

Inspection Mechanisms for Community-based Content Discovery in Microblogs

Nava Tintarev
University of Aberdeen
Aberdeen, UK
nava.tintarev@gmail.com

Byungkyu Kang
Dept. of Computer Science
University of California
Santa Barbara, USA
bkang@cs.ucsb.edu

Tobias Höllerer
Dept. of Computer Science
University of California
Santa Barbara, USA
holl@cs.ucsb.edu

John O'Donovan
Dept. of Computer Science
University of California
Santa Barbara, USA
jod@cs.ucsb.edu

ABSTRACT

This paper presents a formative evaluation of an interface for inspecting microblog content. This novel interface introduces filters by communities, and network structure, as well as ranking of tweets. It aims to improving content discovery, while maintaining content relevance and sense of user control. Participants in the US and the UK interacted with the interface in semi-structured interviews. In two iterations of the same study ($n=4$, $n=8$), we found that the interface gave users a sense of control. Users asked for an active selection of communities, and a more fine-grained functionality for saving individual ‘favorite’ users. Users also highlighted unanticipated uses of the interface such as iteratively discovering new communities to follow, and organizing events. Informed by these studies, we propose improvements and a mock-up for an interface to be used for future larger scale experiments for exploring microblog content.

Author Keywords

Microblogs, visualization, communities, explanations, interfaces, content discovery

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION

Filtering of streaming data such as microblog content is inevitable, even if it is done by showing the most recent content as restricted by screen-size. However our live timelines do often get tailored to us, without transparency or a sense of

control. Getting the selection of the content right is a delicate matter.

Recommender systems address the challenges of finding ‘hidden gems’ which are tailored to individuals from a very wide selection. Implemented well, they hold the key to helping users discover items that are *both* unexpected and relevant, while helping catalog holders sell a wider range of items [3].

In trying to help users make such discoveries, recommender systems walk a thin line between a) making unexpected but risky recommendations (increasing the chances of irrelevant recommendations), and on the other hand b) over-tailoring (resulting in unsurprising recommendations). Over-tailoring can also result in filter bubbles [15], whereby users do not get exposed to items outside their existing interests. For current events, such as content in microblogs, personalization algorithms may narrow what we know, and surround us with information that supports what we already believe. This can result in polarization of views, especially as we have a tendency to self-filter [2].

This paper address these issues by supporting controlled filtering of microblog content. It introduces a novel visualization which supports filtering by allowing a user to control: a) which communities influence their feed b) the network structure relating to these communities and c) different ways of ranking tweets. This visualization is evaluated in two iterations of a qualitative study that assesses the value of such controls, as well as the concrete implementation choices applied. We also discuss the ways these filters and controls are perceived by users, and how they envision that they would use them. We conclude with describing our next steps.

BACKGROUND

Inspectability and Control in Recommender Systems

In the domain of recommender systems there is a growing acceptance and interest in user-centered evaluations [12]. For example, [9] argues for a framework that takes a user-centric approach to recommender system evaluation, beyond the scope of recommendation accuracy. Along the same vein,

Paste the appropriate copyright statement here. ACM now supports three different copyright statements:

- ACM copyright: ACM holds the copyright on the work. This is the historical approach.
- License: The author(s) retain copyright, but ACM receives an exclusive publication license.
- Open Access: The author(s) wish to pay for the work to be open access. The additional fee must be paid to ACM.

This text field is large enough to hold the appropriate release statement assuming it is single spaced.

it has also been recognized that many recommender systems function as *black boxes*, providing no transparency into the working of the recommendation process, nor offering any additional information to accompany the recommendations beyond the recommendations themselves [6].

To address this issue, explanations can be given to improve the transparency and control of recommender systems. Research on textual explanations in recommender systems to date has been evaluated in wide range of domains (varying from movies [19] to financial advice [4]). Increasingly, there has also been a blurring between recommendation and search, making use of information visualization. For example, [20] has looked at how interaction visualization can be used to improve the effectiveness and probability of item selection when users are able to explore and interrelate multiple entities – i.e. items bookmarked by users, recommendations and tags. Similarly [17] found that in addition to receiving transparent and accurate item recommendations, users gained information about their peers, and about the underlying algorithm through interaction with a network visualization.

Inspectability and Control in Microblogs

In order to better deal with the vast amounts of user-generated content in microblogs, a number of recommender systems researchers have studied user experiences through systems that provide transparency of and control over recommendation algorithms. Due to the brevity of microblog messages, many systems provide summary of events or trending topics with detailed explanations [11]. This unique aspect of microblogs makes both inspectability and control of recommender algorithms particularly important, since they help users to more efficiently and effectively deal with fine-grained data. For example, experimental evidence to argue that inspectability and control improve recommendation systems is presented for microblogs in [17], via a commuter traffic analysis experiment, and more generally in [8] using music preference data in their *TasteWeights* system.

