Recommendations beyond the Ratings Matrix

Eirini Ntoutsi
Leibniz Universität Hannover, Germany
ntoutsi@kbs.uni-hannover.de

Kostas Stefanidis
ICS-FORTH, Greece
kstef@ics.forth.gr

ABSTRACT

Recommender systems have become indispensable for several Web sites, such as Amazon, Netflix and Google News, helping users navigate through the abundance of available choices. Although the field has advanced impressively in the last years with respect to models, usage of heterogeneous information, such as ratings and text reviews, and recommendations for modern applications beyond purchases, almost all of the approaches rely on the data that exist within the recommender and on user explicit input. In a rapidly connected world, though, information is not isolated and does not necessarily lie in the database of a single recommender. Rather, Web offers tremendous amount of information on almost everything, from items to users and their tendency to certain items, but also information on general trends and demographics. We envision an out-of-the-box recommender system that exploits the existing information in a recommender, namely, items, users and ratings, but also explores new sources of information out of the database, like user online traces and online discussions about data items, and exploits them for better and innovative recommendations. We discuss the challenges that such an out-of-the-box approach effects and how it reshapes the field of recommenders.

Keywords

Recommendation systems;big data; social data; information hunting; data-driven innovation

1. INTRODUCTION

Recommender systems have become indispensable for several Web sites, such as Amazon, Netflix and Google News, helping users navigate through the abundance of items. Recommender systems facilitate the selection of items by users by issuing recommendations for items they might like. There are different recommendation approaches, like neighborhood-based approaches [5] and model-based ones [16]. There is also a lot of work on specific aspects of recommendations [7], like the cold start problem, the long tail problem and the evaluation of the recommended items in terms of a variety of parameters, like relevance, surprise and serendipity. More recently, a lot of approaches have been proposed that combine numerical ratings with textual reviews [10]. Also, nowadays, recommendations have more broad applications, beyond products, like news recommendations, links (friends) recommendations [19] and more innovative ones like query recommendations [6], medicine recommendations [8] etc.

In all these cases though, it is assumed that the required information for computing recommendations, let it be numerical ratings or reviews, comprises the input to the recommender system and all the different challenges are tackled upon this data. Typically, the user is the *curator* of the information, in the sense that he explicitly adds ratings and/or reviews to the recommender. So the amount, the quality and the up-to-date state of the data depend really on user engagement. Moreover, recommended items are limited on the available options within the recommender.

In a rapidly connected world though, information is not isolated and does not necessarily lie in the data storage infrastructure of a single recommender. Rather, the Web offers tremendous amount of information on almost everything, from items to users and user tendency to certain items, but also information on general trends and demographics. Actually, the Web is what we nowadays take advantage of as humans when looking for information and suggestions, e.g., when we want to buy a book, watch a movie, buy a dress, look for medical advice or search for recipes. On one hand, we do rely on our close network, like family and friends, but pretty often we look online to also consider the crowd (e.g., general trends on clothing) or specific experts (e.g., favorite food bloggers) and reliable sources of information (e.g., Wikipedia). Also, it is often the case that we first search online to get an idea, and then we elaborate further with friends and family to finalize our decision.

Inspired by this human's approach to information hunting, we propose an *out-of-the-box recommendation approach* that exploits the existing information in a recommender, namely ratings, items and users, but explores as well new sources of information out of the database, like users' online traces and online discussion on items, and exploits them for achieving better and out-of-the-recommender box recommendations.

Getting out of the box introduces many new challenges, like what data to look for and where, how to combine the heterogeneous information from the Web and deciding on whether the data are reliable or not. Moreover, such an approach requires a revision of the recommendation process with respect to issues like keeping track of the user history,

DDI'16, May 22-25, 2016, Hannover, Germany

DOI: http://dx.doi.org/10.1145/2911187.2914580

including his within and out of the box behavior, for identifying changes in his tastes and periodicity in his habits, and explaining recommendations to the end-user based on very diverse data sources. Implementation of such an approach also raises a lot of efficiency issues like how often or when information hunting should take place and whether the enriched information should be stored in the recommender. Many challenges also arise from a business perspective as the envisioned system puts the user in the foreground and not the specific business, as it is traditionally done in recommenders.

