. Evaluation used the real-world job application data and demonstrated the effectiveness of their model and the importance of considering the time factor as part of the recommendation process. Olsson et al. [78] expand the scope and define "professional social matching" as the matching of individuals or groups for vocational collaboration and co-creation of value. This covers organizational activities, including recruitment, headhunting, community building, team formation, and individually-driven activities like mentoring, seeking advisory relationships, and general networking.

News Social news aggregators such as Digg, Google Reader, Reddit, and Slashdot, allow users to post and rate news articles and surface the most interesting and trending stories. News recommendation is especially challenging due to the need for freshness. Old stories or stories to which the user has already been exposed will be considered bad recommendations, even when relevant to the user's tastes and preferences. The pace of news appearance is very high, while different users have different news consumption rates, which personalization techniques need to take into account. Digg used to be a popular social news aggregation service, allowing its users to submit links to news stories, vote, and discuss them. Aside from promoting the most popular stories to users (by votes), Lerman [68] described the personalized recommender system implemented for Digg that was based on friends and "diggers like me". Recommendations for another popular news website, Google Reader, were described by Liu et al. [71]. They combined CF techniques with "individual filtering" techniques. Evaluation, based on a live trial, indicated that the hybrid approach performed best and improved 38% over a popularity-based baseline. Pure CF was only able to improve 31% on top of the baseline. An increase in return rate was observed due to the hybrid recommendations, however, interestingly, there was no effect on the overall number of stories read on the homepage.

Microblogs Microblogging, most famously brought into attention by Twitter, allows user to broadcast short messages. Those messages are typically propagated across a network of followers and "followees", built by the user's ability to follow any another user. On twitter, each message is limited to 140 characters and is called a 'tweet'. The high pace of messages (over half a billion tweets per day), their real-time nature, their concise content, and the lack of metadata and structure, make the challenge of filtering and personalizing the Twitter firehose of unique nature. In one of the earlier studies, Chen et al. [17] explored content recommendation through URLs shared in tweets. They compared 12 algorithms that differed in the following aspects: (1) candidate selection was either based on popular tweets or on tweets from followees and followees-of-followees (FoF); (2) topic relevance was based on cosine similarity between the user and the URL. The user's representation was based on self-tweets or on followees' tweets; (3) social voting was based on the number of user's followees who also follow the author and on author's frequency of tweeting. Results, based on a field study with 44 subjects, indicated that social voting worked better than topic relevance; FoF candidate selection outperformed popularity; and using self tweets for user modeling performed better than using the followees' tweets. The introduction of the 'retweet' feature, which allows user to share another user's tweet with their own audience of followers, provided researchers with direct feedback about the level of interest in an individual tweet. Many studies followed that attempted to use this information to predict "good" tweets. For example, Chen et al. [18] suggested a model for personalized tweet recommendation using "collaborative ranking". The model was based on both explicit and latent features and considered a wide variety of topic-level, social relations, and global factors. Evaluation was based on re-tweet prediction and showed the superiority of the collaborative ranking method over various baselines, such as Latent Dirichlet Allocation and Support Vector Machine. It also indicated that all the three factors are important to consider. In more recent work, Alawad et al. [3] studied the recommendation of "novel" tweets that were not posted or shared by anybody in the user's network. To this end, they created the user's egocentric network up to depth two and applied the transitivity property of the friend-of-a-friend relationship to yield recommendations. Piao and Breslin [84] proposed a learning-to-rank approach for recommending tweets that a user might re-tweet based on a deep neural network with Long Short-Term Memory (LSTM). The rank score was based on both the similarity between the embeddings of a user and the tweet and the similarity between the embeddings of the user and the tweet's author. The 2020 RecSys challenge focused on microblogging, with Twitter releasing around 160 Million public tweets obtained by sub-sampling over a period of two weeks. The specific task was to determine the probability that a user is going to engage with the content of another user via reply, retweet, or 'like' [95].

2.2 Group Recommender Systems: Beyond Preference Aggregation

Groups and communities play a central role in social media and often times form the entry gate for participation [89]. This makes group recommendation techniques highly relevant for the SRS domain. Due to this relevance, we briefly review the broad area of group recommendation in this sub-section; in the following section, as part of the enterprise case study, we describe in more detail an example of SRS aimed for communities. Group recommendation targets a group of individuals rather than a single one (Chapter "Group Recommender Systems: Beyond Preference Aggregation"). Example scenarios for group recommendation include friends planning together their "perfect" vacation; a family selecting a movie or a television show to watch together; a group of colleagues choosing a restaurant for an evening outside (or looking for a recipe for a joint meal); or the classic (and less relevant in the era of personal music players) gym problem [73]: selection of a playlist based on the current group of trainees in a fitness center.

Group recommendation poses new challenges compared to individual recommendations. Two of the prominent challenges are the specification of preferences by members and the recommendation generation. Jameson et al. [61] suggested a collaborative interface for members to specify their preferences in a group recommender system for travel, which allowed collaborative editing of the members' preferences. Such an interface holds various benefits: it allows members to persuade others to specify a similar preference to their own, perhaps by giving them information they had previously lacked; it enables to explain and justify a member's preference (e.g., "I can't go hiking due to an injury"); it allows taking into account attitudes and anticipated behavior of other members; and it encourages assimilation to facilitate the reaching of agreement.

The most studied challenge of group recommendation is the generation of recommendations themselves. The two main techniques are profile aggregation and recommendation aggregation. Profile aggregation produces a single profile representative of the group by aggregating the preferences of the different group members. Recommendation aggregation generates a recommendation list for each of the group members and aggregates the list into one single list for the group, typically by using rank aggregation techniques. Berkovsky et al. [10] experimented with these two approaches for recipe recommendation to groups and found that the profile aggregation method was superior over the recommendation aggregation method.

There are various approaches for aggregating member preferences into a single community profile, each with its own pros and cons. Among the prominent approaches are: (1) least misery, which seeks to maximize the minimum ranking of any group member. Obviously, this approach can lead to a recommendation that does not maximize the average rating or the maximum benefit; (2) fairness, which aims at the most equal rating balance across group members. This can lead to a recommendation that gets a low rating by all members of the group; (3) and fusion, which aggregates individual rankings (e.g., by Borda count). Baltrunas et al. [6] compared several techniques for group recommendation using the MovieLens dataset. They examined both profile aggregation and rank aggregation techniques and found the optimal one given a set of parameters, such as the group's size and the similarity among group members. Review of additional studies on group recommendation can be found in a recently published surveys [22, 28].

2.3 Case Study: Social Media Recommendation in the Enterprise

In this section, we review a body of research that explored recommendation of mixed social media items within the enterprise, and included three main studies. The first study [53] focused on recommendation based on social relationships . As previously mentioned, social media enables the exposure of different types of social relationships in a way that has never been possible before. The study explored a rich set of indicators for social relationships based on social media data and compared

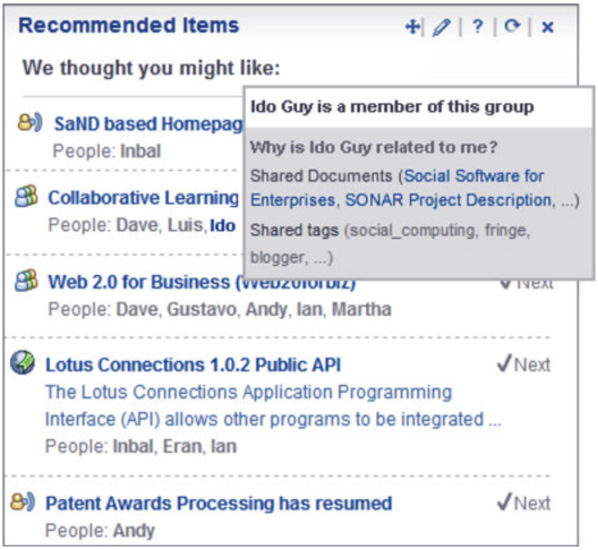

Fig. 1 Widget for social media item recommendation based on related people

two types of networks as basis for recommendation: familiarity and similarity. The familiarity network was built based on explicit and implicit signals from enterprise social media, such as articulated connection within an enterprise SNS, tagging one another, or co-authorship of the same wiki page. The similarity network was based on common activity in enterprise social media, such as membership in the same communities, usage of the same tags, or commenting on the same blog posts. An "overall" network was also examined, combining the two types of relationships.

The recommendation widget, depicted in Fig. 1, presents the recommendations with explanations, which displays the people who served as the "implicit recommenders" and how they were related to both the user and the recommended item. One of the key research questions of the study was whether explanations influence the instant interest in the recommended items. This was examined by comparing recommendations with and without explanations.

The evaluation was primarily based on a user survey with 290 participants. Figure 2 shows the portion of items rated "interesting" for each of the three network types: familiarity, similarity, and overall. Recommendations from familiar people were found significantly more accurate than recommendations from similar people. The overall network did not improve accuracy on top of the familiarity network. That said, recommendations from similar people were found more diverse and less expected, indicating that the similarity network contributes on other dimensions than accuracy to the recommendation quality [75].

Figure 3 displays the effect of explanations. While explanations have been previously shown to have positive effect on recommendation in the long term, by providing transparency and building trust with the user [57], it was found

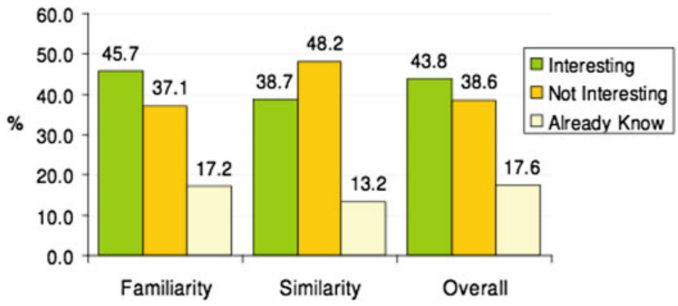

Fig. 2 Rating results across the three network types

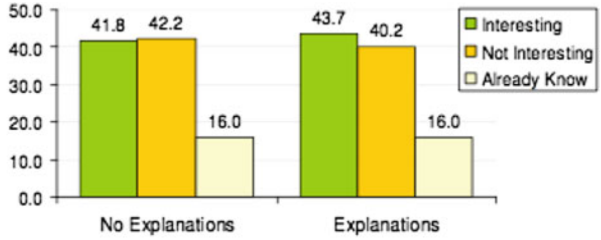

Fig. 3 Rating results with and without explanations

that recommendations with explanations in this case also increase their instant effectiveness: when the people who serve as implicit recommenders were shown, interest rate in the recommendations grew. This was particularly true for familiar people, following the intuition that seeing a familiar person who is related to a recommended item may increase the likelihood of the user's interest in that item (e.g., "if John has bookmarked the page, there must have been something interesting in it").