Community-based Content Discovery

Serendipity is defined as the act of unexpectedly encountering something fortunate. In the domain of recommender systems, one definition has been the extent to which recommended items are both useful and surprising to a user [7]. This paper investigates how exploration can be supported in a way that improves serendipity.

The intuitions guiding the studies in this paper are based on findings in the area of social recommendations, that is based on people’s relationships in online social networks (e.g., [13, 16]) in addition to more classical recommendation algorithms.

The *first intuition* is that weak rather than strong ties are important for content discovery. This intuition is informed by the findings of the cohesive power of weak ties in social networks, and that some information producers are more influential than others in terms of bridging communities and content [5]. Results in the area of social-based explanations also suggest that mentioning which friend(s) influence a recommendation can be beneficial (e.g., [18, 21]). In this case, we

support exploring immediate connections or friends, as well as friends-of-friends.

The *second intuition* is that the intersection of groups may be particularly fortuitous for the discovery of new content. This is informed by exploitation of cross-domain model inspiration as a means for serendipitous recommendations, e.g., [1].

VISUALIZATION

In this study, we designed a web-based visualization that allows users to experience the recommender system we propose (see Figure 1). The first two columns represent “groups” (communities) and “people” (users), allow us to filter ‘tweets’ in the third column by both of these ‘facets’. The system supports therefore support a faceted navigation, with the third column representing the resulting information. In addition, the system supports Pivoting (or set-oriented browsing), in that it allows users to navigate the search space by starting from a set of instances (by selecting which groups they would like to follow).

The rationale for the visualization follows several intuitions with regards to exploring novel and relevant content in social network, as outlined in the section in related work.

The first is that people can find relevant content in the intersection between multiple communities. In the visualization this is represented by the selection of up to three communities to which a user belongs, and color blending to indicate people and content that represents this type of overlap. Another intuition is that weak ties, or friends of friends, are also good candidates for content discovery. In this visualization they are represented as two hops in a network structure. Consequently we included a slider which included 0-hops (do not consider this community), 1-hop (include people who follow a given community), 2-hops (include people who follow people in a given community).

Finally, the ranking of tweets according to a) relevance to a user compared to b) popularity and c) time is also likely to help users find relevant and unexpected content compared to tweets only ordered by time.

Structure and Interaction

Figure 3 shows a snapshot of the interactive visualization used in the study. Information is presented in three columns. From left to right, these are: group/community, people and tweet columns. Users can interact with entities in any of these three columns to highlight associations to entities in other columns. In the people and tweet columns, entities are clustered and colored based on community associations. In the first column, we visualize a set of communities (also referred to as groups), which by design, may have some membership and content overlapping. Within this column, each entity has a widget to control network distance from that entity. This enables the user to specify how that entity contributes users and content to the other columns. In particular, sliders were used for control in Study 1 and radio buttons in Study 2.

In the second column, a ranked list of users related to each community is visualized. These users serve as sources for information recommended in the third column, but the visualiza-

