

Data Portraits and Intermediary Topics: Encouraging Exploration of Politically Diverse Profiles

Eduardo Graells-Garrido*
Telefónica I+D
Santiago, Chile

Mounia Lalmas
Yahoo Labs
London, UK

Ricardo Baeza-Yates
Yahoo Labs
Sunnyvale, USA

ABSTRACT

In micro-blogging platforms, people connect and interact with others. However, due to cognitive biases, they tend to interact with like-minded people and read agreeable information only. Many efforts to make people connect with those who think differently have not worked well. In this paper, we hypothesize, first, that previous approaches have not worked because they have been direct – they have tried to explicitly connect people with those having opposing views on sensitive issues. Second, that neither recommendation or presentation of information by themselves are enough to encourage behavioral change. We propose a platform that mixes a recommender algorithm and a visualization-based user interface to explore recommendations. It recommends politically diverse profiles in terms of distance of latent topics, and displays those recommendations in a visual representation of each user’s personal content. We performed an “in the wild” evaluation of this platform, and found that people explored more recommendations when using a biased algorithm instead of ours. In line with our hypothesis, we also found that the mixture of our recommender algorithm and our user interface, allowed politically interested users to exhibit an unbiased exploration of the recommended profiles. Finally, our results contribute insights in two aspects: first, which individual differences are important when designing platforms aimed at behavioral change; and second, which algorithms and user interfaces should be mixed to help users avoid cognitive mechanisms that lead to biased behavior.

ACM Classification Keywords

H.4.3 Information Storage and Retrieval: Information Search and Retrieval—*Information Filtering*; H.5.2 Information Interfaces and Presentation: User Interfaces—*Graphical user interfaces (GUI)*

Author Keywords

Homophily; Selective Exposure; Recommender Systems; Information Visualization.

*Corresponding author: eduardo.graells@telefonica.com. Work primarily carried out while the first author was a PhD student in the Web Research Group, at Universitat Pompeu Fabra, Barcelona, Spain.

INTRODUCTION

Research from social sciences has shown that, while everyone seems to have a voice on the Web, people tend to listen and connect to those with similar beliefs in political and ideological issues. Such behavior can be explained in terms of homophily [44], a cognitive bias. Homophily is present in many situations and can be beneficial, as communication with culturally alike people is easier to handle. However, with respect to ideological issues it can have serious consequences, both off- and on-line. On one hand, groups of like-minded users tend to disconnect from other groups, polarizing group views [58]. On the other hand, Web platforms recommend and adapt content based on interaction and network data of users, *i. e.*, who is connected to them and what they have liked before. Because one of their main goals is to maximize user engagement, recommendation algorithms often push content that reinforces homophily in behavior by displaying mostly agreeable information. Such biased reinforcement, in turn, makes further recommendations consisting of even more polarized content, confining users to *filter bubbles* [48].

Until now, most approaches have focused on how to motivate users to read challenging information or how to motivate a change in behavior through recommender systems and display of potentially challenging information. This direct approach has not been effective as users often do not value diversity or do not feel satisfied with it [46]. Motivated by this scenario, our work aims at understanding how to encourage exposure to diverse people from an ideological point of view on micro-blogging platforms. Our research questions are:

Is it possible to encourage exploration and acceptance of user profiles recommended on the basis of political diversity? If so, which factors influence this behavior?

We propose an indirect approach using an advanced user interface, because we believe that neither a recommender system nor a user interface alone are enough to encourage unbiased behavior. Our hypothesis is that an indirect, mixed approach helps users to overcome the cognitive dissonance produced by exposure to potentially challenging information.

To achieve this mixed approach, our work combines the output of a recommender algorithm with a visual depiction. The recommender algorithm is aimed at recommending people who may think differently based on a proxy latent space modeling of topics, named *intermediary topics* [29]. Intermediary topics are latent topics that users, with distant views on some sensitive issues, have in common (*e. g.*, a music genre, a cuisine type). These topics are our means to introduce users to each other, knowing that first impressions matter [4], thus giving connection a chance to happen.

The visual depiction is aimed, first, at providing a pleasing and joyful view of the user’s own interests, and then to provide a visual display of recommended profiles. Particularly, our proposal is a visualization of user profiles called *data portrait* [20], with the purpose of making users aware of their own interests and the image they project on a social platform. The visualization of recommendations – of people to “follow” – is based on a hierarchical visualization technique displaying how *recommendees* can be grouped. Both concepts allow users to receive recommendations in the context of their data portrait.