After establishing understanding of people-based recommendation, a second study explored the use of tags for the recommendation task and compared tag-based with people-based recommendation [54]. The people-based recommendations were calculated based on a combined network of familiarity and similarity, with a triple-boost given to the familiarity network based on the results of the previous study.

Based on the results of a preliminary study, the tags used for recommendation included those used by the user and those applied to the user by others via an enterprise people tagging application [87]. Experimentation was made with a pure people-based recommender (PBR), a pure tag-based recommender (TBR), two hybrid people-tag recommenders (or-PTBR and and-PTBR), and a popularity baseline (POPBR).

The main comparison results are shown in Fig. 4. In general, all personalization techniques outperformed the popularity-based recommender. In terms of accuracy (interest rate), tag-based recommenders significantly outperformed people-based recommenders. Yet, people-based recommenders showed other benefits, such as

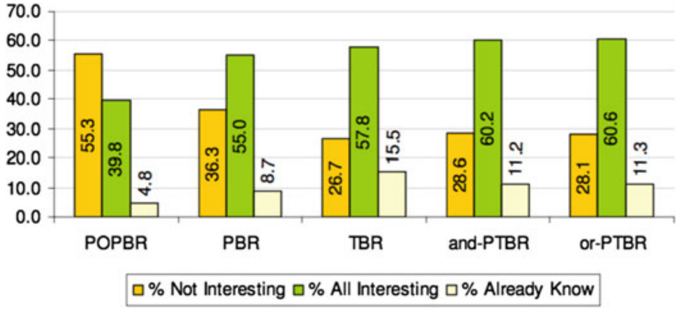

Fig. 4 Rating results for five different recommenders

increased diversity across item types (tags substantially favored bookmarks), less expected results reflected in lower rates of already-known items, and more effective explanations. Specifically regarding explanations, the effect found for people-based explanations in increasing interest rates was not found for tag-based recommenders. Apparently, seeing the related tags to a recommended item does not have the effect (or extra value) that viewing the related people has. Hybrid recommendations, combining people-based and tag-based approaches, were shown to take the good of both worlds and also achieved the best accuracy with a ratio of around 70% interesting items for the top 16 recommendations.

The third study in the series explored recommendation for online communities rather than for individuals [89]. As mentioned before, online communities have become central to social media experience and much of the social media content is created in the context of a community. In that work, recommendations were generated using group recommendation techniques, but were targeted to the community owners (moderators) only, so that they can share the content with the rest of the members as appropriate. Recommendations were generated using two main techniques. The first considered the members of the communities or a subset of them, and applied profile aggregation using the fusion approach (with advanced scoring) to generate a community profile that included both topics and people. These topics and people in turn served as the basis for recommendation: their most related content items were recommended. In particular, three subsets of the members were examined: all members, all owners, and active members. The second technique was content-based (CB): it considered the title, description, and tags of the community to generate recommendations. Hybrid approaches were also considered, by combining the topics and people from the member-based recommenders with the topics extracted by the content-based recommender into one community profile.

Evaluation was conducted using a large user survey of enterprise community owners and results are summarized in Fig. 5. Hybrid recommenders were generally found to perform better than the pure recommenders. For large communities (100 members or more), it was found that the hybrid profile that considered both active members and community's content performed significantly better than all

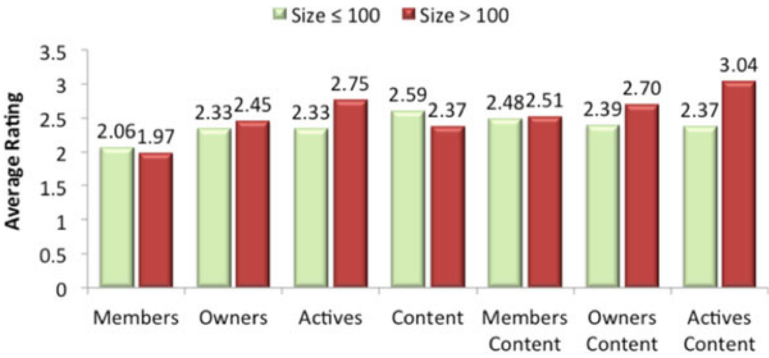

Fig. 5 Average rating for small vs. large communities across seven community profiles

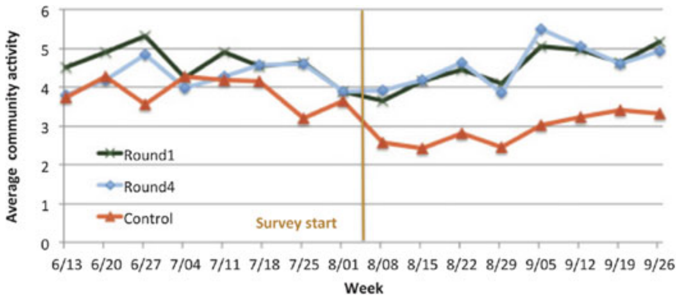

Fig. 6 Community activity before and after a survey where owners could share social media content with their community

other profiles. The pure active member-based profile was second best for large communities. For small communities (less than 100 members), the pure content profile was the best, followed by the hybrid profile considering all members and the content. These results indicate that for small communities, the content is a strong basis for recommendation and all members are a good representative group for profile aggregation. But for large communities, the content is less effective on its own and the group of all members becomes too disparate, while the group of only active members serves as the best basis for profile aggregation.

In a follow-up study, the impact of such recommendations was inspected over four rounds of recommendation to enterprise community owners [47]. Owners could share recommended content, such as blogs, bookmarks, and forum posts with their community members. As can be seen in Fig. 6, such recommendations, in both the first and fourth round, showed significant effect on overall community activity over a period of the following 8 weeks, compared to a control group whose owners received no recommendations.

2.4 Summary

We reviewed different domains for recommendation of social media content and a case study for recommending mixed social media items in the enterprise. We also discussed the importance and relevance of group recommendation techniques when recommending social media content. Below are a few important points we wanted to re-iterate before moving to the next section:

- Articulated social networks play an important role in CF for social media content and enhance traditional CF in various manners.
- Tag-based recommendations are highly effective for producing accurate recommendations and typically outperform regular user-based CF.
- As in Traditional RS, hybrid approaches (e.g., tags + networks, short + long term interests, collaborative + individual filtering) usually enhance recommendation effectiveness.
- A large user-base is desirable and can lead to a strong evaluation on live systems (e.g., A/B testing).
- Accuracy alone is not enough: serendipity, diversity, freshness, and other qualities also play a key role in the success of recommendations.

3 People Recommendation

Social recommender systems span beyond content recommendation. As mentioned in the introduction, social overload originates from both information and interaction overload. Since people are the key element that makes the web "social", recommendation of people is a central pillar within the social recommender system domain. Terveen and Mcdonald [102] coined the term "social matching" for recommender systems that recommend people to people. In their work, they explained why a people recommender is a unique RS, which is different than recommendation of other artifacts, and thus deserves its own special attention (see also Chapter "People-to-People Reciprocal Recommenders"). Among other aspects, trust, reputation, privacy, and personal attraction have greater importance when it comes to people recommendation.

Social media sites and in particular SNSs define different types of explicit (or "articulated") relationships among their users. The main dimensions of the relationship types are:

- Symmetric vs. asymmetric. In some sites, such as Facebook and LinkedIn, a relationship between two users is reciprocated. In such a case, one user typically sends an invitation to connect to another user, who needs to accept the invitation. Once the other user accepts, the two are reciprocally connected on the site [81]. On the other hand, asymmetric relationships, such as on Twitter or Pinterest, allow one user to "subscribe to" or "follow" another user. The other user does not necessarily need to follow the first user back and thus many asymmetric relationships are formed.

- Confirmed vs. non-confirmed. Some of the sites require the other side's agreement for connecting or following, while others do not. Typically, symmetric networks require such confirmation and as long as it has not been received, no connection exists. Asymmetric networks do not usually require a confirmation and any user can choose to follow any other user, however there are exceptions to these norms.
- Ad-hoc vs. permanent. Some of the sites encourage connection for an ad-hoc purpose, such as for people to meet at an event or partner for a joint task, while others encourage a long-term relationship that is meant to last over months and years.
- The site's domain. The domain of the SNS has an important influence on the formed network. For example, Facebook is typically used for maintaining social relationships with friends and acquaintances, while LinkedIn is a professional network meant for maintaining business relationships with colleagues and partners. The goals and characteristics of a connection in each of these sites are therefore different, as they would be in SNSs for other domains, such as travel, art, cooking, question and answering, etc.

The different characteristics of people relationships in the different sites require different recommendation techniques. For example, a recommender for people to connect with on Facebook may seek to recommend familiar people, while a recommender for people to follow on Twitter may recommend people the user is interested in, even if they are not familiar. Recommending "celebrities" or popular people is probably a better strategy for a follower-followee network than for a friendship network. A good summary can be found in [41].

In the remainder of this section, we review three key types of people recommendation: recommending people to connect with, recommending people to follow, and recommending strangers to get to know. We describe the unique challenges and characteristics of each of these recommendation types and demonstrate how existing approaches handle them. Before summarizing the key aspects, we briefly discuss two closely related research areas to people recommendation: link prediction and expertise location.