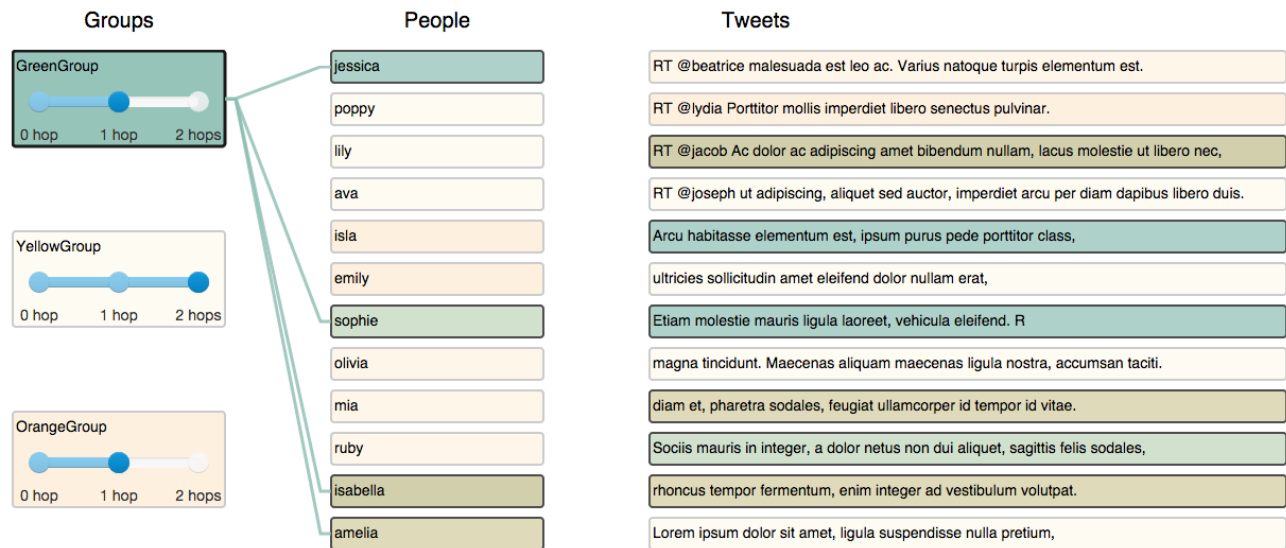


Figure 1. Visualization of the recommendation system used in the study 1.

tion also supports analysis of the connectivity of these users across communities in addition to the content they distribute.

The third column shows the recommended tweets which are by default filtered and ordered according to recency. A user can change the ranking algorithm for this column to either popularity or relevance.

Materials

The materials for the experiment were abstracted: people were given random names of both genders, tweets were short lines from a short Latin text (“Lorem Ipsum. . .”), resulting in a total of 229 tweets. When participants interacted with the system, a random subset of 12 tweets was presented. The top 4 of these tweets included a retweet, to visually increase the similarity with a twitter feed, and was applied consistently across adaptations.

STUDY 1

This section describes a formative study conducted to evaluate the proposed visualization. We used a layered evaluation approach [14], focusing on the decision of an adaptation and how it was applied (in contrast to which data was collected or how it was analyzed). Participants took part in semi-structured interviews, in order to evaluate the user experience (following the guiding scenarios of [10]). More concretely this study aimed to answer the following questions: a) are the three introduced controls (selection of communities, network structure, and ranking of tweets) considered useful for participants? b) is the way they are implemented useful? c) do these controls give users a sense of control? d) do participants use the controls in the way that we envisaged? The version of the system used for this study can be found online¹.

Participants

¹<http://goo.gl/krOvuJ>

Time Popularity Relevance

Tweets

- RT @beatrice malesuada est leo ac. Varius natoque turpis elementum est.
- RT @lydia Porttitor mollis imperdiet libero senectus pulvinar.
- RT @jacob Ac dolor ac adipiscing amet bibendum nullam, lacus molestie ut libero nec,
- RT @joseph ut adipiscing, aliquet sed auctor, imperdiet arcu per diam dapibus libero duis.
- Arcu habitasse elementum est, ipsum purus pede porttitor class,
- ultrices sollicitudin amet eleifend dolor nullam erat,
- Etiam molestie mauris ligula laoreet, vehicula eleifend. R
- magna tincidunt. Maecenas aliquam maecenas ligula nostra, accumsan taciti.
- diam et, pharetra sodales, feugiat ullamcorper id tempor id vitae.
- Sociis mauris in integer, a dolor netus non dui aliquet, sagittis felis sodales,
- rhoncus tempor fermentum, enim integer ad vestibulum volutpat.
- Lorem ipsum dolor sit amet, ligula suspendisse nulla pretium,

4 participants were recruited from research staff at computer science department at a UK university. Their ages ranged from 23-51. They all had twitter accounts, but their experience with twitter ranged from inactive to highly experienced (including the use of twitter management and analytics applications). 1 was female, and 3 male. They all had a native or fluent level of English language skills. Participants varied from PhD students, post-doctoral fellows to teaching staff.

One of the participants had done research with visualizations and twitter, the other three had no experience with either. None knew Latin (one had taken Latin course, but professed a very rudimentary level of knowledge).

Procedure

Participants took part in individual semi-structured interviews, following a user test plan². Following the collection of basic demographic data, participants were given a brief introduction to the system. The various interface components were verbally introduced without interacting with the system. Participants were then given several simple tasks such as including people who are connected to other people for a given community, or ranking tweets by relevance (rather than time). Following each interaction participants were asked how the tweets had changed, if new ones had been added, or if tweets had disappeared. The tasks given were:

- Go to the system online. What are your first impressions?
- Select one of three communities that you are a member of and reflect your interests (if user can not think of any tell them to think of conferences that they attend). Have a look at the tweets that are recommended to you.
- Add tweets (1 hop) for a second community of your choice from the above.

²<https://goo.gl/3KpH9z>

- Is there any relevant tweet from this second community you did not see before? Are there any that have disappeared?
- The tweets are currently ranked by time, change this to rank the tweets by popularity.
- Are there now any tweets you did not see before? Are there any that have disappeared?
- Now, change who you get your tweets from to include people who are linked to (2 hops) people that attend your first community. You may want to remove the second community for this too.
- How about now, are there now any tweets you did not see before? Are there any that have disappeared?