Our system was deployed “in the wild” to evaluate its effects on users, using a proof-of-concept implementation of a recommender system based on the intermediary topics paradigm. We focused on Chile, a Latin-American country with one of the highest Internet penetration rates among developing countries [13], and whose population actively discuss politics on Twitter [60]. As found in a previous study, such discussion is homophilic [30], making it a good candidate for analysis. We analyzed how Chilean users interact with their portraits and explore the recommendations within the realm of *Casual Information Visualization* [52]. We did not frame our evaluation using specific tasks to be performed by users, as these would not reflect scenarios for our system. Instead, we analyzed user behavior employing user engagement metrics [41], as these are more appropriate for studying non-goal oriented user behaviors. The main results of our experiment are:

- The usage of visualization to depict recommendations alongside a data portrait allowed users to explore more recommended profiles, regardless of the recommender algorithm used to generate them.
- Behavioral differences in terms of political involvement influenced how users interacted and engaged with our system.
- On a standalone basis, neither recommender system or visualization help users to exhibit a conscious (unbiased) exploration. However, their combination is effective when users have political content in their profiles.

These results have implications on the design of systems aimed at exploring user generated content: a one-size-fits-all approach misses the opportunity of giving users tools to get the best out of their exploring experience. We discuss design implications in terms of who can be targeted with systems like ours and when to consider presenting diverse recommendations. Our main conclusion is that indirect approaches like these data portraits can help users to make conscious decisions in these biased scenarios.

RELATED WORK

Homophily and Content Recommendation

Homophily is the tendency to form ties with similar others, where similarity can be bound to factors ranging from socio-demographic to behavioral and intra-personal ones [44]. In micro-blogging platforms, the presence of homophily in how individuals interact has been shown to be reflected in the structure of their ego-network, which has allowed to predict user attributes [1, 56], including political leaning [7]. Likewise, its presence has been used to recommend people to interact with [14, 35].

In this paper, we focus on recommendations on micro-blogging platforms based on user similarity. Similarity can be defined in different ways. For instance, two users might be similar if they use the same *hashtags* [12], follow the same accounts [26], mention the same entities [45], or have similar *latent topics* [53] as estimated using Latent Dirichlet Allocation (LDA hereafter) [9]. However, similarity is not the only feature to consider when recommending information or people to follow. Other features include content quality and popularity [15], network relevance (*friend of a friend*) [14], explainability [40], and centrality measures [34].

In our work, we propose *intermediary topics* [29] as a feature to consider when recommending users to follow. The intuition behind intermediary topics is that they focus on homophily on the basis of specific shared latent topics computed using LDA [9]. LDA has been found to be reliable for user classification [51], and in our context, it has successfully identified latent topics that could act as intermediary between people with diverse political profiles [30].

Using topics and relationships between them, represented through the so-called *topic graphs*, for recommendation purposes is not new. For instance, Gretarsson *et al.* [33] visualize topic graphs to ease knowledge discovery. Our context is different. We do not attempt to visualize them; instead, we use the intermediary topics as input features for generating the recommendations. We nonetheless experiment with visual depiction of recommendations, through the paradigm of data portraits, because the visualization of social recommendations has been shown to boost user satisfaction [32].

Encouraging Exposure to Diverse Content

Exposure to exclusively agreeable information or like-minded people reinforces and polarizes individual and group stances on ideological issues [47, 59]. Various works have looked into ways to improve the exposure to challenging information, by employing algorithms for content selection, as well as changing the depiction of this type of information. The degrees of success vary, but results have not been as good as expected in terms of behavioral change.

Opinion Space [23] is a self-organizing interactive visualization of an information space. In it, individual opinions of participants in debate were visualized according to their *opinion profiles*, built automatically for each participant after answering questions about key political issues. Although the visual approach did not reduce selective exposure, it generated more engagement than baseline text-based interfaces and users were more respectful with those having opposite opinions. In *NewsCube* [49], several automatically determined aspects of news stories in political contexts are presented to allow users to access diverse points of view of news events in political contexts. Each aspect is displayed in its own cluster, allowing users to see the diversity of available points of view. This clustered presentation augments the number of interactions with news, but not the number of interactions with different, opposing, clusters [16]. Munson *et al.* [46] tested different ways of altering a user interface without changing its core interaction mechanisms, by changing the sorting order of information and highlighting items pertaining to opposite points

of view with respect to the user. It was found that only a minority of users, called *diversity-aware*, values diversity. In discussion forums, Liao *et al.* [42] added *position indicators* of stance polarization to participants, improving agreement of users with those of opposing views when their positions were not consistently moderate, or when the information seekers were looking for highly accurate information.

