

A new level of social search: discovering the user's opinion before he can make one

Mohammad Raza
Microsoft Research Cambridge

April 2011

Abstract

Current approaches to social search are generally based on the *friends' opinion* and *collaborative filtering* paradigms, which have limited effectiveness in terms of both accuracy and coverage of the information that may be relevant to a user. We propose a new kind of social search signal, based on the idea of *prediction extraction*, which is aimed at improving the quality of information available to current methods, as well as introducing new kinds of search methods.

1 Introduction

The idea of *social search* has been receiving a growing amount of interest recently, as can be seen by the strong efforts of companies such as Google and Microsoft to add “social features” to their search engines, especially in collaboration with social networking websites such as Facebook and Twitter [2, 6, 5, 4, 1]. The hope is that social search may hold the key to addressing the “information overload” problem: helping users discover and decide what is personally relevant to them from the enormous amount of information that may be available. Although there are a number of different interpretations of what social search is [8], it generally refers to the ways in which the collective intelligence of a community, such as people in the social graph of the user, can be used to improve the quality of the search process.

In a certain sense, Google's original *PageRank* algorithm could be considered a primitive form of social search, as it assesses the relative importance of pages according to the link structure of the web, that is, the collective opinion of the webmaster community. Many attempts have since been made to more explicitly utilize social media to improve the search process, especially as social networking becomes ever more prevalent. Some notable examples of current social search applications include Bing's recent incorporation of the Facebook “like” data in search results [2, 1], Google's similar effort with the “+1” button [6], *Delicious* [10] (social bookmarking), *Digg* [9] (news voting), *StumbleUpon* [11] (grouping “like-minded” surfers), *Project Emporia* [12] (personalized news) and Amazon's recommendations [13] (“users who bought x also bought y”). Generally, the approach that is taken in today's applications is to elicit the opinions of people in a community, especially the user's social graph, in the form of votes or annotations (e.g. social bookmarking/tagging [14]), and/or using statistical or machine learning techniques to analyse and predict usage patterns (collaborative filtering [15]). However, these “social signals” based on friends' opinions or usage patterns have limited effectiveness, in terms of both accuracy and coverage.

Firstly, even people in the social network of the user are mostly family, past/present colleagues, or casual acquaintances, and not necessarily people with tastes or opinions similar to the user. But even if friends share certain tastes, there may be a myriad of reasons for the user to have a certain opinion or taste, so they may agree with some friends on certain things but heavily disagree on others for different reasons. The same would be true when it comes to analysing usage patterns: agreeing on certain items does not imply agreeing on others. The essential point is that every person has a unique and complex personality, and there is a limit as to how accurately this can be predicted by classifying them with their friends or other users who had similar opinions on certain things.

The other issue is that of coverage, since the source of social information is usually based on people choosing to provide their opinions. This normally happens either when the person has a strong opinion about the item in

question (such as a movie they thought was really good or really bad), or when the item is generally of significant interest (a movie that everyone is talking about). Part of the difficulty is to motivate people to give more feedback on more mundane items, or items that may be important to different people under different circumstances. For example, some of the reaction to Bing’s inclusion of Facebook “like” results is that “Searching for the same thing with exact keywords for what a friend has liked is kind of rare” [3].

2 Prediction extraction: getting opinions about opinions

The limitations of the friends’ opinion and usage patterns paradigms stem from the underlying assumptions that “your friends are like you” and “people who agree on certain things also agree on others”. We propose a new kind of social signal that is aimed at addressing these weaknesses and improving the quality of social search. It is based on the observation that, although your friend’s are not you and may not have the same tastes and opinions as you, they are actually the people who *know* you best. Moreover, not only do they collectively carry this wealth of information about your personality and tastes, but they are also intelligent (strong-AI-equipped) computational entities that can process this information to make accurate predictions about your taste, as compared to current machine learning techniques. The approach we propose is to elicit *predictions* about the target user’s opinion of a certain item from the user’s friends who have experienced the item, and aggregate these predictions to construct an estimation of the target user’s opinion of the item *before he has experienced it*.

We first give a simple outline of the way in which the process will work, and shall then describe stronger additional features. We propose a service, probably in the form of a Facebook application, called OpinionSquare, which will allow people to give ratings to any items or topics that they have experienced, such as movies, news events, articles, places, fashion, food, games, Youtube videos, images, famous personalities, etc. For any such item that a Facebook user has experienced, he can give a rating (e.g. a score out of 10) reflecting his own opinion, and he can also give a *prediction* for some (or all) of his friends about how he thinks the friend will rate the item in question. These friends may or may not have experienced the item yet, and may not even be aware of the predictions that are being given by their friends, that is, the process is asynchronous.

