

The Opinionated Recommender

Stephen
Bradshaw
College of Engineering and
Informatics
National University of Ireland Galway,
Ireland

Insight Centre for Data Analytics National
University of Ireland Galway, Ireland
s.bradshaw1@nuigalway.ie

Recommender Systems (RSs) are devices that are used to filter data to combat information overload and provide time saving measures to the user. While RSs have traditionally been done using a content or collaborative based approach, recent times have seen a surge in alternative approaches to try and alleviate some of the traditional problems found there such as the filter bubble, matrix scarcity and cold start issues. Many of these new approaches attempt to lever new sources to provide more accurate recommendations and offset some of these issues. In this paper we will outline some of the current flaws and propose a hypothetical system that will exploit external sources to improve upon the state of the art.

Keywords: Recommender Systems, Topic Extraction, Social Media

1. INTRODUCTION

Traditionally recommendations have been approached in two ways, either on a content basis or through collaborative means. Content based recommendations is done by comparing the attributes of an item and recommending items of a similar nature. This system has been used very successfully for years by companies like Amazon [4]. Issues with content based recommender systems (CBRs) is that of *overspecialisation* [13]. This is when a recommender becomes so attuned to the customer profile that it only recommends things that the customer already has knowledge of; and fails to produce content that is novel or contains serendipitous value. [1] notes that serendipity has two facets, the degree of surprise for the user and usefulness. In addition RSs can suffer from *matrix scarcity*, where there are very few ratings on items/people from which to base recommendations on.

Collaborative recommender systems (CFs) recommend items based on their similarity to those known to be of interest to similar users. They do this by constructing a summary of interests or profile and find other users with similar profiles in order to make further recommendations. The result of using a community of people is that it can offset the issue of *overspecialisation* and can introduce elements of serendipity and novelty in the recommendations, however traditional issues with this approach include

cold start problems [10]. The cold start problem is where new users have not provided enough feedback from which to base new recommendations on. User feedback is itself a big problem in the area of RSs as it can be difficult to get users to interact with a system that does not produce instantaneous and accurate results. In addition, CF approaches are prone to create scenarios where the user is exposed to a *filter bubble* [7]; where a user only hears back things that support her own views, and are not exposed to conflicting points of view.

Our aim is to investigate and design a recommender system that does not fall prey to the issues stated above. We believe that through the application of text content analysis, and harvesting social media we can improve upon the quality of returned information and improved information assimilation can be made through improved presentation on areas of interest. Improvements can come in the form of highlighting sections that have garnered a lot of attention or perhaps through a graph format. To foster a degree of serendipity we aim to incorporate views found on social media platforms (like Reddit and boards.ie). Our work focuses on identifying topics in a text and linking those to topics expressed in opinions. In Section 2 we will talk about some recently proposed solutions to the problems stated and their shortcomings. Section 3 will deal with some approaches to deal with returning personalised content. In Section 4 we will outline a hypothetical

approach that we believe is capable of tackling these issues as well as dealing with traditional flaws with recommender systems.

2. TRUST BASED RECOMMENDER SYSTEMS (TBRS)

One approach to offset some of the issues stated above is proposed by [6]; who argue that CFs treat all profiles as independent entities and fail to acknowledge that they might contain an element of interconnectivity amongst the users. They coin their model *social trust ensemble* which aims to get at the core of CFs principles: which is that we will accept the recommendations of a friend over that of a stranger. To prove their hypothesis they create a use case which recommends films to users. They use Epinions.com as source from which to make their predictions. Epinions is a site where people can give ratings of 1 - 5 on items. In addition the users provide a trusted list of associates, which the authors use to identify trusted connections. They can then increase the coverage of the recommendations by incorporating the trusted friend list into the process. This they state increases the coverage of the system and offsets the sparsity issue. Finally, they use a probabilistic factor analysis model (PFAM) to produce the final results. The strength of this approach is that PFAM is not computationally expensive, and so can be extended to a larger dataset. Their system shows that coverage can be increased using fewer recommendations, however it requires the creation of a trusted list. One could argue that they are merely moving the sparsity problem to that of the trusted friends list. Without having a sufficiently large set of trusted users the system will succumb to the same issue.