The rest of the paper is organized as follows. Section 2 overviews the existing ways for producing recommendations, and Section 3 introduces the "getting-out-of-the-box" approach. Section 4 discusses the new challenges that arise, while Section 5 analyzes their effect on recommendations. Conclusions and outlook are presented in Section 6.

2. IN-THE-BOX RECOMMENDATIONS

A recommender system consists typically of a set of items I, a set of users U and the ratings of users for certain items. A user $u \in U$ might rate an item $i \in I$ with a score $rating(u, i)$ in $[0.0, 1.0]$. Typically, the cardinality of the item set I is very high and users rate only a few items. For the items unrated by the users, recommender systems estimate a relevance score: $relevance(u, i)$, $u \in U$, $i \in I$. There are different ways to estimate the relevance score of an item for a user. In the *content-based approaches*, the estimation of the relevance score of a user for an item is based on the ratings that the user itselft has assigned to similar items, whereas in *collaborative filtering approaches*, the relevance score is predicted using previous ratings for the item by similar users. Similar items and users are located via *similarity functions* that evaluate the proximity between items and users, respectively. The simplest approach for finding similar users or similar items is by linear scanning the whole database, which is a costly process. More efficient approaches exist, for example, approaches that employ user models, derived through, e.g., clustering [11], [12], for prediction.

Apart from recommendations for single users, there are cases, such as visiting a restaurant or selecting a holiday destination, where a group of people participates in an activity. For such cases, *group recommendations* try to satisfy the preferences of all the group members (e.g., [15, 11]). Three different designs are, in general, employed for aggregating the preferences of the users within a group: (i) the *least misery* design, capturing cases where strong user preferences act as a veto (e.g., do not recommend steakhouses to a group when a member is vegetarian), (ii) the *fair* design, capturing more democratic cases where the majority of the group members is satisfied, and (iii) the *most optimistic* design, capturing cases where the more satisfied member of the group acts as the most influential one (e.g., recommend a movie to a group when a member is highly interested in it and the rest have reasonable satisfaction).

A different aspect of group recommendations appears when specific *constraints* apply to the members of the group; constraints refer to preferences that the members of the group express for the other participants (e.g., [13]). For example, a vacation package may seem more attractive to a user, if the other members of the group are of a similar age, whereas a course may be recommended to a group of students that have similar or diverse backgrounds depending on the scope of the course. Constraints may describe limitations from the user/customer or the system/company perspective. In the latter, constraints refer to a set of properties that the group under construction must satisfy, expressing the requirements of the company concerning the group that an item is targeting on.

Since users usually have different preferences under different circumstances, for both single and group recommendations, context-based recommendations have been proposed [2].

3. GETTING OUT OF THE BOX

Most of the existing recommendation applications are closed, in the sense that they profile the user behavior only within their application. For example, Netflix exploits the watching history of its users, Amazon the purchase history of its users and so forth. However, Netflix or Amazon comprise just one of the many stops of a user in the Web journey and therefore their view of the user is quite limited to his within-the-system behavior. Users typically use different tools to fulfill their needs; in terms of music for example, they might use Spotify, YouTube, MusicLoad etc. However, except for the dedicated music websites, users share, like and comment on music videos in Facebook, Twitter or Google+. To profile a user therefore, one should follow a more holistic approach by looking at the overall online user's presence, instead of within a single application/website. A similar approach should be followed for the recommended items; a user might be interested in items out of the recommender-box.

Such a holistic approach does not only offer a better profiling of the users and a wider range of item choices, but it also allows to deal with traditional recommendation problems, like data sparsity and the cold start problem. Users that are new to the application and therefore not profiled yet, are probably "old" Web users and might have left their "traces" in the Web through e.g., information sharing, comments, posts etc. Old to the application users might also suffer from data sparsity as it is quite common nowadays for users to use multiple systems to fulfill their needs, e.g., one can watch movies at both Netflix and Amazon. Moreover, users use multiple networks to share information on their preferences, e.g., for movies or concerts. For instance, a user might tweet about his favorite movie or comment on its YouTube trailer. So, focusing only on the within-the-recommender information, is very restrictive as users' behavior is manifested mainly outside of the specific application due to the multiple user presence in the Web and the typically low user engagement with the application. Increasing user engagement might be a solution however it is costly and difficult due to the variety of products and services in the Web. What is more feasible is a change of perspective from the recommender. Instead of waiting for the user to actively participate in the application (the typical existing passive approach), the recommender should take a more active approach by getting out of the box to find user/item/rating traces in the Web.