3.1 Recommending People to Connect With

The first study that focused on people recommendation in an SNS introduced the "do you know?" (DYK) widget [49]. The widget recommended people to connect to within an enterprise SNS. The action the widget was targeting was clicking a 'connect' button that would trigger an invitation to connect within the SNS, which the other side would need to confirm for the connection to become public. Recommendations were made based on a variety of familiarity signals: org-

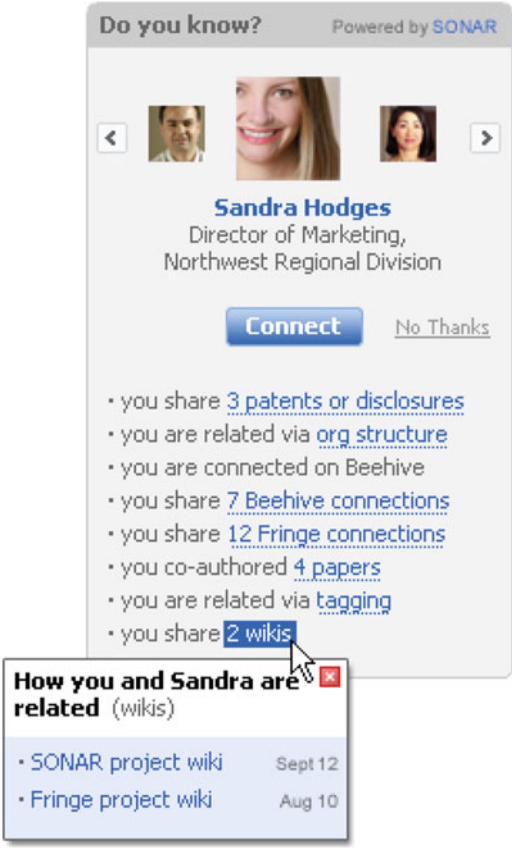

Fig. 7 The "Do You know?" (DYK) widget

chart relationships (peers, manager-employee, etc.), paper and patent co-authorship, project co-membership, blog commenting, person tagging, mutual connections, connection on another SNS, wiki co-editing, and file sharing. Figure 7 illustrates the widget, which included detailed explanations for each recommendation. The explanations indicated the counts per each of the signals mentioned above and further hovering over an evidence line allowed seeing the specific details (e.g., the wiki pages co-edited) and getting to the actual page of the evidence pieces.

The evaluation of the widget was based on a field study of its use within the Fringe enterprise SNS. Fringe had the "friending" feature before, but did not have a people recommender. The inspected effect on the site was dramatic. Both the number of invitations sent and the number of users who send invitations significantly increased, as can be seen in Figs. 8 and 9. One of the users of the site explained: *"I must say I am a lazy social networker, but Fringe was the first application motivating me to go ahead and send out some invitations to others to connect."* Explanations increased user trust in the system and made them feel more comfortable sending invitations, as one user described: *"If I see more direct connections I'm more likely*

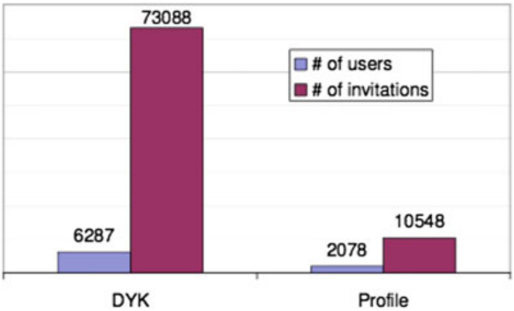

Fig. 8 DYK vs. Profile usage throughout the inspected period

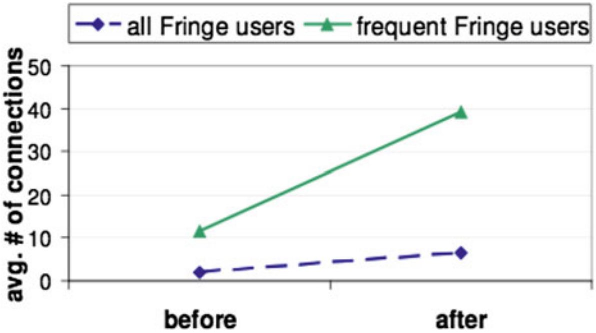

Fig. 9 Average number of invitations per user before and after the inspected period

to add them [. . .] I know they are not recommended by accident.” Overall, there was a substantial increase in the number of connections per user on Fringe. However, a sharp decay of the widget usage was found over time, as excitement of the feature dropped and potential connections were exhausted.

In a follow-up study [16], conducted within a different enterprise SNS, nicknamed Beehive, the aggregation algorithm used by the DYK widget (termed ‘SONAR’) was compared with three other algorithms for people recommendation: (1) Content Matching (CM)—based on cosine similarity of the content created by both users: profile entries, status messages, photos’ text, shared lists, job title, location, description, and tags. Word vectors were created by a simple TF-IDF procedure. Latent semantic analysis (LSA) was not shown to produce better results and was not applied since it does not yield intuitive explanations; (2) Content plus Link (CplusL)—combined CM with social links. A social link was defined as a sequence of 3 or 4 users, where for each pair of users in the sequence u1 and u2, either u1 connects to u2, u2 connects to u1, or u1 commented on u2’s content; (3) Friend of Friends (FoF)—based on the number of mutual friends, as done in many of the popular SNSs. The FoF algorithm was able to produce recommendations for only 57.2% of the users (compared to 87.7% for SONAR). Figure 10 shows the recommendation widget.

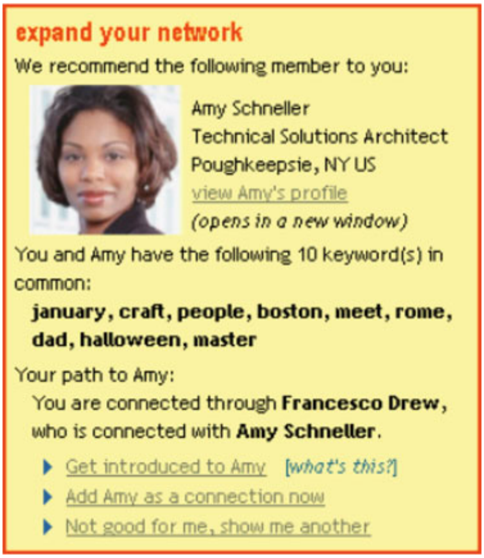

Fig. 10 People recommender widget showing a person recommended using the CplusL algorithm

Evaluation was based on a user survey and a controlled field study. Figure 11 shows the main survey results. CM and CplusL produced mostly unknown people, while SONAR and FoF produced mostly known individuals. As could be expected, a higher portion of the recommended people who were familiar to the user were rated as good recommendations and resulted in a "connect" action. Yet, the unknown recommended individuals may help discover new potential friends. The overall superiority of algorithms that involve social links over content was clear: only 30.5% of the CM recommendations resulted in a connect action, compared to 40% for CplusL, 47.7% for FoF, and 59.7% for SONAR.

A later study examined the recommendation impact on the network structure [21]. Since recommendations play such a key role in building the network during its early stages, they also substantially influence the structure of the generated network, its characteristics, and measurements. For example, Fig. 12 shows the average degree of recommended connections for each of the four algorithms. FoF is the most biased towards high-degree connections, while CM does not have such bias: it often recommends users with few connections or even none at all. The high-degrees of FoF recommendations lead to a network with fewer nodes and higher average degree compared to the network created by CM recommendations. Another aspect of the effect of recommendations on the network is betweenness centrality, which measures the importance of nodes in the graph [12]: CM and SONAR generate the highest delta in betweenness compared to CplusL and FoF. Regarding demographic characteristics, CM is most biased towards the same country, but least biased towards the same organizational unit, while SONAR substantially increases cross-country and intra-unit connections. The network effects of people recommendations

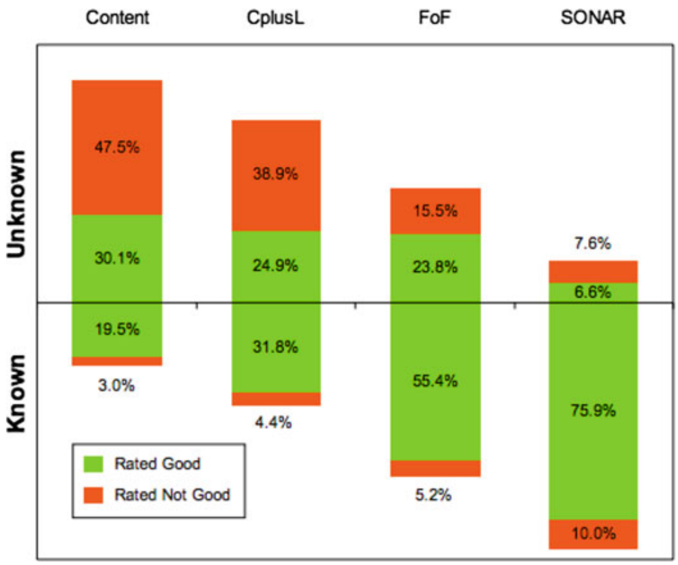

Fig. 11 Survey results for the four algorithms

Fig. 12 Degree of recommended connections across the four algorithms

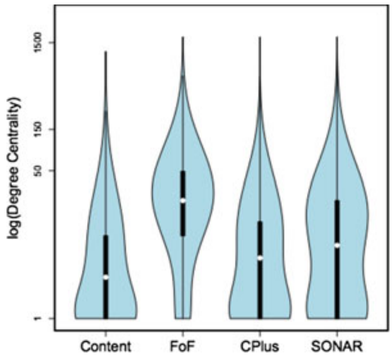

are an important global aspect of a people recommender and need to be considered when designing a new people recommender system.

Another related study by Freyne et al. focused on recommendation as a means to increase new users' engagement within an enterprise SNS [32]. That study used aggregated data external to the SNS in question to recommend both people and content to new users. Even brand new employees could still get recommendations based on their initial data, such as their org-chart information (indicating their peers), location, or organizational unit. The results indicated that combined recommendations have a significant effect in increasing users' visits to the site as well as their viewing activity and actual contributions to the site (the latter is depicted in Fig. 13). Interestingly, people recommendations were most effective when focusing on recommending the most active users, even if they had less familiarity signals

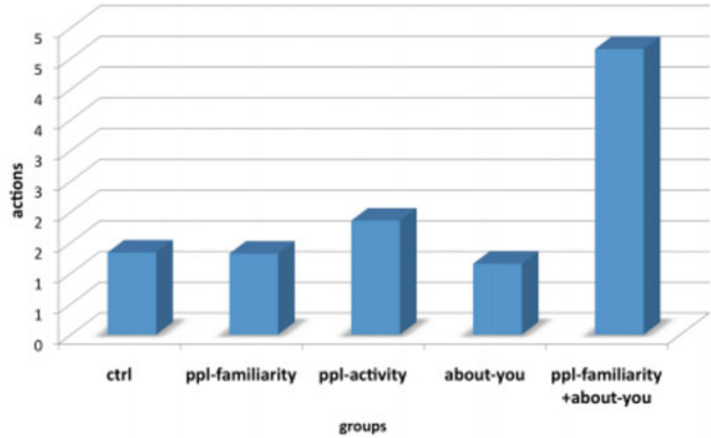

Fig. 13 Actions over four months

with the user. Yet, as discussed, such recommendations can have a long-term effect on the network structure and lead to a less balanced degree distribution.