One possible reason of why direct approaches have not worked is the *selective exposure* [37] mechanism. Selective exposure makes users discard potentially challenging information, either by not reading it, not accepting the recommendation, or preferring an agreeable alternative, regardless of the factual value of the agreeable information item. Through selective exposure, users avoid *cognitive dissonance* [24], a state of discomfort that affects persons confronted with conflicting ideas, beliefs, values or emotional reactions. Thus, if users were to receive recommendations of users who think differently, the selective exposure mechanism would likely prevent them from even browsing the recommendations.

In our work, we apply an indirect approach, where we use intermediary topics [29] to recommend people with potentially distant opinions (as captured by LDA), and yet with specific interests in common. In contrast to the approaches above mentioned, the recommendations are not based on their (non-)alignment on politics or sensitive issues. Instead, we build a data portrait of users of micro-blogging platforms, and show recommendations in that context, emphasizing the similarity of users in terms of their alignment on intermediary topics.

Information Visualization

Our work is related to the field known as *Casual Information Visualization*, defined by Pousman *et al.* [52] as “*the use of computer mediated tools to depict personally meaningful information in visual ways that support everyday users in both everyday work and non-work situations.*” The focus on everyday situations means that there does not need to be a concrete task to be completed, nor a specific analytic insight to be expected.

Yi *et al.* [67] identify four cognitive processes that lead to insights gained through visualization: *provide overview*, *adjust*, *detect pattern*, and *match mental model*. Those relevant to our context are *provide overview* and *match mental model*. To provide an overview of profiles and match mental models of portrayed users, we use *word clouds* as the primary element of our proposed visualization. Word clouds have a long history in information visualization [63]. Although not appropriate for analytical tasks, word clouds are expressive, familiar, and popular with users, as they help them expressing themselves [64]. We use word clouds both to provide an overview of a profile and as a navigational tool to explore it, in a coordinated view with other visual elements of the portrait.

Other visualization techniques such as *WordTrees* [65] and *PhraseNets* [61] have been employed to depict structure in text. Although we model data portraits as bipartite graphs between user interests (keywords) and user generated content (micro-posts), we do not focus on relations between words nor

the text structure, and thus, we use *word clouds* instead of any of the above mentioned techniques.

Visualization of micro-blog data covers a wide range of applications, including event monitoring [21, 43], visual analysis [19], group content analysis [3], and ego-networks [38]. We use *data portraits* [20] to visually represent user profiles. Data portraits are “*abstract representations of users’ interaction history*” [66], and have been built for personal informatics systems [5], Twitter profiles [22] and discussion forums [66], among others. Particularly, we borrow heavily from the work of Viégas *et al.* [62] in visualizing e-mail archives. E-mail archives are “*known datasets*”, and users have expectations of what they will find on such portraits. If those expectations are fulfilled, the enjoyment of the application is greater. In addition, by reflecting on their past, users discover unknown patterns about themselves and their relationships. We rely on this potential enjoyment and discovery to provide a positive experience for users, as a way to mitigate the cognitive dissonance that could be caused by our recommendations.

Graphical techniques used in recommendation contexts include controllable Venn-diagrams [50], network graphs [32, 33], *dust and magnet* [2], and compound graphs [27]. These are targeted at either expert users who know how to control such visualizations, or task-based systems. Instead, similarly to the *Hax* application [57], we propose *circle packing*, a hierarchical visualization technique [17], to create a casual and user friendly depiction of our generated recommendations, as it does not rely on user-controllability nor user expertise. However, differently from *Hax*, which builds a user interface to find audiences to broadcast information, our user interface allows users to find people to interact with.