We envision a recommender system that, in contrary to existing systems that rely solely on within-the-recommender information, exploits the vast amount of information available in the Web about users, items and the interest of users for certain items. Getting out of the box and building a more complete and holistic profile for the users, as well as a more comprehensive view for the data items, and the user

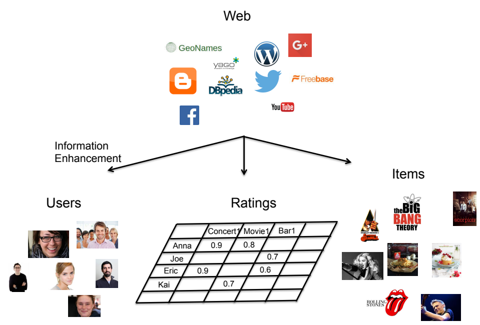

Figure 1: An overview of the envisioned out of the box recommender approach.

preferences for certain items will also allow for better "in-the-box" recommendations. Also, such a solution will benefit all the evolved recommender applications as cross-application recommendations are possible.

An overview of the envisioned framework is given in Figure 1: the bottom part of the figure displays the existing approaches. The connections to the Web depict endless possibilities for data enrichment at all levels: items, users, ratings. Such an open system would potentially result in high-quality recommendations and also in innovative types of recommendations, like cross-application recommendations.

4. INFORMATION HUNTING

Getting out of the box and exploiting the vast amount of information in the Web for recommendations introduces several interesting challenges and therefore, research opportunities. Among the core challenges are: what data to look for and how (Section 4.1) and how to integrate this heterogeneous information (Section 4.2).

4.1 Data enrichment

Lack of data regarding user preferences for the different items, known as the data sparsity problem, is one of the key problems in recommendations. Information hunting aims at filling the missing data and enriching existing ones through external sources like the Web.

In particular, all main entities involved in a recommendation application, namely users, items and ratings, can be semantically enhanced with information implicitly gathered from the Web. Regarding users, we can employ information aggregated from external sources (e.g., social networks). That is, instead of using the plain information that a user gives for himself to a recommender system, we can exploit information available in numerous external sources, such as Facebook, LinkedIn, Fourthsquare and Amazon [18]. The motivation behind this approach, is that a user describes and expresses himself differently in different networks, depending on the domain. Therefore, by examining the different social networks, we can identify different user interests, user activities, information about places he visited, and so forth. In that sense, the complete user profile is revealed by combining information from all the different sources. User's social profile can be also extended to bring together different social networks. The extended graph can be used to infer the user's neighborhood, for instance, when following a collaborative filtering approach.

To enrich items' descriptions, we can exploit (semantic) information retrieved from the Web, such as Web pages, thesauri or ontologies, and published results and reports.This way, items can be annotated using, for example, terms from ontologies and other semantic resources to enhance their description quality. Furthermore, the plethora of well-organized information over the Web in collectively maintained knowledge repositories, such as Wikipedia and LibraryThing, can be used for correlating and computing similarities between data items. Online discussions about certain items could be also employed, as people nowadays tend to provide detailed comments on the different aspects of the items.

Regarding user preferences for certain items, we can rely on his online activities that involve the specific aspect. For example, a user may express an interest in Twitter about a movie, he might rate movie's trailer in YouTube and share the link in Facebook. These actions might reveal user's preference for the item. Except for filling missing ratings, existing ratings could be also enhanced with respect to their context (e.g., time or place) and ratings criteria (e.g., director counts more than actors). Contextual recommendations have been already studied [2], however they also rely on user explicit feedback. We propose here an indirect inference of the context through, e.g., the user reviews for certain items and possible geo-location and network information (e.g., geo-annotated posts with co-tagged friends). Multi-criteria ratings [1] rely also on user explicit subratings for different aspects of the items and are suffering from lack of data as the item space is further expanded due to the sub-criteria. With the abundance of free text reviews nowadays, we can implicitly extract both the aspects of the items that are of interest to a user and their associated ratings/sentiment, using NLP and sentiment analysis techniques.