Another system that uses opinions as a source to create a trust based recommender system is that one proposed in [5]. They argue that their system is built on a *web of trust* where greater weight is given to recommendations that are coming from trusted associates. One strength of their approach is that it reduces the changes of malicious recommendations by favouring those that come from a trusted source. Like [6] they feel that a trust based recommendation system will alleviate the sparsity problem and new user issue. It is worth noting that Epinions.com is significantly smaller database 83,509 ratings than some other movie review database. Eachmovie contains 2, 811, 983 ratings and Movielens has 1,000,209 ratings [6]. We argue that while there are merits for incorporating trust into recommendations, existing databases that contain both trust metrics and reviews are more sparsely populated than standard movie review databases and thus are not a viable solution to the issue.

[5] conduct a study on the strengths of incorporating trust into recommender systems through a study of Cyworld. Cyworld is a Korean friend site of a similar nature to Facebook. The authors conducted a user study that has 42 members select their favourite skine. A skine is a profile picture accessory that can be purchased to alter the appearance of your profile page. They determined 'trust' levels by measuring the amount of interaction between two users. Interaction is determined from number of messages left on profile walls. The results from their study show that by incorporating additional social data into standard CF approach one can improve the quality of the recommendations.

3. PERSONALISED SEARCH

An approach that aimed to achieve a similar outcome to our own was conducted by Teevan [11]. Her approach automatically created two profiles for the user, one with previous searches and the other on pages visited. She experimented with trying various additional inputs to improve results, namely; processing emails exchanges, calendar dates and documents stored on the users computer. She evaluated her work on a test group and found that user satisfaction can be achieved through personalising the search results returned.

An additionally approach that aimed to gauge how well general opinion can be factored into a personalised recommender was done through the aid of Amazons Mechanical Turk. The Mechanical Turk is a market place run by Amazon that matches workers to employers in the performance of simple repetitive tasks. General tasks can include subjective analysis like which colour is nicer. Workers who perform these tasks are referred to as turk workers. Experimentation was performed on doing two tasks *taste matching* and *taste grokking*. The first applied CF techniques on groups of turk workers and formed groups of people with similar interests, from which to determine new items of interest for a user. The second was to give the turk workers a number of sample items and see if they can recommend additional items that the user would like. The strengths of the authors approach is that they created a novel personalisation system that can make subjective decisions dynamically. They also demonstrated that harvesting the opinion of the crowd can be very beneficial in designing a recommender. Another strength was that the system required little user knowledge. In the next section we will sketch our own proposed approach which also aims to incorporate opinion as a factor to improve recommendations.

4. THE OPINIONATED RECOMMENDER

To utilise content found on social sites we propose applying information retrieval methods to data found on chat forums (Boards.ie / Reddit). Our assumption is that people express views on contemporary issues like those expressed by journalists there. In addition we assume that these views are influenced by content found in daily newspapers. We propose a system that mines news websites for the topics contained within. In our proposed system a user inputs to the system a URL to a news article on sport that she finds interesting. Topic are mined from this article in order to identify topics of interest. Topics are a useful item to determine from stories as they contain a large amount of semantic information.

Topic extraction can potentially be performed using natural language processing techniques. An example for such is Part Of Speech Tagging (POS). POS is where nouns, adjectives and verbs are identified and used for syntax analysis [2]. A sliding window approach [12] can be used to see which noun noun, noun adjective pairings are occurring most frequently and they are in turn used to ascertain the intended topic of the text. Rousseau et al [9] propose an interesting variation on the standard tf-idf approach [8] called *tw-idf* where instead of counting term occurrences *tf*, they count term co-occurrences *tw* and graph them as edges on a node, which might offer an interesting platform to expand upon for determining topics and their interconnectivity. Open questions in this regard include; can the number of topics in an article be determined from density of clusters, and how dense does a cluster have to be before it is considered a topic? Finally is there a means of identifying one term that can accurately express the topic of the assembled words?

An external source such as DBpedia is then used to obtain additional information on the topics. The same process is applied to the chat forums to see if the topics identified are being discussed there. Once the topics have been identified we propose graphing the topics discussed in the chat forums and applying distance metrics to see what additional topics are most closely related to the users' preference topics. A vector is then created storing the additional topics weighted by their proximity to known topics of interest. The vector might be used to inform a decision tree on the users interest; where topics may be the leaves of the tree and the presence of a certain combination of leaves will indicate a user's interest or otherwise. Newspaper conglomerations (such as Google News) are then crawled to see if new hitherto unknown articles can be found. These articles are returned to the user in a hierarchical order of interest. In addition, hot topics or topics that