Friend recommendation has also become popular on mobile devices, where location often plays a role and makes the recommendations more transient or ad-hoc. Quercia et al. [86] discussed "friendsensing", sensing friends based on Bluetooth information on mobile devices. Friends were recommended based on co-location, while two basic approaches were attempted, taking into account the duration of co-location and its frequency, respectively. A weighted graph was built accordingly and recommendations were generated using that graph based on link analysis (shortest path, page rank, k-markov chain, and HITs). Simulation-based evaluation indicated both basic approaches perform similarly well and way beyond a random baseline.

Gurini et al. [40] proposed a matrix factorization model with temporal dynamics to provide people recommendations. Each dimension in a three-dimensional matrix factorization model represented an "attitude": sentiment, volume, and objectivity. Recommender's accuracy and diversity was shown to increase with attitudes and temporal features.

3.2 *Recommending Strangers*

The focus of the work discussed thus far has been on recommending familiar people one can connect to. As already implied, there could also be value in recommending people the user does not know. StrangerRS [43] attempted to recommend people who are unknown yet interesting within the organization. Such recommendations can be useful in many potential manners, such as, for getting help or advice, reach

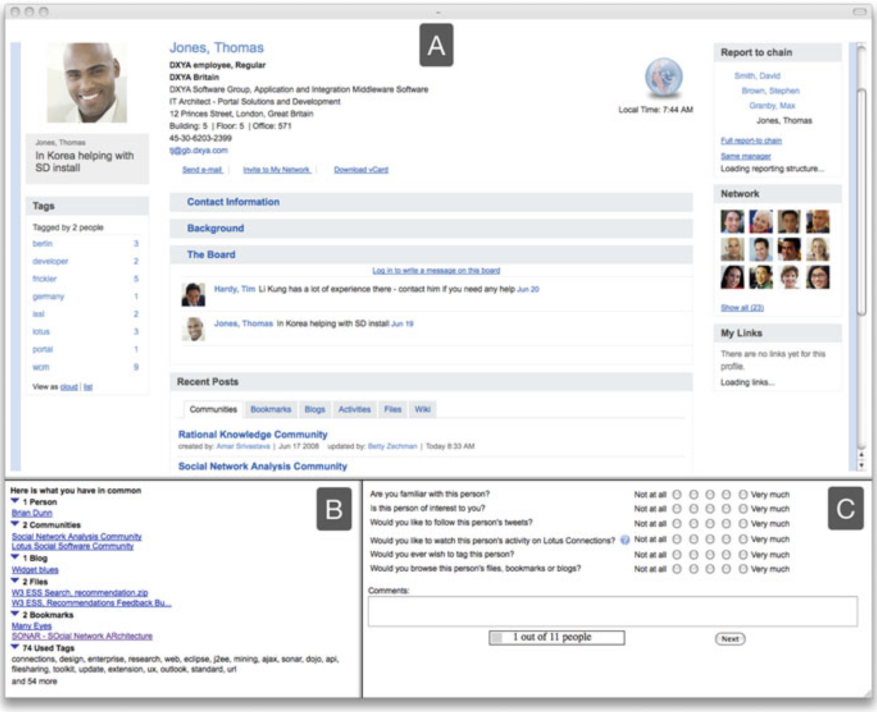

Fig. 14 User interface of the stranger recommender system

new opportunities, discover new routes for career development, learn about new assets that can be leveraged, connect with subject-matter experts and influencers, cultivate one's organizational social capital, and grow own reputation and influence within the organization. As mentioned before, recommendation of people to connect to within an SNS is mostly effective for the network-building phase. Afterwards one's recommendations become staler, as the network becomes more stable and connection to others becomes less frequent. This is where stranger recommendation can become more relevant and complement the recommendation of familiar individuals, by suggesting interesting people the user does not know, but may want to start getting acquainted with.

Figure 14 shows the user interface of StrangerRS. Since it aimed at recommending strangers, more information about each person was presented, in the form of their full profile page (part A). Evidence for why this person may be interesting was also presented (part B). It included similarity points with that individual, such as common tags, common communities, common files, and others. The action suggested by the recommender was not a connection within the SNS, since it is likely to be too soon to connect to a stranger, but rather it was suggested to view the person's profile, read their blog, or follow them (part C).

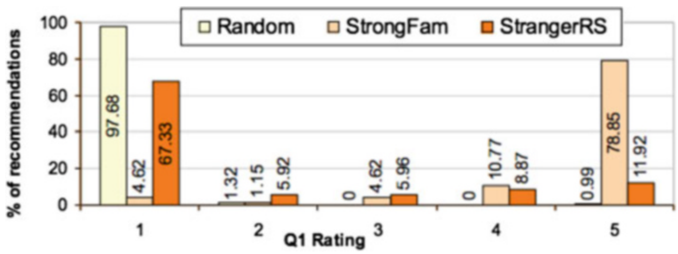

Fig. 15 Rating of "strangerness" for StrangerRS and two baselines: random and strong familiarity

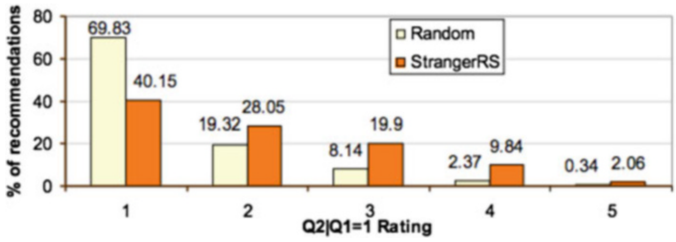

Fig. 16 Rating of interest in strangers for RandomRS vs. the random baseline

A successful recommendation by StrangerRS was considered a recommendation of a stranger who might be interesting to the user. These two, almost contradicting, goals were not easy to satisfy and led to a much lower accuracy level than usual familiar people recommendation. Yet, supposedly, the value of a successful recommendation in this case is much higher, since this is no longer just about facilitating connection to a known person, but rather about exposing the user to a new interesting person s/he was not even aware of. The method used for producing the recommendations was based on network composition: the extracted familiarity network was subtracted from the extracted similarity network to produce the recommendations. Jaccard index was the main measure used for similarity between two individuals. Results, depicted in Figs. 15 and 16, indicated that two thirds of the recommended individuals were indeed strangers, yet strangers who were significantly more interesting than a random stranger. Out of 9 recommendations presented to each user, 67% included at least one stranger rated 3 or above in terms of the user's interest, on a 5-point Likert scale.

Stranger recommendation is also a common feature of online dating websites. Pizzato et al. [85] introduced RECON, a reciprocal recommender for online dating. Similarly to the original social matching framework, they specified a few special characteristics of reciprocal recommendations, where people are both the subject and object of recommendations. These included the fact that success is dependent on both sides; the need for both sides to provide their profiles so that matching can occur; and the typical requirement that one individual will not be recommended to too many others. Their evaluation, conducted on a major Australian dating site, was based on four weeks of training and two weeks of testing, where success was determined based on previous user interaction. Generally they found that accounting for reciprocity features improves recommendation accuracy and helps address the cold-start problem. Zheng et al. [112] studied speed dating data, where the user's expectations are well defined, and showed that addressing fairness by trading off utility and performance yields a better recommender system.

3.3 Recommending People to Follow

Two studies were the earliest to explore recommendation of people to follow. Hannon et al. [55] used a CB-CF hybridization to recommend "followees" on Twitter. They examined several ways to generate user profiles, based on the user's own tweets, the user's followers, the user's followees, the user's followers' tweets, and the user's followees' tweets. The open source search engine Lucene was used to index users by their profile, after applying TF-IDF to boost distinctive terms or users within the profile. They applied an offline evaluation using a dataset with 20,000 Twitter users. 19,000 were used as a training set and the remaining 1000 were the test users. The different methods were compared based on their ability to predict the user's followees. A slight advantage was observed to profiles that were based on followers and followers' tweets. Hybrid profiles further improved the precision. A small-scale live trial was also conducted where users indicated whom they were likely to follow. On average, hybrid approached reached about 7 out of 30 accurate recommendations.

A second study was performed by Brzozowski and Romero [13], who experimented with the WaterCooler enterprise SNS. During a 24-day live trial period, they observed patterns of 110 users who followed 774 new individuals. The strongest pattern found was of the form $A \leftarrow X \rightarrow B$, meaning that sharing an audience (follower) with another person is a strong reason to follow that person. Most-replied was found as a strong global signal. Similarity and most-read were found as weaker signals for followee recommendation.

Gupta et al. [39] revealed some details about the followee recommender systems in use by Twitter. From an architectural perspective, they noted the decision to process the entire Twitter follower-followee graph in memory using a single server, which contributed to the performance of the feature. They developed an open-source in-memory graph processing engine to traverse the Twitter graph and generate recommendations. The algorithm used was a combination of a random walk and SALSA [67], comparing two approaches: the first gives each user the same influence regardless of the number of users they follow or are followed by and the second gives equal influence to each follower-followee edge. Sanz-Cruzado and Castells [92] study diversity in people recommender systems, linking to prior diversity notions in recommender systems. They focus on the notion of weak ties to enhance structural diversity and demonstrate, using Twitter data, how state-of-the-art recommendation methods compare in such diversity and its tradeoff with accuracy. They show that diverse recommendations result in a corresponding diversity enhancement in the flow of information through the network, with potential alleviation of filter bubbles.