For each user and a given item, a weighted average of the ratings predicted for him by his friends is maintained, where the weights are initially uniform over all friends. This measure, based on the collective intelligence of the user’s friends and their knowledge of him, is taken as an estimation of the target user’s rating of the item. It will be used to improve the target user’s search results, which could be in the form of ranking results, annotations showing friend’s predicted ratings, or could be made more directly available as new events on OpinionSquare.

When this target user eventually experiences the item at some later time, he can give his *actual* rating of the item on OpinionSquare. At this point, the accuracy of each of his friend’s predictions will be measured to determine how close their prediction was to the actual rating. The measured accuracy will be used for two purposes. The first is to update the friend’s weight in future calculations of the weighted average of predictions. This will improve the accuracy of the estimation, as well as work against malicious use such as friends purposely making inaccurate predictions (such predictions will get diminishing weights in future estimations). The second use of the accuracy measure will be to calculate various kinds of *scores* reflecting a user’s prediction ability, such as

- an overall average score indicating the user’s prediction ability in general
- a score with respect to each friend indicating “how well the user knows this friend”
- a score with respect to different categories averaged over friends, indicating the user’s prediction ability in specialised areas, e.g. entertainment, sports, fashion, world events, science, etc

Such scores which may be displayed in a user’s profile, in leaderboards, or shared between friends, are aimed at bringing a gamification aspect to the process.

3 Motivation

An important part of the process defined above is the motivation for users to make predictions about their friends' opinions. We discuss some of the aspects of OpinionSquare that are expected to induce such motivation.

Enjoyment. There are a number of examples of successful Facebook applications such as “How well do you know me?” that indicate that people find it entertaining to attempt to guess each other's personalities and tastes. This is expected to become more interesting in this case, since users are allowed to choose any items of interest that they encounter, as they encounter them, e.g. films they have recently seen or articles they have read. The process can be made much simpler and quicker with a browser plug-in that allows people to easily input OpinionSquare predictions for their Facebook friends as they visit various websites, as well as a mobile phone app that makes it easier to enter opinion predictions as places are visited and events encountered.

Curiosity. People are generally interested to know their friends' opinions and responses about things that they encounter, and these usually form topics of discussion in social interactions. OpinionSquare may provide a platform for triggering such discussions or social interactions with respect to current events. It will also be interesting for people to everyday find new estimations of their opinions about things they are not even aware of.

Altruism. People like to help their friends if they come across things which they feel may be relevant to them. OpinionSquare will provide a simple, quick and non-intrusive way of giving personalised recommendations to friends, with the unexpected sense of reward if predictions (recommendations) will be appreciated by their friends at a later time. For example, if a user is buying a car and comes across various options, and makes predictions for his friend who is currently not buying, then the friend may be very thankful for those predictions if he does decide to buy at some future date.

Competition. Finally, there is the sense of competition and boasting, brought about by the scoring feature. People may compete with one another to maximise their scores, and may boast about their prediction abilities. However, a possible drawback is that people may also resort to cheating by informing each other of their actual ratings in order to mutually increase each other's scores. This could be countered by gamifying the activity in different ways, for example in a “one-on-one” manner focusing on scores between pairs of friends to see who knows who better. However, in general it may be better to not stress the competition aspect too much and present the application more as an interesting, enjoyable and helpful social activity.

4 Elaborating opinions

So far we have described the process with a simple format for opinion expression which is a rating out of 10. We shall now discuss two stronger features of the approach that will lead to more powerful information extraction and search improvement. The first of these is to include elaboration on opinions. As well as giving a predicted score out of 10 for each of his friends, the user may also provide additional tags expressing the specific response that he expects the friend to have about the item in question. These could be in the form of descriptive words or short phrases, e.g. “dreadful”, “nice”, “fascinating”, “boring”, “stylish”, “cringe-worthy”, “silly”, “cliche”, “nostalgic”, “scary”, “funny”, “shocking”, etc. The accuracy measure and the reward factor (apart from the standard motivational aspects) may be quantified as a score based on similarity of the predicted and actual responses given by the target user, perhaps based on syntactic and semantic similarities of words using resources such as Wordnet. Higher reward could be assigned for accurate predictions that are less common amongst friends, or those that are less general.