Following the interaction with the system, participants took part in an exit interview where they were asked about their perceived control of the system, the usefulness of various functionalities, and how they would use them for exploration. More concretely the questions asked included

- How did it feel? What was your impression? (Positive impressions? Negative impressions?)
- Would you have liked more training on how to interact with the visualization before you got started?
- How helpful did you find the following functionalities (1-7, unhelpful to helpful), and how could they be improved?
 - Tweets organized by community;
 - Changing how the tweets are ordered/ranked
 - Changing who I get tweets from (0,1,2 hops)
 - Being able to interact with the system to specify different preferences
 - The links between different parts of the interface (people, groups, tweets).
- Do you think these functionalities would help you find new and relevant information you would not find otherwise? How would you use them to do this?
- Does the filtering give you a sense of what you might be missing, or does it hide information that you need?
- Did you feel like you had control over which information was presented to you?
- Would you liked to have had any controls that are not present in this interface?

Results

Are the introduced controls (communities, network structure, and ranking of tweets) considered useful by participants?

The scores given to the various controls was generally high (5 or above). There were three exceptions. Participant3 did not find tweet ranking by relevance and popularity useful at all. Participant4 gave low scores to the hop control for network structure, and the links, but this was due to the way they were implemented, and is discussed below.

Is the way they are implemented useful?

All the participants noted that the interface was simple and clean, and had a good first impression. Participant4 noted that it would be well suited for a mobile interface.

- **Hop control** All of the participants found it difficult to understand the control for the network structure. When thinking aloud, several said that pulling the slider further to the right would increase the number of tweets on a certain topic, rather than widen the network (which potentially would dilute the focus of the tweets).
- **Community selection** Participant1 wanted to ‘activate’ a community by selecting its box. This seems more intuitive than selecting 0 hops for the communities they did not want to follow.
- **People** In addition to filtering on community structure and inclusion, several participants wanted a finer grain control of which users were included in the selection of tweets. Some users wanted to activate users somehow, by either adding them to favorites at the top of the person list, or activating through selection. These participants felt that this should influence the ranking of tweets.
- **Tweets** Participants felt that tweets belonging to the same community should not only have the same color, but be grouped together. Participant3 (experienced twitter user) felt that ranking of tweets by any other measure than recency (time) was not useful.
- **Links** Participant3 found the links and colors between the columns inconsistent. The relationship between the first two columns used links, whereas the relationship between the second two columns used colors.
- **Color-interleaving** Participant1 mistook the color-interleaving to imply significance, as they varied in hue. However, the other participants interpreted this correctly although did ask if the interpretation was correct.

Do these controls give users a sense of control?

All of the participants felt that the interface improved their control over their tweets. They also consistently agreed that they would be missing some content, and that they were not in complete control, but that they were happy with the balance in the trade-off.

However, Participant3 felt that they wanted to be able to scroll through all of their tweets, especially because they did not have the finer grained control of which individuals appeared in their feed.

Do participants use the controls in the way that we envisaged?

All of the participants completed the simple tasks given to them. They all stated that they would find new and relevant content using the interface, although the highly experience twitter user felt they already find novel content using tools such as TweetDeck. When asked how they would you use the functionalities to find new and relevant information, participants suggested two uses we had not initially considered:

Organizing events Participant3 felt that the groups could be defined by other characteristics rather than membership of

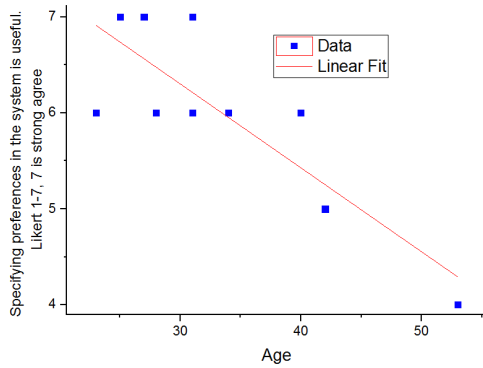


Figure 2. Plot showing correlation between participant age and reported importance of “Being able to interact with the system to specify different preferences”.

a community, such as geographic location. This participant suggested that they would use this functionality to identify and coordinate groups of people when organizing events on the topics they were interested in.

Discover new groups Participant2 was confident that they would find new relevant communities when looking at the intersection of existing communities that they follow. This participant listed three music bands that they listen to and would follow on twitter. They would use the system to discover new bands, and would then add them as a new group as a “seed” for further discovery.