METHODS AND DESIGN RATIONALE

For our work, we consider the micro-blogging platform Twitter, where users publish micro-posts known as *tweets*, each having a maximum length of 140 characters. Each user can *follow* other users, making their tweets available in his/her own *timeline*. Because of homophily [44], people tend to connect with like-minded individuals and do not think of the possible benefits of connecting with people of different (including opposing) opinions in sensitive issues. We design an intelligent system that recommends people to a target user, to study: (1) what factors affect the exploration and acceptance of recommendations built having politically diverse profiles into consideration, and (2) how to encourage unbiased behavior.

We introduce a data portrait design used in conjunction with a recommendation algorithm to recommend politically diverse profiles who share intermediary topics [29, 30]. Although the data portrait is about each user’s own data, we use it as a context to visually present the generated recommendations.

Portraying People’s Data

Data portraits [20] can be used to create a self-image to present the target user. Based on the enjoyment and self-reflection experimented in “*known data*” scenarios [62], our rationale is that we reinforce the non-conflicting interests of users when browsing their own profiles, while allowing to contextualize recommendations according to these interests.

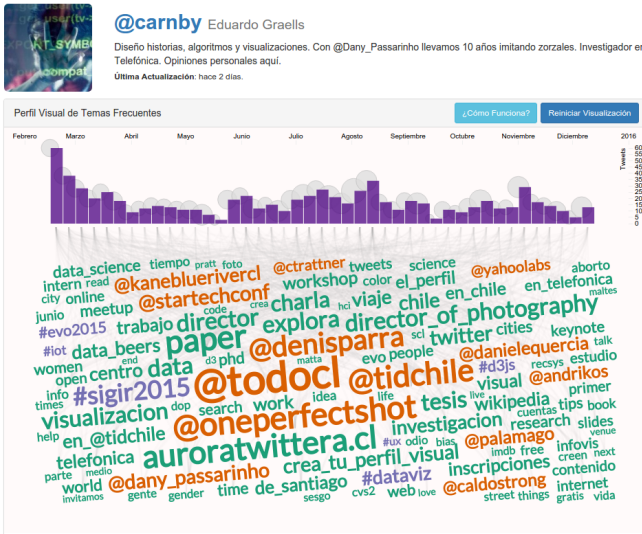


Figure 1: Our data portrait design. In the image, the portrait of the Twitter account *@carnby* (of Eduardo Graells-Garrido).

Our proposed design is displayed in Figure 1. User interests are estimated by counting the frequencies of the n -grams (with $n \leq 3$) of the words appearing in their timeline. The word cloud layout is based on *Wordle* [64], allowing for a tighter yet flexible representation of words.¹ Note that common word clouds follow two patterns of rotation: random angle or 90° . The first pattern makes reading the elements of the word cloud hard, whereas the second provides some sense of structure usually not present in the data. To promote a playful appearance, we decided to use rotated text. We used a fixed rotation for all words of -7° . This value was chosen arbitrarily after manual experimentation. As a way to compensate for the slight text rotation, to maintain readability we used a sans-serif font [55].

The color coding of word cloud elements is based on the type of keyword. We consider three categories: *hashtags* (#7570b3), *mentions* (#d95f02) and *regular words* (#1b9e77). This palette is based on the color-blind friendly *Set2* palette by Harrower *et al.* [36], although we darkened the colors to provide a better contrast. Each word has an invisible box that serves as clickable area, and as indicator when a particular word is highlighted when the box is visible.

To provide temporal information, we include a histogram of tweeting activity. This histogram encodes the number of tweets published or retweeted during a given time window. The number of bins is computed using Sturge’s formula ($k = \lceil \log_2 n + 1 \rceil$) by the *d3.js* library [10]. Each bin is accompanied with a circle positioned on its upper-left corner, which is a turn-on/off switch for a tweet to be displayed in an overlay window. Although all circles are similar in size, their ratios vary slightly according to the popularity of the most popular tweet in the bin. Circles allow the user to select a bin regardless of its size, a feature particularly useful for time windows with low activity.

¹We use the implementation by Jason Davies available at <http://www.jasondavies.com/wordcloud/>.

Finally, although users may want to change the background image of the data portrait, just like they can change the background on Twitter, we did not consider background customization to keep full control of aesthetics. Instead, we added their avatars and self-reported descriptions.