3.4 Related Research Areas

Link prediction in social networks is a fertile research area that is closely related to people recommendation and has often been offered to enhance it. The seminal work by Liben-Nowell and Kleinberg [70] formalized it as a task to predict new interactions within a social network based on the existing set of interactions. Experimentation with paper co-authorship networks showed, using an unsupervised learning approach, that the network topology can be effectively used to predict future collaboration. Moving to the social media domain, Leskovec et al. [69] developed models to determine the sign of links (positive or negative) in SNSs where interactions can be positive or negative (Epinions, Slashdot, Wikipedia). Fire et al. [29] experimented with five social media sites, including Facebook, YouTube, and Flickr, and proposed a set of graph-topology features for identifying missing links. This technique was shown to outperform common-friends and Jaccard's coefficient measures, implying it can be useful for recommending new connections. Scellato et al. [93] focused on location-based social networks and suggested a supervised learning framework to predict new links among users and places. In another study of mobile networks, Wang et al. [105] showed that combining network-based features with human mobility features (e.g., user movement across locations) can significantly improve link prediction performance using supervised learning.

It is also worth mentioning the research area of expertise location [64, 74] in the context of people recommendation. Expertise location deals with the problem of finding an expert in a given domain or technical area. It thus falls within the broad search domain, since it is triggered by a user query. Similarly to the difference between content and people recommendation, in expertise location the results are people rather than documents as in content search. For similar reasons to those already discussed in this section, the case of searching for people bears some unique characteristics compared to other content search scenarios and therefore forms its own area of research. Social media serves as a particularly good source for expertise mining and analysis[42] and for providing explanation in the form of expertise evidence [111]. Despite its pertinence to the search field, expertise location is sometimes mixed with people recommendation and in many cases is termed "expert recommendation". It should be noted that a people recommender should be considered as such only when it does not involve a user query and is initiated by the system rather than by the user.

3.5 Summary

We have seen that people recommendation is a complex field of study. The fact that it deals with recommending people to themselves bring many interesting aspects to the table. For example, explanations may serve in this case to make users feel more comfortable accepting a recommendation and sending an invitation to connect or start following (in most cases, knowing that the user who has been followed will get a notification about it, even if their approval is not required). We reviewed three types of people recommendations: recommendation of familiar people, for example to connect with on an SNS; recommendation of interesting people, for example to follow in a social media site; and recommendation of strangers, for dating or for getting to know within a community or organization. A social website may transfer between these types of recommendations according to the user's phase within the site. For example, it may be desirable to recommend familiar and interesting people for users in their early stages, so they can build their network of friends or followees. In a later stage, when users start to exhaust their connections, stranger recommendations can help users get to know new individuals and increase their social capital.

4 Discussion

In this section, we summarize key SRS-related topics that were brought up throughout the previous two sections on people and content recommendation and suggest directions for future work.

Explanations The public nature of social media data enables to provide more transparency into recommendations by showing how they were formed. In some of the enterprise examples we reviewed for both content and people recommendations, explanations were found to have a key role in increasing the instant acceptance rate of recommendations [49, 53]. Studies in the Web domain have shown a similar effect [66, 103]. Beyond that, explanations in RS have been shown to have longer-term effects of building trust relationships with the user [57] (Chapter "Beyond Explaining Single Item Recommendations").

There are also a few challenges with regards to explanations. First, as we have seen, explanations do not always increase accuracy. For example, in our mixed content recommendation study, tag-based explanations did not increase recommendations' ratings. Second, not every recommendation method can provide intuitive explanations; there is usually a trade-off between the method's complexity and the clarity of explanations it can provide. For instance, recommendations that are based on clustering techniques are usually harder to explain. Third, explanations pose challenges in terms of privacy. For example, the YouTube explanations [23] explicitly show videos previously watched by the user, which directly expose information that might be sensitive if watched by another person. Fourth, explanations require extra real-estate on the user interface, which might be particularly challenging on mobile devices; therefore their cost-to-value ratio should be carefully considered when designing the recommender system.

Privacy As mentioned several times throughout this chapter, one of the key benefits of social media data is that large portions of it are public and thus can be used for analysis without infringing user privacy, as is the case, for example, with email or file system data. It should be noted, however, that in some countries, public social media information is still considered personal information (PI), when linked to an identity of a real person. This means that analysis and inference from such data may still require explicit user consent. Indeed, aggregation of public data, even if it was previously accessible, may reveal sensitive information the user did not intend to expose. In addition, as just mentioned, explanations aimed for a specific user might reveal very sensitive data, such as browsing or viewing history, when exposed to another person who may watch the screen alongside. Finally, there is much social media data that is still access-restricted. Recommender systems should pay special attention not to infringe the privacy model of the data, to avoid the exposure of sensitive information [25].

Tags The work we reviewed indicated that tags, a mechanism introduced by social media to annotate content, such as web pages, photos, or people, can be particularly effective as a basis for recommendation. Tags' ability to concisely summarize user perspective over large content pieces make them a highly valuable resource for producing recommendations [94]. Aside from recommendations, tags have been shown to be useful for other purposes, such as enhancing search or generating "tag clouds" that summarize the common topics of a group of items to the user [63]. Unfortunately, despite their value, tag usage is on the decrease in recent years, with sites such as Delicious becoming less popular and other sites giving less prominence to tags. Tag recommendation techniques [62, 96], which are another type of SRS not discussed in this chapter, should be used to promote tag usage and close the loop: tag recommendations help generate more tags, while these tags, in turn, used to produce other recommendations.

Social Relationships One of the most important contributions of social media to recommender systems is the introduction of the explicit (articulated) network . Social network sites , such as Facebook, LinkedIn, and Twitter, allow people to explicitly articulate their connections. As mentioned, there are two main types of connections, one expresses familiarity and the other expresses interest. Both of these articulated networks are very useful for content recommendation, and were shown to enhance traditional CF techniques. They also have other benefits: (1) sparing the need for explicit feedback in the form of ratings to determine the network of similarity, (2) help coping with the new-user cold start problem, in case the network can be used across social media websites, and (3) helping users judge the recommendations, since they originate from people they know or are interested in (also making explanations more effective). On the other hand, as we have seen, recommendations of people to connect with or to follow are essential for enhancing the formation of such explicit relationships. This is a classic demonstration of the mutual relationship between recommender systems and social media discussed in the introduction: on the one hand social media introduces a new type of data that enhances RS; on the other hand RS are essential for generating this type of data.

Trust and Reputation The topic of trust has a tremendous importance in the RS domain. Obviously, the best recommendations come from a trusted person. But on the other hand, trust is very challenging to compute as it represents a very abstract and subjective quality between two individuals. Reputation represents a more general concept about a person's perception by others [59]. One way to define it is the aggregation of trust in this person across the entire set of users. Social media and the "wisdom of the crowd" enable to estimate trust and reputation in ways that have not been possible before. Online social relationships and content feedback forms (comments, 'likes', etc.) introduce more signals that can be used to calculate trust and reputation. That said, many of the studies still use rough estimations that are based on controversial assumptions, for example, that a friend on an SNS is a trustworthy individual. Evaluation of trust and reputation is also particularly challenging, as even in the real world people have hard time figuring out who they trust or who has a good reputation. Assuming a network of trust is given, there are growing amounts of research that explore how to use it to enhance CF. The early work of Golbeck [37] suggested to adapt the CF formula in a way that would boost similar users whom the user trusts. More advanced approaches incorporate trust in matrix factorization techniques [60]. Finally, Wu et al. [108] extend collaborative topic regression to model social trust ensemble that reflects true friends in SRS.

Evaluation As reviewed throughout this chapter, evaluation of SRS typically uses the common methods in the broader RS domain (Chapter "Evaluating Recommender Systems"). These include offline evaluation, user studies (especially common for SRS), and live field studies or A/B testing. Evaluation measures include RMSE, NDCG, precision, and other commonly used metrics from the RS domain. Looking forward, since social media is characterized by the "wisdom of the crowd", it will only be natural to see more crowdsourcing techniques used for evaluation of SRS. These have become common in many domains in the recent years, including information retrieval (e.g., [4, 14, 65]), however they are not as common yet in RS evaluation. Evaluation that goes beyond accuracy to include serendipity ("surprise"), diversity, novelty, coverage, and other factors is also due in the SRS area [34]. Finally, evaluation over time, which also examines the broader effect of the recommendation on the surrounding ecosystem of users, as demonstrated in [21], is a highly desirable direction. Rather than focusing mostly on recommendation effectiveness, their broader and longer-term influence on the environment should also be considered. As another example to such research, Said and Bellogin [91] started to explore the effect of recipe recommendation within the Allrecipes.com SNS on users' health. This kind of research requires new tools and creative thinking to be brought into the existing set of evaluation methods.

Recommending Content to Produce We extensively discussed content recommendation in Sect. 2. Our examples focused on content the user consumes: video, news, questions, social media items, etc. As explained in the introduction, one of the key characteristics of social media is that users are not just the consumers, but also the producers of content. There is a body of research that attempts to recommend users content they may want to produce. Question recommendation in CQA sites, which has already been mentioned in Sect. 2, has a role in encouraging users to produce content in the form of answers. Other works attempted to encourage users to create more profile entries [35], inspire users to write blogs [24], and prompt them to edit articles on Wikipedia [19]. Recommending content to generate is a particularly challenging task since the entry barrier is higher as many social media users are lurkers (only consume content). It is rooted in the area of persuasive technologies and theories such as self determination [90] and behavioral models [30]. Clearly, recommending content to produce has a central role in the symbiosis between recommender systems and social media.

5 Emerging Domains and Open Challenges

We conclude this chapter by pointing out potential emerging domains for SRS and a few open challenges on top of the topics discussed throughout this chapter and summarized in the previous section.

5.1 Emerging Domains

We enumerate four domains, which we think can serve as a fertile ground for SRS research in the years to come.

Mobile and Wearables Recommendations for mobile devices, such as PDAs, have been suggested since the beginning of the millennium. As smartphones and tablets with advanced technologies, such as high-resolution cameras, GPS, and touch screens started to prevail, recommendation technologies adapted themselves, for example, by taking into account the user's location. The combination of mobile and social (sometimes referred to as SoLoMo—social, mobile, and location) holds new opportunities for SRS, which will combine the advanced capabilities of mobile devices with social interaction across these devices. Looking further into the future, wearable devices, such as glasses and watches, are likely to have access to even more personal information that on the one hand will provide more data for SRS to work with, and on the other hand will require more advanced recommendation techniques, so these devices can work appropriately with minimum input from the user.