Other suggestions

Participants suggested several features they would expect in an interface that was integrated with twitter. For example, they would want to be able to view the profiles (or at least, the first 50 characters) of the people they are receiving tweets from. Others wanted to be able to reply to tweets directly from the feed.

Another suggestion was to introduce separate columns for different communities. This may be related to the request by other users to be able to group tweets by community.

STUDY 2

The first study identified several limitations of the system, which were addressed for a second iteration of evaluation. Improvements included: a) using buttons rather than a slider to control the number of hops; b) sorting people by group affinity, e.g. greenGroup people were listed at the top, rather than mixed throughout the list; c) identifying how many people were filtered (i.e. “Showing 12 of 1307”). The improved interface can be seen in Figure 3, with annotations to highlight each improvement. The version of the system used for this study can be found online³.

Participants

8 participants were recruited from research staff at computer science department at a US university. Their ages ranged from 20 to 45. 5 participants were female and 3 were male. Participants varied from PhD students, post-doctoral fellows

³<http://penguinkang.com/intRS/>

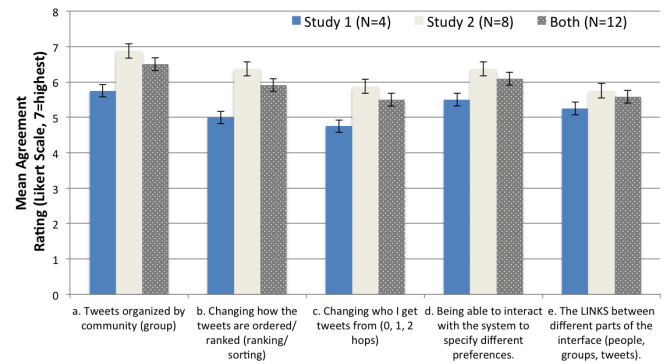


Figure 4. Analysis of subjective results in exit interviews for the two studies. Error bars show standard error.

to teaching staff in computer science, engineering, media-arts and physics. They all had a native or fluent level of English language skills. 6 of the participants had Twitter accounts, and one person had done research with Twitter data in the past. 5 had done research with visualization. As with Study1, no participants knew Latin.

Procedure

As in Study1, participants took part in individual semi-structured interviews. Studies were conducted in a computer science lab on campus using two notebook computers. The participant interacted with the UI on one, and the experimenter/interviewer took notes on the other. On average, studies lasted 35 minutes (min 28 minutes, max 43 minutes).

Results

In this section, we revisit questions from Study1 and add additional comments and discussion based on the new participants interacting with the improved UI in Study2. Figure 4 shows a comparison of participants’ opinions on the different features of the system between Study1 (N=4) and Study2 (N=8), along with the combined score (N=12). We note that the combined score is based on two slightly different UI designs, and it is only used as a rough estimate of the overall group evaluation.

Are the introduced controls (communities, network structure, and ranking of tweets) considered useful by participants?

The scores shown in Figure 4 range between 5.58 and 6.87 for Study2, shown in the middle column of each group, an average of approximately one point on the 7 point scale. Compared to Study1, the interface modifications appear to have had a positive impact on user experience with the system. While this is a promising side result, the purpose of the study was to provide a formative evaluation of the interface.

Participants reported the best score for the feature to organize Tweets by community, which is a core contribution of the system. This is encouraging feedback as the authors are designing a larger-scale quantitative evaluation with this as a central feature. The features that elicited the lowest scores were the hop-distance selector and the edge visualizations between the columns.