4.2 Data integration

Even if we are able to identify additional information for users, items and preferences, appearing outside the recommender system, it is challenging to understand which pieces of information refer to same entities, so as to integrate them, in order to manage and further process them. Especially recently, numerous bases on the Web offer comprehensive, machine-readable descriptions of a large variety of real-world entities (e.g., users, items, or even profiles and user preferences) published, typically, as Linked Data. Such bases (e.g., DBpedia, YAGO and Freebase) may provide multiple, non-identical descriptions for the same entities. Entity resolution aims to identify different descriptions that refer to the same entity, and emerges as a central data-processing task for an *entity-centric organization* of Web data [4]. It is needed to enrich interlinking of data elements describing entities, so that the Web of data can be accessed by machines as a *global data space* allowing the use of standard languages.

Although entity resolution has attracted significant attention from many researchers in information systems, database and machine-learning communities, there are new challenges stemming from the Web openness in describing a multitude of entity types across domains. The *scale*, *diversity* and *graph structuring* of descriptions challenge the core entity resolution tasks, namely, (i) how descriptions can be *effectively compared for similarity* and (ii) how resolution algorithms can *efficiently filter the candidate pairs* of descriptions that need to be compared.

In a multi-type and large-scale entity resolution, we need to examine whether two descriptions are *somehow similar* without resorting to domain-specific similarity functions and mapping rules. Furthermore, the resolution of some entity descriptions might influence the resolution of other neighbourhood descriptions. This setting clearly goes beyond deduplication (or record linkage) of collections of descriptions usually referring to a single entity type that slightly differ only in their attribute values. It essentially requires leveraging similarity of descriptions both on their *content* and *structure*. It also forces us to revisit traditional resolution workflows consisting of separate *indexing* (for pruning the number of candidate pairs) and *matching* (for resolving entity descriptions) phases.

5. RECOMMENDATIONS REVISITED

Getting out of the box and exploiting the huge amounts of information available outside the data repository of a recommender system introduces interesting research opportunities, and requires revising the purposes and goals towards producing recommendations. In this section, (i) we investigate cases in which the recommendations output is not restricted within a domain, (ii) we discuss the lifelong tracking of users, which implies an extensive knowledge about user tastes and preferences, but also imposes new challenges for the recommenders, like dealing with changes in user profiles, (iii/iv) we come upon techniques for exploration and visualization that facilitate and guide users to focus on the relevant aspects of their search needs, and (v) we study the notion of re-finding recommendations, that is, a specific way for exploring suggestions, produced and consumed by a user in the past.

5.1 Cross-domain recommendations

Traditionally, recommendations are produced within a domain, i.e., when asking for movies or job vacancies, the suggestions consist only of movies or jobs. When a database contains data items from different domains, this paradigm can be extended, so as to support *cross-domain recommendations*. For example, *packet recommendations* (e.g., [14]) produce composite items consisting of a central item, possibly in the main domain of interest for a user, and a set of satellite items from different domains compatible with the central item. Compatibility can be assumed either as soft (e.g., other books that are often purchased together with the movie being browsed) or hard (e.g., battery packs that must be compatible with a laptop or a travel destination that must be within a certain distance from the main destination).

Interestingly, the notion of cross-domain recommendations can be extended, so as to support recommending packages, consisting of data items from different domains, appearing in different systems or warehouses, which are located by exploiting new sources of information integrated with the data available at the recommender system. Examples are similar as in the previous scenario. When for instance selecting CDs, you might also find proposals about books. However, in this case, books can be originated from a company outside the record company, and identified by taking into consideration the user's likes and posts in social networks, in addition to the traditional used information, i.e., similarities between users and data items. Moving forward, given that the granularity of a user's taste that is captured by his profile might be, many times, too coarse, recommender systems could help users to express their needs by allowing them to provide examples originated from different domains, based on which the system's suggestions will be identified.