Interactions and Component Linkage. We link user interests and time (bins) using *Bézier* curves. The links are always visible, to make the structure behind the data portrait explicit to the user. Those links are displayed in a non-highlighted state. To highlight links and change the state of the portrait, the following interactions are available:

- When users click on a specific word, the corresponding bins are highlighted and connected through *Bézier* curves (see Figure 2, right).
- A click on a specific bin has two consequences:
 - A tweet overlay is displayed with the most popular tweet in it (see Figure 2, left). This tweet is context-dependent: if no word was selected before, it displays the overall most popular tweet; otherwise, it displays the most popular tweet relevant to the corresponding user interest. When a tweet is overlaid, the circle assigned to the current bin is highlighted.
 - The words related to all tweets in the bin are highlighted (see Figure 2, center).
- When displaying a tweet overlay, if the user clicks its corresponding highlighted circle, the circle is desaturated (*i. e.*, colored in gray), and the tweet overlay is hidden (see Figure 2, center).

The visualization can be reset by clicking on the “Reset Portrait” button (“Reiniciar Visualización”). Additionally, we display a “How it Works?” button (“¿Cómo Funciona?”), which then displays a pop-up window with instructions.

Influence and Pilot Study. Our proposed design is built upon the results of a pilot study [28] and the design of *Themail* [62], a data portrait of e-mail conversations. In the pilot study, a data portrait was built following an organic design to depict user profiles with recommendations injected inside it. The design was positively received by users, because displaying their interests allowed them to discover new things about themselves. It had, however, no impact on user behavior in terms of recommendations. In addition, several issues with the interface were raised, which were addressed in the new design (as shown in Figure 1), namely: consideration of time, a meaningful color palette for words, personalization of the data portrait, and better readability.

Recommending People Using Intermediary Topics

Our proposed approach aims at recommending people that, while having a potentially distant position in political issues, have shared interests with the target user. We do so by defining a proof-of-concept recommender algorithm. In this algorithm, the scoring of each recommendation is calculated by weighting two scores: 1) a content-based distance using LDA; and 2) an user similarity score based on *intermediary topics* [29], which we use as a proxy of shared interests. By weighting factors, this allows selecting a candidate that has intermediary topics with the target user, but is distant in terms of profile-wide

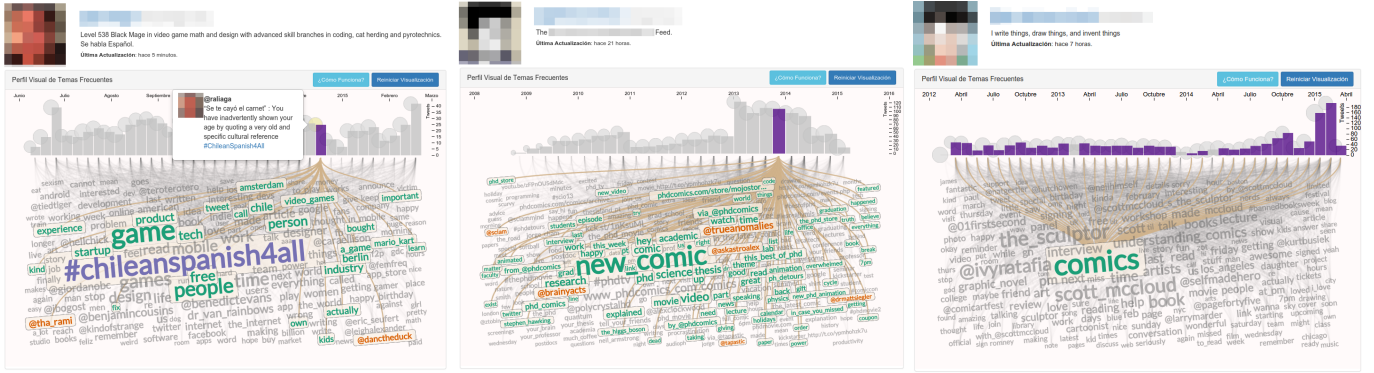


Figure 2: State of the data portrait after several interactions. Left: a click on the histogram will display a tweet overlay, with links to all related keywords to the corresponding bin. Center: a click on the bin circle will deactivate the tweet overlay to ease exploration. Right: a click on a keyword will link related histogram bins. The source profile identities have been pixelated to maintain anonymity.

topics as captured by LDA if the target user is politically vocal. On the other hand, users that are similar in political content but with no intermediary topics would not be recommended to each other, as recommendations are based on shared interests characterized through intermediary topics.