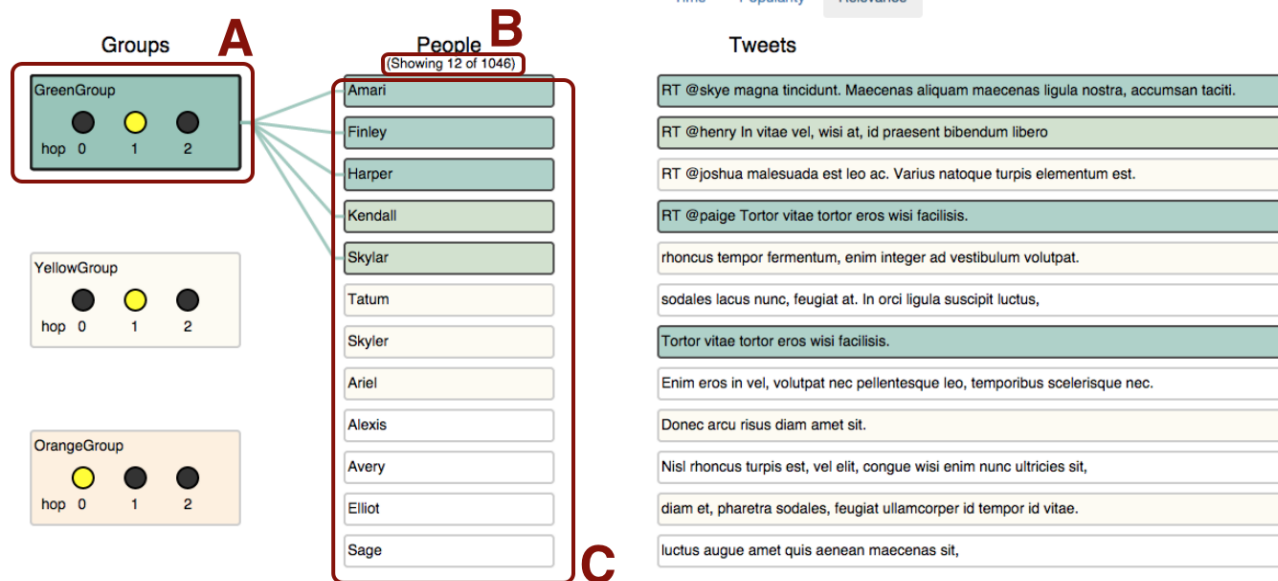


Figure 3. Improved visualization design used in the main user study. Annotation (A) shows changes to the number-of-hops selection. (B) shows the number of filtered users interactively in the form “m of n”, and (C) shows connectivity-based clustering and associated coloring of nodes in the “People” column.

Participants also reported that they liked the ability to change how Tweets were ordered and ranked through the interface. One participants commented that “I can’t do this in Facebook or Twitter – this is great!”. Support for expressing real-time preferences through interactive interface components met with strong positive feedback, with all users reporting a sense of increased control over the information feed.

Is the way they are implemented useful?

Similarly to Study1 study, all participants commented that the interface was clean and well organized. One participant complained that it was too complex and could benefit from having less data. 50% of the participants pointed out an issue with the node-coloring in column 2, shown in Figure 3. Note that this figure needs to be viewed in color to see the true effect (see link to system above).

- **Hop control** Some participants did not realize that the 0 position essentially turned the group node off. There were also multiple comments that when hop control was set to 0, showing the nodes opaquely was not a good design choice. One participant explicitly mentioned that it would be better to remove these nodes completely, noting that the visual effect of setting the hop-control to 0 would be much shorter. Unlike Study1, no participants confused the hop slider with a weighting mechanism, and all understood that it sourced users from n-hops farther away in the Twitter network.
- **Community selection** Most participants commented that community selection and analysis was a strong point of the system. Suggested communities included musical artists, pet fan clubs, and conferences or meetings.
- **People** A few participants reported having trouble understanding the coloring and community-based group-

ing/clustering in this column. All participants understood the data flow correctly by the end of the sessions, but this feature took longer than others for them to master. The main cited reason for this was that the colors – added to distinguish the groups, were too similar, as mentioned above. Two participants mentioned that it would be useful to select or weight people of interest.

- **Tweets** Two participants suggested that a ranking score would be useful to distinguish between tweets in the right column. Participants also requested that when a change is made in the system, the source of that change’s effect on the list should be visualized. Our proposed solution to this is shown in Figure 5 as a ranking source indicator for each tweet.
- **Links** Participants were slightly dissatisfied with how links were shown in the system. Three people commented that links should be shown across all columns when a particular group is selected in the left column, or when any other node is selected, to visually communicate the associations of that node. Other participants commented that the on-demand design was a good idea to avoid cluttering the view.
- **Color-interleaving** Half of the users complained that this was too subtle and needed to be made more explicit. This has been addressed through the use of colored icons next to people to signal group memberships. The color palette has also been changed to make clearer distinction between groups.

Do these features give users a sense of control?

In keeping with Study1, all of the participants felt that the interface improved their control over their tweets. They also consistently agreed that they would be missing some content,

and that they were not in complete control, but that they were happy with the balance in the trade-off. Similar to the Study1, two participants suggested use of scrolling or similar mechanism to view filtered-out tweets in case they wanted to.

Do participants use the features in the way that we envisaged? Generally, participants reported that they would find the system useful for discovering new content and exploring community structure in the domains that they chose (music, conferences, pet fan clubs etc.). In particular, they felt that real-time preference feedback, community selection and algorithm selection (time, relevance or popularity) gave them a good sense of control. Many commented that such features would be useful on everyday social media streams such as in Twitter and Facebook.

Participants suggested similar uses of the controls as in Study1. Many suggested using the system for organizing events and advertising across relevant communities, and for discovering new groups. Echoing the comments of Study1, one participant mentioned that they would like to use the system for exploring a broader network of musical artists. They described selection of three fan club communities as in our experimental setup, but went on to describe iteratively replacing them with new nodes that were discovered on the right column, thereby applying the interface (theoretically) as a network traversal and discovery tool. This is an example of a reported use that was not in our design.