Usefulness. A straightforward application of this feature will be to include the predicted response tags as annotations in the target user's search results, to give him more insight as to how his friends think he may react to the item. However, a more interesting reverse-application is personalised *response-based* search. A standard search process is item-based, that is, the user has some idea about the kind of items he is searching for, which he tries to discover by entering keywords that describe characteristics of those items. In a response-based search,

the search process would be driven by the personal responses that a user desires to be evoked in him. He could enter the “responses” in a search query to get results about items that are likely to evoke those responses in him, e.g. he may be looking for an article that is “intriguing” for him, a game that is “difficult” for him, a movie he will find “scary”, a YouTube video he will find “funny”, or an image or song that may be “nostalgic” for him, etc. Such queries are highly subjective in nature and the answers will be different for different people, but could be accurately predicted by the people who know the user.

5 Explicating opinions

As well as specifying the kinds of responses that will be expected from friends, users may also supply the reasons for why they expect those responses. These could again be entered in the form of single words or short phrases. Taking the example from [2], a user may predict a very high rating for the movie “Inception” for a certain friend, and give the reason as “Leonardo Dicaprio” if he knows that his friend really likes this actor. The quantified reward factor and accuracy measure could again be based on similarities between the predicted reasons and actual reasons that the target user eventually provides, with higher points to accurately predicted reasons that are less common among friends.

Usefulness. The main advantage of this feature will be the extraction of highly user-relevant information about the item in question. In standard tagging-based approaches, items are usually tagged with keywords that describe their essential characteristic properties, which helps people to search for them or to get a summary of their important characteristics. However, this does not usually describe the properties of the item that may be personally relevant to the user. For such user-centric information,

- the user usually has to manually perform research on the item, such as reading reviews or descriptions.
- the information may not even explicitly exist anywhere on the web, as such properties may not be important characteristics of the item that are relevant for people in general.
- the information may be missed by the user, as he may not think of searching for a particular property that the item may have that would interest him, and therefore would not be able to discover the relevance of the item for him.

The *reason tags* supplied by friends will provide keywords about the item in relation to the target user, that is, they will extract those properties of the item that are most personally or idiosyncratically relevant for the target user. These may be properties of the item that are not significant in general, or worth mentioning in a review, or properties that the target user may not even think of searching for in the item. For example, the target user may have an interest in any movies shot in Paris. The movie “Inception” has many scenes shot in Paris, but this may not be stated explicitly in any of the websites about the movie, and is not something that the user would think of searching for when deciding whether or not he should see the movie. But his friend, having seen the film, can predict “Paris” as one of the reasons for the person to be interested in the movie.

6 Advantages and comparison

We now discuss some of the advantages we expect to gain from the prediction extraction approach, especially in comparison with existing approaches.

- **Accuracy**
 - **Quality of information.** A user’s friends, especially family, childhood friends and colleagues would be well aware of the user’s personality, including tastes, idiosyncrasies, views, and circumstances. Most of this information would not even be available in an explicit form on the web, much less in a form that can be used by search methods. Considering the example from [2], friends and family members may know that a user likes Leonardo Dicaprio but hates movies with guns, which may be the relevant information needed to help the user determine if he should see the film “Inception”.

- **Quality of prediction.** Apart from having in depth knowledge about the user, the other issue is how best to process this information to anticipate the user’s response to something. Even if information such as the user “likes Leonardo Dicaprio”, but “hates movies with guns” is available, these different factors all need to be combined with various degrees of importance and context to form an overall assessment of whether the user will like the movie. Such intuitive judgment-making is a kind of AI-complete problem to which human intelligence is much better suited than existing machine learning techniques.

- **Coverage**

- **More items.** Users generally do not give opinions on items that are not very popular or those they do not have a strong opinion of, but all of the motivational aspects of OpinionSquare described above may yield more predictions about more items that are encountered by users every day.
- **Discovery vs. Search.** Encouraging users to always keep their friends in mind with respect to things they encounter is expected to strengthen the discovery aspect, that is, people learning more about things that they are not searching for but may find interesting and relevant. This will especially be amplified by the elaboration and explication features.
- **Latent and volatile knowledge.** The approach may also provide extraction of more latent and volatile knowledge that people carry about their friends but does not become explicit or useful due to circumstantial reasons, and can be easily lost. For example, if a person experiences a restaurant in a foreign country that a particular friend may really like, then he may not consider recommending it since it is not currently relevant for his friend, or he may not even think of the friend in that context. Such knowledge may be volatile (easily forgotten), so at a time when it becomes relevant (if the friend visits the country), it may no longer be accessible or an opportunity for its communication may not even arise. The asynchronous nature of OpinionSquare will mean that such knowledge can be inputted as it is detected, and be used at a later time when it becomes relevant.
- **Combined extended coverage.** The prediction extraction technique of OpinionSquare may be combined with the existing paradigms to further improve coverage. For example, if a user A trusts the opinion of a certain friend B, but B has not yet experienced or commented on the item in question, then A can still be informed of B’s *predicted opinion* about the item, as judged by B’s friends (who may or may not be A’s friends). Similarly, the predicted opinion information can be used to improve the accuracy of extrapolation methods that use statistical or machine learning techniques to account for missing information.