Smart TVs RS have been quite popular in the TV domain for many years. The Netflix prize advanced this domain even further [9]. However, as TVs continue to evolve into "smart TVs", they enable many more social elements, such as sharing and interaction between watchers, which make the new TVs a social medium on its own. This provides a highly interesting opportunity for SRS to make this new generation of televisions even smarter.

Automotive The automotive domain is also evolving in recent years. Self-driving cars is arguably the most exciting challenge on the table, but new car models allow more collaboration between cars and their drivers. Being such an advanced instrument, the car itself plays a special role and can sometimes be treated similarly to a person, given all the information gathered through its sensors. As more collaboration is expected to characterize the new generation of smart cars, SRS can play a key role in sparing extra work from drivers and providing cars with more necessary information. We start to see this in social navigation technologies, such as Waze, but this is likely only the beginning.

Healthcare The healthcare domain has always been slow to adopt "social", among other things due to the special privacy concerns it entails. On the other hand, it is not hard to imagine how much this domain can benefit from more sharing and collaboration, both among patients and among doctors. In recent years, we start to see a movement towards more openness to medical data sharing. For instance, Yang et al. [109] present a case study for healthcare based on a framework for a social recommender system using both network structure analysis and social context mining. First, microblog users are recommended using exponential graph models. Then, a recommendation list is created by analysing the micro-blog network structure. Finally, sentiment similarities are used to filter the recommendation list and find users who have similar attitudes to the same topic. The recommendation results of diabetes accounts over the Weibo network demonstrate high performance. As it seems that "social healthcare" is taking off, the SRS community should consider how recommendations should be used in this domain, with all the complexities involved and the critical implications of a successful versus wrong recommendation.

5.2 Open Challenges

We finally highlight three more challenges for researchers in the SRS area to consider.

Social Streams Social streams, such as Twitter or the Facebook newsfeed, syndicate user activity within a social media site or a set of sites. Millions of users who share and interact in social media create a firehose of data in real-time that poses new types of challenges in terms of filtering and personalization. There are different types of streams in terms of the data they contain (homogenous as in

Twitter or heterogeneous as in Facebook), the source of data (a single site or a group of sites), its access-control (public or friends-only), and subscription model (following or "friending"). As demonstrated in the Twitter-related work reviewed in this chapter, the stream's data is different than "traditional" social media content: it represents an activity rather than an artifact or an entity; it is more intensive as one entity (e.g., a wiki page) may have a large amount of activities (e.g., edits); it may be very noisy (e.g., multiple wiki edits might not be of interest); its freshness is key: items that are few days old might already be irrelevant; and it is sparse in content and metadata (e.g., Twitter messages are limited to 140 characters). Due to all these unique characteristics, recommending social stream items becomes a challenge on its own within the SRS domain, and as social information continues to grow, handling this task is becoming both more challenging and more important [2, 31, 44, 48, 52, 80]. On the other hand, the stream data can also be used to model users' interests. Its fresh and concise nature can help build a user model that is up-to-date, identify changes in users' tastes and preferences in real-time, and detect global trends that may influence the recommendation strategy [33, 83].

Beyond Accuracy and Evaluation over Time Many of the studies we reviewed focused on measuring the effectiveness of recommendation by their accuracy. As social recommendation proliferate, it is more important than ever to consider the bigger picture when evaluating the value of recommendation. Typical beyond-accuracy measures should be considered, including serendipity, diversity, novelty, and coverage [34, 75]. In addition, the effectiveness of recommendation should be compared against the case where no recommendation would have been provided [8]. Recommendations that can make the user discover and take action regarding an item s/he would not have noticed otherwise, are obviously more valuable. In many of the works we reviewed, evaluation was based on a one-time user survey. Longer term evaluation is required as the results may substantially change over time. Techniques that learn and adapt over time based on user behavior are going to be essential. Additionally, evaluation that examines the broader effect of the recommendation on the surrounding ecosystem of users, as demonstrated in [21, 31, 47, 91] is a highly desirable direction for SRS evaluation. This requires new tools and creative thinking to be brought into the existing evaluation methods.

Cross-Domain Analysis As we discussed, migrating data from one social media service to another may go a long way enhancing recommendations and help deal with the cold start problem for new users. Indeed, using another site's network, tags, and other types of information have been performed by various previous systems as mentioned in this chapter. Yet, social media sites differ in many aspects. It is not certain that one's travel network can serve as a reliable source of recommendation for recipes. Similarly, the tags used in a news site context are not necessarily valuable for video recommendation. More research is due to explore the common and different among social media systems and when information can effectively port from one application to another to be used for recommendation. Cross-domain recommendations in RS have always been harder to explore since they require richer datasets and involve more complex use cases and research questions (Chapter “Design and Evaluation of Cross-Domain Recommender Systems”). As social media continues to evolve, it will be more important to explore and better understand these complexities.

Extraction from User-Generated Content User-generated content (UGC) is abundant across social media sites, in various forms, such as blog posts, forum threads, community question-answering, reviews, and others. Due to its sheer volume, it is often hard for social media users to find the information they seek for within UGC. Common methods allow searching UGC or sorting it by different criteria such as number of votes or date. Recent work also suggested the extraction of specific information types from UGC to provide specific types of information needs. The extraction of tips—short and practical pieces of advice—has been proposed for CQA content (“how to” questions, in particular) [107], products [58] and, prominently, in the travel domain [46, 113]. Other types of extracted information include personal experiences [82], fun trivial facts [104], locations [15], and product descriptions [76, 77]. Such extractions practically yield recommendations of new content types, such as tips, experiences, descriptions, or trivia facts, to social media users, allowing them to enjoy different aspects unlocked in UGC, which they have no realistic chance to discover by merely traversing or searching. Further use of state-of-the-art NLP methods, such as summarization and translation, can help transform free-form UGC, which is one of the most prominent characteristics of social media, into a source of multiple information types that address different needs of social media users.

Multistakeholder Recommendation As most recommender systems apply personalization techniques, they tend to focus on the preferences and needs of a single user. Multistakeholder recommendation has emerged d as a unifying framework for describing and understanding recommendation settings where the end user is not the sole focus [1] (Chapter “Multistakeholder Recommender Systems”). System objectives, such as fairness [112], balance, and profitability receive attention, as well as concerns from other stakeholders, such as the providers or sellers of items being recommended. This extension spans beyond people and group recommendation [41], and requires new optimization targets and evaluation metrics. While multistakeholder issues have surfaced regularly in the history of recommender systems research and have been a constant constraint in fielded applications, the recognition of common threads and research questions has been a more recent occurrence.

References

1. H. Abdollahpouri, G. Adomavicius, R. Burke, I. Guy, D. Jannach, T. Kamishima, J. Krasnodebski, L. Pizzato, Multistakeholder recommendation: survey and research directions. User Model. User-Adapt. Interact. **30**(1), 127–158 (2020)
2. D. Agarwal, B.-C. Chen, Q. He, Z. Hua, G. Lebanon, Y. Ma, P. Shivaswamy, H.-P. Tseng, J. Yang, L. Zhang. Personalizing LinkedIn feed, in *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '15* (Association for Computing Machinery, New York, 2015), pp. 1651–1660
3. N.A. Alawad, A. Anagnostopoulos, S. Leonardi, I. Mele, F. Silvestri, Network-aware recommendations of novel tweets, in *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16* (Association for Computing Machinery, New York, 2016), pp. 913–916
4. O. Alonso, S. Mizzaro, Can we get rid of TREC assessors? Using Mechanical Turk for relevance assessment, in *Proceedings of the SIGIR 2009 Workshop on the Future of IR Evaluation*, vol. 15 (2009), p. 16
5. J. Arguello, J.L. Elsas, J. Callan, J.G. Carbonell, Document representation and query expansion models for blog recommendation, in *Proceedings of the Second AAAI Conference on Weblogs and Social Media - ICWSM '08* (2008)
6. L. Baltrunas, T. Makcinskas, F. Ricci, Group recommendations with rank aggregation and collaborative filtering, in *Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10* (ACM, New York, 2010), pp. 119–126
7. A. Barua, S.W. Thomas, A.E. Hassan, What are developers talking about? An analysis of topics and trends in stack overflow. Empir. Softw. Eng. **19**(3), 619–654 (2014)
8. T. Belluf, L. Xavier, R. Giglio, Case study on the business value impact of personalized recommendations on a large online retailer, in *Proceedings of the Sixth ACM Conference on Recommender Systems, RecSys '12* (ACM, New York, 2012), pp. 277–280
9. J. Bennett, S. Lanning, The Netflix prize, in *Proceedings of KDD Cup and Workshop*, vol. 2007 (2007), p. 35
10. S. Berkovsky, J. Freyne, Group-based recipe recommendations: analysis of data aggregation strategies, in *Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10* (ACM, New York, 2010), pp. 111–118
11. D.M. Boyd, N.B. Ellison, Social network sites: definition, history, and scholarship. J. Comput. Mediat. Commun. **13**(1), 210–230 (2007)
12. U. Brandes, A faster algorithm for betweenness centrality. J. Math. Sociol. **25**(2), 163–177 (2001)
13. M.J. Brzozowski, D.M. Romero, Who should i follow? Recommending people in directed social networks, in *ICWSM* (2011)
14. M. Buhrmester, T. Kwang, S.D. Gosling, Amazon's Mechanical Turk a new source of inexpensive, yet high-quality, data? Perspect. Psychol. Sci. **6**(1), 3–5 (2011)
15. J. Capdevila, M. Arias, A. Arratia, Geosrs: a hybrid social recommender system for geolocated data. Inf. Syst. **57**, 111–128 (2016)
16. J. Chen, W. Geyer, C. Dugan, M. Muller, I. Guy, Make new friends, but keep the old: recommending people on social networking sites, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '09* (ACM, New York, 2009), pp. 201–210
17. J. Chen, R. Nairn, L. Nelson, M. Bernstein, E. Chi, Short and tweet: experiments on recommending content from information streams, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10* (ACM, New York, 2010), pp. 1185–1194
18. K. Chen, T. Chen, G. Zheng, O. Jin, E. Yao, Y. Yu, Collaborative personalized tweet recommendation, in *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '12* (ACM, New York, 2012), pp. 661–670