Another participant proposed to use the system to analyze which community produced the most popular content on Twitter, by using the popularity ranking algorithm and traversing the edge connections back to the groups.

Other suggestions

Participants suggested a variety of ways to improve the interface. These included addition of multimedia content to the tweet column, and visually distinguishing retweets (compared to original tweets) by color. Participants also suggested creating visually distinct colorings for blended color groups, and displaying links to all group memberships upon clicking a user node (rather than upon hover). Another request was for an indication of how much data has been filtered in all the columns (currently only for the people column).

Participants also suggested measuring the usefulness of the system for getting an overview of a new community or topic.

Demographics Analysis

A brief analysis of demographics and responses showed an interesting correlation between participant age and the perceived importance of specifying preferences on-the-fly in the user interface. Figure 2 shows a plot with the Likert-scale responses for the dynamic preferences shown on the Y-axis and participant age shown on the X-axis. The data follows a negative linear trend, with younger participants specifying a higher perceived importance of specifying preferences.

CONCLUSION AND FUTURE WORK

In this paper we evaluated a visualization which allowed users to explore and filter microblog content for communities to

which they belong. The ability to organize Tweets by community, the core contribution of the visualization, was rated the most highly. Users also stated that the interface gave them enough control over their content, even if they felt some information would inevitably be hidden – the trade-off was considered acceptable.

We also found several unexpected uses of the system. For example two separate participants, in different experimental settings (one in the UK and one in the US) applied the interface (theoretically) as a network traversal and discovery tool for music.

Figure 5 introduces an improved mock-up with a number of changes. In addition to these improvements, we are planning larger-scale quantitative evaluations. One of these will explore the use of community-based filters, and the other controls introduced in this paper, on existing twitter feeds.

ACKNOWLEDGMENTS

This research has been carried out within the project Scrutable Autonomous Systems (SAsSY), funded by the UK Engineering and Physical Sciences Research Council, grant ref. EP/J012084/1.

This work was also partially supported by the U.S. Army Research Laboratory under Cooperative Agreement No. W911NF-09-2-0053; The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of ARL, NSF, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

1. André, P., m.c. schraefel, Teevan, J., and Dumais, S. T. Discovery is never by chance: Designing for (un)serendipity. In *CC* (2009).
2. Bakshy, E., Messing, S., and Adamic, L. A. Exposure to ideologically diverse news and opinion on facebook. *Science* 348 (2015), 1130–1132.
3. Castells, P., Hurley, N., and Vargas, S. *Recommender Systems Handbook (second ed)*. (in press), ch. Novelty and Diversity in Recommender Systems.
4. Felfernig, A., Teppan, E., and Gula, B. Knowledge-based recommender technologies for marketing and sales. *Int. J. Patt. Recogn. Artif. Intell.* 21 (2007), 333–355.
5. Granovetter, M. S. The Strength of Weak Ties. *The American Journal of Sociology* 78, 6 (1973), 1360–1380.
6. Herlocker, J. L., Konstan, J. A., and Riedl, J. Explaining collaborative filtering recommendations. In *ACM conference on Computer supported cooperative work* (2000), 241–250.
7. Herlocker, J. L., Konstan, J. A., Terveen, L., and Riedl, J. T. Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.* 22, 1 (2004), 5–53.
8. Knijnenburg, B. P., Bostandjiev, S., O’Donovan, J., and Kobsa, A. Inspectability and control in social