- **New methods.** Apart from improvements in the quality of socially-extracted information, we have also described the possibilities for new kinds of search methods, such as *response-based search* using elaborated opinion predictions, and *user-specific keywords* based on explicated predictions.

7 Conclusion

We have proposed a new dimension to social search, based on the idea of prediction extraction. This extends the current approaches of “your friends are like you” and “people who agree on certain things also agree on others” with the new dimension of “your friends know you and can best predict your taste”. We have discussed the potential of this approach for improving the quality of the social information available to current methods, as well as introducing new methods for social search. Hopefully, the proposed application will also be an enjoyable and helpful social activity that will encourage people to keep their friends and their feelings in mind in everything that they encounter. We end by noting the seemingly impossible challenge presented in [1]:

“I dont think Bing is going to discover my dislike for artichokes anytime soon. It would be freaky though. If I searched for italian appetizers and it warned me not to bother with antipasto salads because they contain artichoke hearts”.

As we have discussed, something like this may in fact be conceivable with the approach proposed here: if your friend searches for salads or tries a new antipasto salad at any time, they can give a low rating on your behalf about it and the reason as “artichokes”, possibly leading to the desired outcome the next time you search for italian appetizers.

References

- [1] Cynthia Boris. “Bing Partners with Facebook for Personalized Search”. *The Pilgrim Network*, 2010. <http://www.marketingpilgrim.com/2010/10/bing-partners-with-facebook-for-personalized-search.html>
- [2] Satya Nadella. “New Signals in Search: The Bing Social Layer”. *Bing Search Blog*, 2010. <http://www.bing.com/community/site.blogs/b/search/archive/2010/10/13/new-signals-in-search-the-bing-social-layer.aspx>.
- [3] P. Yiu. “Bing’s new social features arrive today”. *Bing Search Blog*, 2010. <http://www.bing.com/community/site.blogs/b/search/archive/2010/11/02/search-blog-bing-s-new-social-search-features-arrive-today.aspx?wa=wsignin1.0>.
- [4] Maureen Heymans. “Introducing Google Social Search: I finally found my friend’s New York blog!”. *The Official Google Blog*, 2009. <http://googleblog.blogspot.com/2009/10/introducing-google-social-search-i.html>.
- [5] Doug Sherrets. “Google’s Marissa Mayer: Social search is the future”. *VentureBeat*, 2008. <http://venturebeat.com/2008/01/31/googles-marissa-mayer-social-search-is-the-future>.
- [6] Rob Spiro. “+1’s: the right recommendations right when you want them in your search results”. *The Official Google Blog*, 2011. <http://googleblog.blogspot.com/2011/03/1s-right-recommendations-right-when-you.html>.
- [7] Ed H. Chi. “Information Seeking Can Be Social”. *Computer*, Volume 42 Issue 3. *IEEE Computer Society Press*, 2009.
- [8] Michael McDonnell, Ali Shiri. “Social search: A taxonomy of, and a user-centred approach to, social web search”. *Program: electronic library and information systems*, Volume 45 Issue 1. *Emerald Group Publishing*, 2011.
- [9] Digg. <http://www.digg.com>.
- [10] Delicious. <http://www.delicious.com>.
- [11] StumbleUpon. <http://www.stumble.com>.
- [12] ProjectEmporia. <http://www.projectemporia.com>.
- [13] Amazon Recommendations. <http://www.amazon.com/gp/help/customer/display.html?ie=UTF8&nodeId=13316081>
- [14] P. Heymann, G. Koutrika and H. Garcia-Molina. “Can social bookmarking improve web search?”. *Proceedings of the International Conference on Web Search and Web Data Mining*. *ACM*, 2008.
- [15] X. Su and T. M. Khoshgoftaar. “A survey of collaborative filtering techniques”. *Advances in Artificial Intelligence*. *Hindawi Publishing Corp*, 2009.