5.2 Lifelong recommendations and learning

The majority of recommendation approaches considers a static case, i.e., it requires the whole set of data (users, items and preferences) as input, which is obsolete nowadays, due to the huge amount of generated data and the lifelong tracking of users presence online. Keeping track of the user history, including his within the box and out of the box behavior, does not only result in more data for recommendations (useful for e.g., tackling the sparsity problem), but it also allows for the study of possible changes in user tastes and identifying periodicity in his habits. Change is natural over time; in terms of the recommender this means that it should adapt to such changes in order to provide up to date recommendations to its users. Data ageing is a typical way to deal with drifts and shifts in user behavior by downgrading historical data as obsolete and paying more attention to recent ones that reflect the current user profile best [17]. More work should be devoted to this direction, as the changing pace nowadays is very fast. Existing results on the effect of time in the quality of recommendations are some times contradictory. Approaches that discard past instances have been criticized as loosing too much signal and more elaborate methods exist that separate transient factors from lasting ones [9]. For a detailed observation and tracking of the manner in which a user behaves, especially when we care for particular occasions or circumstances, the notion of context, such as location and accompanying people, can been employed as well. Recently, [3] present the first attempt for implicitly extracting such sort of information, by employing online reviews. However, we should also take into account that a long term user monitoring implies an extensive knowledge about his tastes and preferences, which might result in privacy risks for the user.

5.3 Interactive exploration

Plainly, huge amounts of data are available today to a more diverse and less technically oriented audience; notably, the Web represents the largest and arguably the most complex repository of content approachable to a wide spectrum of users. Taking into account this observation, new forms of data exploration and interaction become increasingly more attractive to aid users navigate through the information space and overcome the challenges of *information overload*. The interaction between the user and the data repositories can be driven directly by the users interpretation of their information need and their information foraging constraints. Alternatively, a search engine can mediate the user-data interactions; the process starts with the user entering query-terms that act as surrogates for the user information goals. Free-text queries allow end-users to start expressing in a simple way their needs, independently from the underlying data model and structure, and from a specific query language. Given a query, the most common strategy has been to present the results as a ranked list. Users have to subsequently peruse the list to satisfy their information needs through browsing the links and by issuing further queries. However, all available data pieces get rapidly diversified both in terms of its complexity, and of the media through which the information is encoded, spanning from large amounts of unstructured and semi-structured data to semantically rich information. This way, increasing demands for sophisticated discovery capabilities are now being imposed by numerous applications in various domains. Although long challenged by works, such as the *berrypicking model*, common systems still assume that the user has a static information need, which remains unchanged during the seeking process. Thus, there is a need to develop novel paradigms for exploratory user-data interactions that emphasize user context and interactivity with the goal of facilitating exploration, interpretation, retrieval, and assimilation of information. Recommendation applications tend to anticipate user needs by automatically suggesting the information which is most appropriate to the users and their current context. In this direction, exploratory search could be fueled by the growth of online social interactions within social networks and Web communities. Many useful facts about entities (e.g., people, locations) and their relationships can be found in a multitude of semi-structured and structured data sources, such as Wikipedia, DBPedia, Freebase, and many others. In overall, novel discovery methods are required to provide highly expressive discovery capabilities over large amounts of entity-relationship data, which are yet intuitive for the end-users.

6. CONCLUSIONS

Recommendations have always been an important area for both research and industry. What we present here is the need to overcome the current restrictions that the recommendation-related data should lie in the data infrastructure of a single recommender and that the user should be the only curator of information in such a system. Rather, the huge amount of heterogeneous data that are continuously collected from the Web call for a reshaping of the recommenders. Data can be enriched from many diverse sources (social media, ontologies, etc.) and at every level (user, item, ratings), providing a much wealthier data input to the recommender and allowing for better recommendations, improving the user experience with the recommender and dealing with traditional problems in the field, like the cold-start problem and data sparsity. Moreover it opens new opportunities for applications involved in the recommendation process as the users can be better profiled and therefore the quality of recommendations can be improved but also cross-platform recommendations are possible.