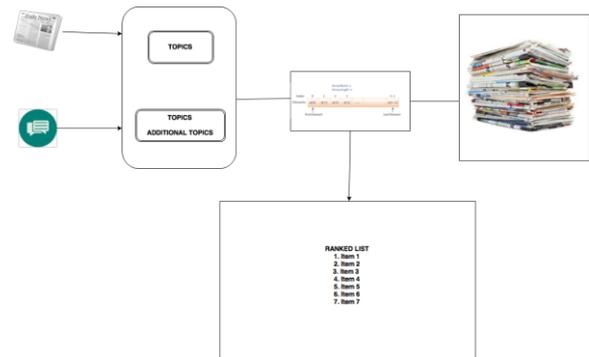


Figure 1: System Architecture

are producing a lot of 'chat' can be highlighted so the user gains an extra level of depth on the issues.

5. DISCUSSION

Traditional approaches to recommender systems have reached the zenith of their ability and there is been a marked increase in the number of researchers looking for additional ways to improve them [9] [11] [12]. External sources such as social media sites are seen as having big potential to be utilised to improve upon recommendations. In addition making recommendations informed by user feedback has shown to be insufficient on its own [3] as the context of a user's situation influences what numeric feedback she will apply. The user may rate a mediocre film more generously if she had watched a very bad film prior and inversely a good film may only receive a middling rating if the film the rater had watched prior was superb. We argue that content based measures can be exploited in the form of content analysis, and that a collaborative-like approach can increase the diversity of recommendation. We propose a system that is based on opinion expressed in chat forums, which can be exploited by applying information retrieval approaches to evaluate the information found there and use it to augment RSs. Our approach aims to avoid some of the standard RS issues such as *cold start*, *information sparsity* and add *serendipity* to the recommendations. Future work includes building a system that incorporate the above outlined approaches and evaluating whether such an approach can improve the coverage as well as satisfaction of a user.

ACKNOWLEDGEMENTS

I would like to thank my supervisor Matthias Nickles for his assistance and feedback. I also would like to thank the reviewers for their insights and helpful pointers.

REFERENCES

- [1] Mouzhi Ge, Carla Delgado-Battenfeld, and Dietmar Jannach. Beyond accuracy: evaluating recommender systems by coverage and serendipity. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 257–260. ACM, 2010.
- [2] Thomas L Griffiths, Mark Steyvers, David M Blei, and Joshua B Tenenbaum. Integrating topics and syntax. In *Advances in neural information processing systems*, pages 537–544, 2004.
- [3] Jonathan L Herlocker, Joseph A Konstan, Loren G Terveen, and John T Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)*, 22(1):5–53, 2004.
- [4] Greg Linden, Brent Smith, and Jeremy York. Amazon. com recommendations: Item-to-item collaborative filtering. *Internet Computing, IEEE*, 7(1):76–80, 2003.
- [5] Fengkun Liu and Hong Joo Lee. Use of social network information to enhance collaborative filtering performance. *Expert systems with applications*, 37(7):4772–4778, 2010.
- [6] Hao Ma, Irwin King, and Michael R Lyu. Learning to recommend with explicit and implicit social relations. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):29, 2011.
- [7] Tien T Nguyen, Pik-Mai Hui, F Maxwell Harper, Loren Terveen, and Joseph A Konstan. Exploring the filter bubble: the effect of using recommender systems on content diversity. In *Proceedings of the 23rd international conference on World wide web*, pages 677–686. ACM, 2014.
- [8] Juan Ramos. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*, 2003.
- [9] François Rousseau and Michalis Vazirgiannis. Graph-of-word and tw-idf: new approach to ad hoc ir. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, pages 59–68. ACM, 2013.
- [10] Guy Shani and Asela Gunawardana. Evaluating recommendation systems. In *Recommender systems handbook*, pages 257–297. Springer, 2011.
- [11] Jaime Teevan, Susan T Dumais, and Eric Horvitz. Personalizing search via automated analysis of interests and activities. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 449–456. ACM, 2005.
- [12] Ville H Tuulos and Henry Tirri. Combining topic models and social networks for chat data mining. In *Proceedings of the 2004 IEEE/WIC/ACM international Conference on Web intelligence*, pages 206–213. IEEE Computer Society, 2004.
- [13] Sung-Shun Weng and Mei-Ju Liu. Feature-based recommendations for one-to-one marketing. *Expert Systems with Applications*, 26(4):493–508, 2004.