19. D. Cosley, D. Frankowski, L. Terveen, J. Riedl, SuggestBot: using intelligent task routing to help people find work in wikipedia, in *Proceedings of the 12th International Conference on Intelligent User Interfaces, IUI '07* (ACM, New York, 2007), pp. 32–41
20. G. Costa, R. Ortale, Model-based collaborative personalized recommendation on signed social rating networks. ACM Trans. Internet Technol. **16**(3), 1–21 (2016)
21. E.M. Daly, W. Geyer, D.R. Millen, The network effects of recommending social connections, in *Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10* (ACM, New York, 2010), pp. 301–304
22. S. Dara, C.R. Chowdary, C. Kumar, A survey on group recommender systems. J. Intell. Inf. Syst. **54**(2), 271–295 (2020)
23. J. Davidson, B. Livingston, D. Sampath, B. Liebald, J. Liu, P. Nandy, T. Van Vleet, U. Gargi, S. Gupta, Y. He, et al., The YouTube video recommendation system, in *Proceedings of the Fourth ACM conference on Recommender Systems - RecSys '10* (2010), pp. 293–296
24. C. Dugan, W. Geyer, D.R. Millen, Lessons learned from blog muse: audience-based inspiration for bloggers, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10* (ACM, New York, 2010), pp. 1965–1974
25. C. Dwyer, Privacy in the age of Google and Facebook. IEEE Technol. Soc. Mag. **30**(3), 58–63 (2011)
26. M. Eirinaki, J. Gao, I. Varlamis, K. Tserpes, Recommender systems for large-scale social networks: a review of challenges and solutions. Future Gener. Comput. Syst. **78**, 413–418 (2018)
27. W. Fan, Y. Ma, D. Yin, J. Wang, J. Tang, Q. Li, Deep social collaborative filtering, in *Proceedings of the 13th ACM Conference on Recommender Systems, RecSys '19* (Association for Computing Machinery, New York, 2019), pp. 305–313
28. A. Felfernig, L. Boratto, M. Stettinger, M. Tkalčič, *Group Recommender Systems: An Introduction* (Springer, Berlin, 2018)
29. M. Fire, L. Tenenboim, O. Lesser, R. Puzis, L. Rokach, Y. Elovici, Link prediction in social networks using computationally efficient topological features, in *2011 IEEE Third International Conference on Privacy, Security, Risk And Trust (Passat), and 2011 IEEE Third International Conference on Social Computing (SocialCom)* (IEEE, Piscataway, 2011), pp. 73–80
30. B. Fogg, A behavior model for persuasive design, in *Proceedings of the 4th International Conference on Persuasive Technology, Persuasive '09* (ACM, New York, 2009), pp. 40:1–40:7
31. J. Freyne, S. Berkovsky, E.M. Daly, W. Geyer, Social networking feeds: recommending items of interest, in *Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10* (ACM, New York, 2010), pp. 277–280
32. J. Freyne, M. Jacovi, I. Guy, W. Geyer, Increasing engagement through early recommender intervention, in *Proceedings of the Third ACM Conference on Recommender Systems, RecSys '09* (ACM, New York, 2009), pp. 85–92
33. S. Garcia Esparza, M.P. O'Mahony, B. Smyth, On the real-time web as a source of recommendation knowledge, in *Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10* (ACM, New York, 2010), pp. 305–308
34. M. Ge, C. Delgado-Battenfeld, D. Jannach, Beyond accuracy: evaluating recommender systems by coverage and serendipity, in *Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10* (ACM, New York, 2010), pp. 257–260
35. W. Geyer, C. Dugan, D.R. Millen, M. Muller, J. Freyne, Recommending topics for self-descriptions in online user profiles, in *Proceedings of the 2008 ACM Conference on Recommender Systems, RecSys '08* (ACM, New York, 2008), pp. 59–66
36. J. Golbeck, Generating predictive movie recommendations from trust in social networks, in *Proceedings of the 4th International Conference on Trust Management, iTrust'06* (Springer, Berlin, 2006), pp. 93–104
37. J.A. Golbeck, Computing and applying trust in web-based social networks. PhD thesis, College Park, 2005, AAI3178583

38. G. Groh, C. Ehmig, Recommendations in taste related domains, in *Proceedings of the 2007 International ACM Conference on Conference on Supporting Group Work - GROUP '07* (2007), pp. 127–136
39. P. Gupta, A. Goel, J. Lin, A. Sharma, D. Wang, R. Zadeh, WTF: the who to follow service at twitter, in *Proceedings of the 22Nd International Conference on World Wide Web, WWW '13, Republic and Canton of Geneva*. International World Wide Web Conferences Steering Committee (2013), pp. 505–514
40. D.F. Gurini, F. Gasparetti, A. Micarelli, G. Sansonetti, Temporal people-to-people recommendation on social networks with sentiment-based matrix factorization. Future Gener. Comput. Syst. **78**, 430–439 (2018)
41. I. Guy, People recommendation on social media, in *Social Information Access* (Springer, Berlin, 2018), pp. 570–623
42. I. Guy, U. Avraham, D. Carmel, S. Ur, M. Jacovi, I. Ronen, Mining expertise and interests from social media, in *Proceedings of the 22Nd International Conference on World Wide Web, WWW '13, Republic and Canton of Geneva*. International World Wide Web Conferences Steering Committee (2013), pp. 515–526
43. I. Guy, D. Carmel, Social recommender systems, in *Proceedings of the 20th International Conference Companion on World Wide Web - WWW '11* (2011), pp. 283–284
44. I. Guy, R. Levin, T. Daniel, E. Bolshinsky, Islands in the stream: a study of item recommendation within an enterprise social stream, in *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '15* (Association for Computing Machinery, New York, 2015), pp. 665–674
45. I. Guy, V. Makarenkov, N. Hazon, L. Rokach, B. Shapira, Identifying informational vs. conversational questions on community question answering archives, in *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM '18* (Association for Computing Machinery, New York, 2018), pp. 216–224
46. I. Guy, A. Mejer, A. Nus, F. Raiber, Extracting and ranking travel tips from user-generated reviews, in *Proceedings of the 26th International Conference on World Wide Web, WWW '17, Republic and Canton of Geneva, CHE*. International World Wide Web Conferences Steering Committee (2017), pp. 987–996
47. I. Guy, I. Ronen, E. Kravi, M. Barnea, Increasing activity in enterprise online communities using content recommendation. ACM Trans. Comput. Hum. Interact. **23**(4), 1–28 (2016)
48. I. Guy, I. Ronen, A. Raviv, Personalized activity streams: sifting through the river of news, in *Proceedings of the Fifth ACM Conference on Recommender Systems, RecSys '11* (ACM, New York, 2011), pp. 181–188
49. I. Guy, I. Ronen, E. Wilcox, Do You Know?: recommending people to invite into your social network, in *Proceedings of the 14th International Conference on Intelligent User Interfaces, IUI '09* (ACM, New York, 2009), pp. 77–86
50. I. Guy, I. Ronen, N. Zwerdling, I. Zuyev-Grabovitch, M. Jacovi, What is your organization "like"? A study of liking activity in the enterprise, in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI '16* (Association for Computing Machinery, New York, 2016), pp. 3025–3037
51. I. Guy, B. Shapira, From royals to vegans: characterizing question trolling on a community question answering website, in *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR '18* (Association for Computing Machinery, New York, 2018), pp. 835–844
52. I. Guy, T. Steier, M. Barnea, I. Ronen, T. Daniel, Swimming against the streamz: search and analytics over the enterprise activity stream, in *Proceedings of the 21st ACM International Conference on Information and Knowledge Management, CIKM '12* (Association for Computing Machinery, New York, 2012), pp. 1587–1591
53. I. Guy, N. Zwerdling, D. Carmel, I. Ronen, E. Uziel, S. Yogev, S. Ofek-Koifman, Personalized recommendation of social software items based on social relations, in *Proceedings of the Third ACM Conference on Recommender Systems, RecSys '09* (ACM, New York, 2009), pp. 53–60

54. I. Guy, N. Zwerdling, I. Ronen, D. Carmel, E. Uziel, Social media recommendation based on people and tags, in *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '10* (ACM, New York, 2010), pp. 194–201
55. J. Hannon, M. Bennett, B. Smyth, Recommending Twitter users to follow using content and collaborative filtering approaches, in *Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10* (ACM, New York, 2010), pp. 199–206
56. M.R. Hasan, A.J. Kumar, Y. Liu, Excessive use of online video streaming services: impact of recommender system use, psychological factors, and motives. Comput. Hum. Behav. **80**, 220–228 (2018)
57. J.L. Herlocker, J.A. Konstan, J. Riedl, Explaining collaborative filtering recommendations, in *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work, CSCW '00* (ACM, New York, 2000), pp. 241–250
58. S. Hirsch, S. Novgorodov, I. Guy, A. Nus, Generating tips from product reviews, in *Proceedings of the Fourteenth ACM International Conference on Web Search and Data Mining, WSDM '21* (Association for Computing Machinery, New York, 2021)
59. M. Jacovi, I. Guy, S. Kremer-Davidson, S. Porat, N. Aizenbud-Reshef, The perception of others: Inferring reputation from social media in the enterprise, in *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW '14* (ACM, New York, 2014), pp. 756–766
60. M. Jamali, M. Ester, A matrix factorization technique with trust propagation for recommendation in social networks, in *Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10* (ACM, New York, 2010), pp. 135–142
61. A. Jameson, S. Baldes, T. Kleinbauer, Two methods for enhancing mutual awareness in a group recommender system, in *Proceedings of the Working Conference on Advanced Visual Interfaces, AVI '04* (ACM, New York, 2004), pp. 447–449
62. R. Jäschke, L. Marinho, A. Hotho, L. Schmidt-Thieme, G. Stumme, Tag recommendations in folksonomies, in *Knowledge Discovery in Databases: PKDD 2007* (Springer, Berlin, 2007), pp. 506–514
63. O. Kaser, D. Lemire, Tag-cloud drawing: algorithms for cloud visualization. Preprint, arXiv cs/0703109 (2007)
64. H. Kautz, B. Selman, M. Shah, Referral web: combining social networks and collaborative filtering. Commun. ACM **40**(3), 63–65 (1997)
65. A. Kittur, E.H. Chi, B. Suh, Crowdsourcing user studies with Mechanical Turk, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '08* (ACM, New York, 2008), pp. 453–456
66. J. Kunkel, T. Donkers, L. Michael, C.-M. Barbu, J. Ziegler, Let me explain: impact of personal and impersonal explanations on trust in recommender systems, in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI '19* (Association for Computing Machinery, New York, 2019), pp. 1–12
67. R. Lempel, S. Moran, SALSA: the stochastic approach for link-structure analysis. ACM Trans. Inf. Syst. **19**(2), 131–160 (2001)
68. K. Lerman, Social networks and social information filtering on digg, in *Proceedings of the first AAAI Conference on Weblogs and Social Media - ICWSM '07* (2007)
69. J. Leskovec, D. Huttenlocher, J. Kleinberg, Predicting positive and negative links in online social networks, in *Proceedings of the 19th International Conference on World Wide Web, WWW '10* (ACM, New York, 2010), pp. 641–650
70. D. Liben-Nowell, J. Kleinberg, The link-prediction problem for social networks. J. Am. Soc. Inf. Sci. Technol. **58**(7), 1019–1031 (2007)
71. J. Liu, P. Dolan, E.R. Pedersen, Personalized news recommendation based on click behavior, in *Proceedings of the 15th International Conference on Intelligent User Interfaces - IUI '10* (2010), pp. 31–40