"There are not enough data to profile a user/item" is not the case anymore; maybe there are not enough within the recommender, but there is an abundance of related data in the Web, or one can employ crowdsourcing to raise data resources. What is challenging now is acquiring useful data (signal out of the noise in the Web) and employing them in combination to the existing data withing the recommender to raise the quality of recommendations and improve the overall user experience with the recommender.

7. REFERENCES

[1] G. Adomavicius, N. Manouselis, and Y. Kwon. Multi-criteria recommender systems. In *Recommender Systems Handbook*. 2011.

[2] G. Adomavicius, R. Sankaranarayanan, S. Sen, and A. Tuzhilin. Incorporating contextual information in recommender systems using a multidimensional approach. *ACM Trans. Inf. Syst.*, 23(1):103–145, 2005.

[3] G. Chen and L. Chen. Recommendation based on contextual opinions. In *UMAP*, 2014.

[4] V. Christophides, V. Efthymiou, and K. Stefanidis. *Entity Resolution in the Web of Data*. Synthesis Lectures on the Semantic Web: Theory and Technology. Morgan & Claypool Publishers, 2015.

[5] C. Desrosiers and G. Karypis. A comprehensive survey of neighborhood-based recommendation methods. In *Recommender Systems Handbook*, pages 107–144. 2011.

[6] M. Eirinaki, S. Abraham, N. Polyzotis, and N. Shaikh. Querie: Collaborative database exploration. *IEEE TKDE*, 26(7):1778–1790, 2014.

[7] P. B. Kantor, L. Rokach, F. Ricci, and B. Shapira. *Recommender systems handbook*. Springer, 2011.

[8] H. Kondylakis, L. Koumakis, E. Kazantzaki, M. Chatzimina, M. Psaraki, K. Marias, and M. Tsiknakis. Patient empowerment through personal medical recommendations. In *MEDINFO*, 2015.

[9] Y. Koren. Collaborative filtering with temporal dynamics. In *SIGKDD*, 2009.

[10] J. McAuley and J. Leskovec. Hidden factors and hidden topics: Understanding rating dimensions with review text. In *RecSys*, 2013.

[11] E. Ntoutsi, K. Stefanidis, K. Nørvåg, and H. Kriegel. Fast group recommendations by applying user clustering. In *ER*, 2012.

[12] E. Ntoutsi, K. Stefanidis, K. Rausch, and H. Kriegel. Strength Lies in Differences: Diversifying friends for recommendations through subspace clustering. In *CIKM*, 2014.

[13] A. G. Parameswaran, P. Venetis, and H. Garcia-Molina. Recommendation systems with complex constraints: A course recommendation perspective. *ACM Trans. Inf. Syst.*, 29(4):20, 2011.

[14] S. B. Roy, S. Amer-Yahia, A. Chawla, G. Das, and C. Yu. Constructing and exploring composite items. In *SIGMOD*, 2010.

[15] S. B. Roy, S. Amer-Yahia, A. Chawla, G. Das, and C. Yu. Space efficiency in group recommendation. *VLDB J.*, 19(6):877–900, 2010.

[16] J. J. Sandvig, B. Mobasher, and R. D. Burke. A survey of collaborative recommendation and the robustness of model-based algorithms. *IEEE Data Eng. Bull.*, 31(2):3–13, 2008.

[17] K. Stefanidis, E. Ntoutsi, M. Petropoulos, K. Nørvåg, and H. Kriegel. A framework for modeling, computing and presenting time-aware recommendations. *T. Large-Scale Data- and Knowledge-Centered Systems*, 10:146–172, 2013.

[18] M. Wischenbart, S. Mitsch, E. Kapsammer, A. Kusel, B. Pröll, W. Retschitzegger, W. Schwinger, J. Schönböck, M. Wimmer, and S. Lechner. User profile integration made easy: model-driven extraction and transformation of social network schemas. In *WWW*, 2012.

[19] Z. Yin, M. Gupta, T. Weninger, and J. Han. LINKREC: a unified framework for link recommendation with user attributes and graph structure. In *WWW*, 2010.