These topics are obtained by running LDA over a corpus of user documents or *microblogs* [53], and then creating a topic graph where nodes are topics. Two topics are connected through a weighted edge if both topics contribute content to at least one user. Edge weights are based on the fraction of users that contribute to each edge. Then, weighted closeness centrality is estimated in the graph, as a measure of how topics can make people closer. Intermediary topics are defined as those in the top 50% central topics, and have been found to be shared among a politically diverse set of people [30].

Notations. We represent a user u as a vector:

$$\vec{u} = [p_0(u), p_1(u), \dots, p_k(u)],$$

where k is the number of latent topics, and $p_i(u)$ is $P(t_i | u)$ as defined by the LDA model for a topic t_i . Next, given two users, u_1 and u_2 , we define their topical distance as the normalized *Kullback-Leibler Symmetric Distance* [8]:

$$KLD(u_1 || u_2) = \sum_{i=0}^k \{ \vec{u}_1[i] - \vec{u}_2[i] \} \log \left(\frac{\vec{u}_1[i]}{\vec{u}_2[i]} \right).$$

To normalize a distance into the range $[0, 1]$, given a set of distances, we divide each one by the maximum distance found.

Similarity Considering Intermediary Topics. We estimate *similarity with respect to intermediary topics* and not with respect to all topics, as we want to recommend users who might not be close in distance terms (*e. g.*, they could have ideological differences), but share intermediary topics. The set of intermediary topics for user u is defined as:

$$IT(u) = \{ i : \vec{u}[i] \geq \varepsilon, t_i \text{ is an intermediary topic} \},$$

where ε is a threshold for topical significance, which depends on the context. For instance, the default value used in the *gensim* library is 0.01 [54].

We define similarity with respect to intermediary topics as the *Jaccard Similarity* between two users:

$$JIT(u_1, u_2) = \frac{|IT(u_1) \cup IT(u_2)|}{|IT(u_1) \cap IT(u_2)|}.$$

Using this formula, when two users share all intermediary topics, $J(u_1, u_2) = 1$, and when users do not share any intermediary topic, $J(u_1, u_2) = 0$.

Algorithm Formalization. Each candidate for recommendation is scored using a *F-Score* [6] of *latent topical distance* and *similarity with respect to intermediary topics*:

$$\text{score} = (1 + \gamma^2) \times \frac{S(u_1, u_2) \times (1 - D(u_1, u_2))}{\gamma^2 \times (1 - D(u_1, u_2)) + S(u_1, u_2)}.$$

where S is similarity, D is distance, and the balance factor γ indicates the importance given to the distance in comparison to the importance given to similarity. For instance, $\gamma = 1$ gives equal importance to both factors, $\gamma = 0.5$ gives more importance to distance, and $\gamma = 2.0$ gives more importance to similarity.

Having estimated a measure of how close two users are, as well as how similar their sets of intermediary topics are, we can formalize our algorithm to recommend people with intermediary topics as follows. Given a target user u , a candidate set of recommendations C , a balancing factor γ , and the number of desired recommendations n , we estimate the score for all candidates and return the top- n scored candidates.

Displaying Recommendations with Circle Packing

The set of recommendations is displayed below the main data portrait as a separate unit, although both are clearly part of the same system.

Our algorithm generates a list of recommended accounts to follow. Each recommended account contains an avatar, a username, a biography, and a link to the full profile on Twitter. For each of those accounts we know the set of latent topics according to the LDA model. To prepare recommendations for visualization, we cluster accounts based on their common latent topics. We implemented a simple scheme, where two

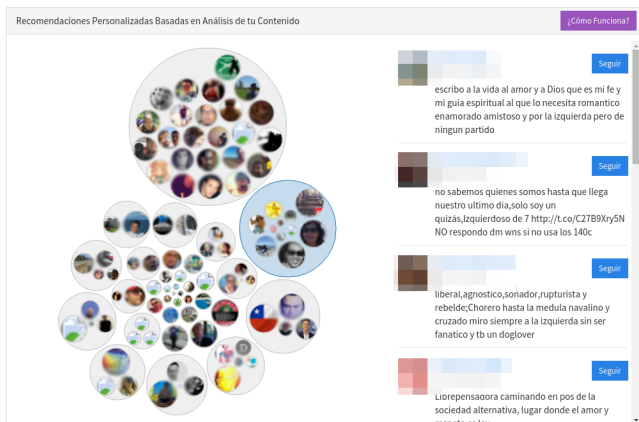


Figure 3: Display of recommendations using Circle Packing. The recommended profile identities have been pixelated to maintain anonymity.

users are in the same cluster if their most contributing latent topic is the same. More complex clustering methods could be explored, but this is not the focus of this paper.