72. C. Macdonald, I. Ounis, The TREC blogs06 collection: creating and analysing a blog test collection. Department of Computer Science, University of Glasgow Tech Report TR-2006-224, vol. 1 (2006), pp. 3–1
73. J.F. McCarthy, T.D. Anagnost, MusicFX: an arbiter of group preferences for computer supported collaborative workouts, in *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work, CSCW '98* (ACM, New York, 1998), pp. 363–372
74. D.W. McDonald, M.S. Ackerman, Just talk to me: a field study of expertise location, in *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work, CSCW '98* (ACM, New York, 1998), pp. 315–324
75. S.M. McNee, J. Riedl, J.A. Konstan, Being accurate is not enough: how accuracy metrics have hurt recommender systems, in *CHI '06 Extended Abstracts on Human Factors in Computing Systems, CHI EA '06* (ACM, New York, 2006), pp. 1097–1101
76. S. Novgorodov, I. Guy, G. Elad, K. Radinsky, Generating product descriptions from user reviews, in *The World Wide Web Conference, WWW '19* (Association for Computing Machinery, New York, 2019), pp. 1354–1364
77. S. Novgorodov, I. Guy, G. Elad, K. Radinsky, Descriptions from the customers: comparative analysis of review-based product description generation methods. ACM Trans. Internet Technol. **20**(4), 1–31 (2020)
78. T. Olsson, J. Huhtamäki, H. Kärkkäinen, Directions for professional social matching systems. Commun. ACM **63**(2), 60–69 (2020)
79. T. O'Reilly, *What Is Web 2.0* (O'Reilly Media, Sebastopol, 2009)
80. T. Paek, M. Gamon, S. Counts, D.M. Chickering, A. Dhesi, Predicting the importance of newsfeed posts and social network friends, in *AAAI*, vol. 10 (2010), pp. 1419–1424
81. I. Palomares, C. Porcel, L. Pizzato, I. Guy, E. Herrera-Viedma, Reciprocal recommender systems: analysis of state-of-art literature, challenges and opportunities on social recommendation. Preprint, arXiv:2007.16120 (2020)
82. K.C. Park, Y. Jeong, S.-H. Myaeng, Detecting experiences from weblogs, in *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics* (2010), pp. 1464–1472
83. O. Phelan, K. McCarthy, B. Smyth, Using Twitter to recommend real-time topical news, in *Proceedings of the Third ACM Conference on Recommender Systems, RecSys '09* (ACM, New York, 2009), pp. 385–388
84. G. Piao, J.G. Breslin, Learning to rank tweets with author-based long short-term memory networks, in *Web Engineering*, ed. by T. Mikkonen, R. Klamma, J. Hernández (Springer, Berlin, 2018), pp. 288–295
85. L. Pizzato, T. Rej, T. Chung, I. Koprinska, J. Kay, RECON: a reciprocal recommender for online dating, in *Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10* (ACM, New York, 2010), pp. 207–214
86. D. Quercia, L. Capra, FriendSensing: recommending friends using mobile phones, in *Proceedings of the Third ACM conference on Recommender systems - RecSys '09* (2009), pp. 273–276
87. D.R. Raban, A. Danan, I. Ronen, I. Guy, Impression management through people tagging in the enterprise: implications for social media sampling and design. J. Inf. Sci. **43**(3), 295–315 (2017)
88. P. Resnick, H.R. Varian, Recommender systems. Commun. ACM **40**(3), 56–58 (1997)
89. I. Ronen, I. Guy, E. Kravi, M. Barnea, Recommending social media content to community owners, in *Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '14* (ACM, New York, 2014), pp. 243–252
90. R.M. Ryan, E.L. Deci, Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. Am. Psychol. **55**(1), 68 (2000)
91. A. Said, A. Bellogín, You are what you eat! tracking health through recipe interactions, in *6th RecSys Workshop on Recommender Systems and the Social Web, RSWeb '14* (2014), p. 4

92. J. Sanz-Cruzado, P. Castells, Enhancing structural diversity in social networks by recommending weak ties, in *Proceedings of the 12th ACM Conference on Recommender Systems, RecSys '18* (Association for Computing Machinery, New York, 2018), pp. 233–241
93. S. Scellato, A. Noulas, C. Mascolo, Exploiting place features in link prediction on location-based social networks, in *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '11* (ACM, New York, 2011), pp. 1046–1054
94. S. Sen, J. Vig, J. Riedl, Tagommenders: connecting users to items through tags, in *Proceedings of the 18th International Conference on World Wide Web, WWW '09* (ACM, New York, 2009), pp. 671–680
95. S. Sidana, A combination of classification based methods for recommending tweets, in *Proceedings of the Recommender Systems Challenge 2020* (Association for Computing Machinery, New York, 2020), pp. 1–5
96. B. Sigurbjörnsson, R. van Zwol, Flickr tag recommendation based on collective knowledge, in *Proceedings of the 17th International Conference on World Wide Web, WWW '08* (ACM, New York, 2008), pp. 327–336
97. R.R. Sinha, K. Swearingen, Comparing recommendations made by online systems and friends, in *DELOS Workshop: Personalisation and Recommender Systems in Digital Libraries*, vol. 106 (2001)
98. P. Sun, L. Wu, M. Wang, Attentive recurrent social recommendation, in *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR '18* (Association for Computing Machinery, New York, 2018), pp. 185–194
99. I. Szpektor, Y. Maarek, D. Pelleg, When relevance is not enough: promoting diversity and freshness in personalized question recommendation, in *Proceedings of the 22nd International Conference on World Wide Web, WWW '13, Republic and Canton of Geneva*. International World Wide Web Conferences Steering Committee (2013), pp. 1249–1260
100. H. Tahmasebi, R. Ravanmehr, R. Mohamadrezaei, Social movie recommender system based on deep autoencoder network using twitter data. Neural Comput. Appl. **33**, 1607–1623 (2021)
101. A. Taneja, A. Arora, Modeling user preferences using neural networks and tensor factorization model. Int. J. Inf. Manag. **45**, 132–148 (2019)
102. L. Terveen, D.W. McDonald, Social matching: a framework and research agenda. ACM Trans. Comput. Hum. Interact. **12**(3), 401–434 (2005)
103. C.-H. Tsai, P. Brusilovsky, Explaining recommendations in an interactive hybrid social recommender, in *Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI '19* (Association for Computing Machinery, New York, 2019), pp. 391–396
104. D. Tsurel, D. Pelleg, I. Guy, D. Shahaf, Fun facts: automatic trivia fact extraction from wikipedia, in *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining, WSDM '17* (Association for Computing Machinery, New York, 2017), pp. 345–354
105. D. Wang, D. Pedreschi, C. Song, F. Giannotti, A.-L. Barabasi, Human mobility, social ties, and link prediction, in *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '11* (ACM, New York, 2011), pp. 1100–1108
106. J. Wang, Y. Zhang, C. Posse, A. Bhasin, Is it time for a career switch? in *Proceedings of the 22Nd International Conference on World Wide Web, WWW '13, Republic and Canton of Geneva*. International World Wide Web Conferences Steering Committee (2013), pp. 1377–1388
107. I. Weber, A. Ukkonen, A. Gionis, Answers, not links: extracting tips from yahoo! answers to address how-to web queries, in *Proceedings of the Fifth ACM International Conference on Web Search and Data Mining, WSDM '12* (Association for Computing Machinery, New York, 2012), pp. 613–622
108. H. Wu, K. Yue, Y. Pei, B. Li, Y. Zhao, F. Dong, Collaborative topic regression with social trust ensemble for recommendation in social media systems. Knowl. Based Syst. **97**, 111–122 (2016)
109. D. Yang, C. Huang, M. Wang, A social recommender system by combining social network and sentiment similarity: a case study of healthcare. J. Inf. Sci. **43**(5), 635–648 (2017)

110. X. Yang, Y. Guo, Y. Liu, H. Steck, A survey of collaborative filtering based social recommender systems. Comput. Comm. **41**, 1–10 (2014)
111. A. Yogev, I. Guy, I. Ronen, N. Zwerdling, M. Barnea, Social media-based expertise evidence, in *ECSCW 2015: Proceedings of the 14th European Conference on Computer Supported Cooperative Work, 19–23 September 2015, Oslo* (Springer, Berlin, 2015), pp. 63–82
112. Y. Zheng, T. Dave, N. Mishra, H. Kumar, Fairness in reciprocal recommendations: a speed-dating study, in *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization, UMAP '18* (Association for Computing Machinery, New York, 2018), pp. 29–34
113. D. Zhu, T. Lappas, J. Zhang, Unsupervised tip-mining from customer reviews. Decis. Support Syst. **107**, 116–124 (2018)