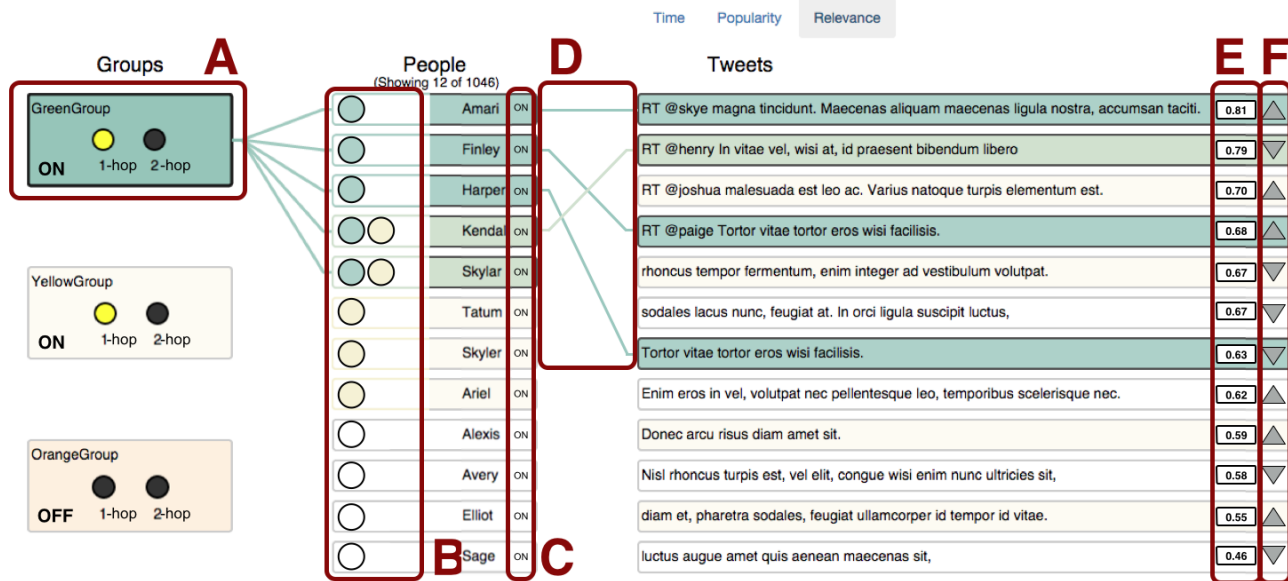


Figure 5. Mock-up of improved UI and interaction design based on study results and analysis: (A) improved representation of the hop-distance controls, (B) iconization to show group memberships, (C) Activation (on/off) control of nodes, (D) visualization of dynamic edges, (E) addition of a ranking score for recommended content, and (F), addition of a provenance arrow to show what the previous interaction did to the ranking of each recommendation.

- recommenders. In *Proceedings of the Sixth ACM Conference on Recommender Systems, RecSys '12*, ACM (New York, NY, USA, 2012), 43–50.
- Knijnenburg, B. P., Willemsen, M. C., Gantner, Z., Soncu, H., and Newell, C. Explaining the user experience of recommender systems. *User Modeling and User-Adapted Interaction* 22, 4-5 (2012), 441–504.
 - Lam, H., Bertini, E., Isenberg, P., Plaisant, C., and Carpendale, S. Empirical studies in information visualization: Seven scenarios. *IEEE Transactions on Visualization and Computer Graphics* 18(9) (2012), 1520–1536.
 - Marcus, A., Bernstein, M. S., Badar, O., Karger, D. R., Madden, S., and Miller, R. C. Twitinfo: Aggregating and visualizing microblogs for event exploration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11*, ACM (New York, NY, USA, 2011), 227–236.
 - McNee, S. M., Riedl, J., and Konstan, J. A. Being accurate is not enough: How accuracy metrics have hurt recommender systems. In *Extended Abstracts of the 2006 ACM Conference on Human Factors in Computing Systems (CHI 2006)* (2006).
 - Nagulendra, S., and Vassileva, J. Providing awareness, understanding and control of personalized stream filtering in a p2p social network. In *Conference on Collaboration and Technology (CRIWG)* (2013).
 - Paramythis, A., Weibelzahl, S., and Masthoff, J. Layered evaluation of interactive adaptive systems: Framework and formative methods. *User Modeling and User-Adapted Interaction* 20 (2010).
 - Pariser, E. *The filter bubble: What the Internet is hiding from you*. Penguin Books, 2011.
 - Pizzato, L., Rej, T., Akehurst, J., Koprinska, I., Yacef, K., and Kay, J. Recommending people to people: the nature of reciprocal recommenders with a case study in online dating. *User Modeling and User-Adapted Interaction* 23 (2013), 447–488.
 - Schaffer, J., Giridhar, P., Jones, D., Höllerer, T., Abdelzaher, T., and O'Donovan, J. Getting the message?: A study of explanation interfaces for microblog data analysis. In *Proceedings of the 20th International Conference on Intelligent User Interfaces, IUI '15*, ACM (New York, NY, USA, 2015), 345–356.
 - Sharma, A., and Cosley, D. Do social explanations work? studying and modeling the effects of social explanations in recommender systems. In *World Wide Web (WWW)* (2013).
 - Tintarev, N., and Masthoff, J. Personalizing movie explanations using commercial meta-data. In *Adaptive Hypermedia* (2008).
 - Verbert, K., Parra, D., Brusilovsky, P., and Duval, E. Visualizing recommendations to support exploration, transparency and controllability. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces, IUI '13*, ACM (New York, NY, USA, 2013), 351–362.
 - Wang, B., Ester, M., Bu, J., and Cai, D. Who also likes it? generating the most persuasive social explanations in recommender systems. In *Twenty-Eighth AAAI Conference on Artificial Intelligence* (2014).