Visual depictions have the potential to change how users perceive recommendations. First, visualization of social recommendations has been shown to increase user satisfaction [32]. Second, explaining recommendations is important, as explanations increase user involvement and acceptance [40]. Using visualization techniques to display recommendations allows us to depict the underlying structure behind them, hence providing an *implicit* explanation. Conversely, when using text only, recommendations have to be explained in natural language, since something like “Topic 5” is meaningless. Hence, visualization is a natural way to overcome this.

We employ an *enclosure diagram* built using *Circle Packing* [17] to display clustered recommendations. Circle Packing (CP hereafter) is a hierarchical visualization technique, where each node of the hierarchy is represented as a circle, with nesting according to the hierarchy. Even though this kind of diagram does not use effectively all the space available (for instance, unlike treemaps [11]), CP “*effectively reveals the hierarchy*” [39], and has an attractive organic appearance. Additionally, circles maintain aspect ratio (unlike cells in a treemap), which is useful to display avatars of different sizes, as seen on previous applications in social query systems [57]. An example visualization is shown in Figure 3.

Interactivity. Our design shows the CP visualization of the hierarchical structure of recommendations, but no actual detailed recommendation is shown at first. A message indicating that users can interact with the visualization is displayed instead. Then, when users click on a cluster, the cluster is highlighted and a list on the right of the visualization displays a detailed list of the corresponding recommendations. In the detailed list, each recommendation profile contains the account name, the full name of the recommended user, and the self-reported biography. The account name is linked to his/her profile on Twitter, and a “Follow” (“*Seguir*”) button allows the user to directly follow the recommended account.

DEPLOYMENT “IN THE WILD”

We detail how we tested our platform, by deploying a proof-of-concept implementation in an uncontrolled setting [18].

Aurora Twittera: A Chilean News Aggregator Platform

We implemented the visual designs using the *d3.js* library [10], and the recommendation algorithms using the LDA implementation from the *gensim* library [54]. We incorporated both, the data portrait design and the recommender system, in a Web platform named *Aurora Twittera* [31] (available at <http://auroratwittera.cl>), *AT* hereafter. *AT* is targeted at Chilean users (hence the `.cl` domain). It is a news aggregator that constantly crawls tweets about Chilean news, specially with respect to political events, but also news on other topics (*e. g.*, sports, cultural events).

Building Portraits

In *AT*, users could create their “Visual Profiles” (“*Perfiles Visuales*”) by connecting their Twitter accounts with the site. Since users had to log-in to the site to browse their portraits, we could gather rich interaction data, which we then analyzed.

After pressing the “*Create Your Profile*” button, users were redirected to the Twitter website, which asked for login credentials and permission to modify their accounts. We asked for these permissions to have a “*Follow*” button next to each recommendation. A scheduler service processed queued portraits, both, those newly created and those queued for update. Tweets were downloaded using the Twitter API using the user credentials. Then, we estimated the user interests according to the methodology described in the previous section and identified the intermediary topics using LDA. We considered the top-300 user interests according to frequency, a number chosen based on the width and height of the data portraits, which is suitable for a 1024x768 screen resolution.

From the up-to-date dataset of users in *AT* we generated, every day, a list of candidate people who tweeted in the previous 48 hours, and for whom we estimate their latent topics using the entire corpus of users who published tweets in those 48 hours. This regularly updated list of candidates allowed us to present fresh recommendations to users.

Promoting the System

AT has a social bot in Twitter, with username *@todocl*. The social bot *@todocl* published tweets mentioning users when their portraits were ready (usually within less than one minute after sign-up), and every three days when their portraits were updated. Although portrait updates were daily, notification was limited to every three days to avoid spamming. Additionally, to promote our system we performed a number of actions:

- Created several *demo portraits* for people to browse, and publicized them on *@todocl*’s timeline. The demo portraits were of popular user accounts. Sometimes the portrayed users, when being mentioned to notify them about the availability of their portraits, retweeted our announcements.
- Created a campaign on <http://ads.twitter.com> aimed at Chilean desktop users in Twitter who were active for at least one month. As result, 42,190 promoted tweets were

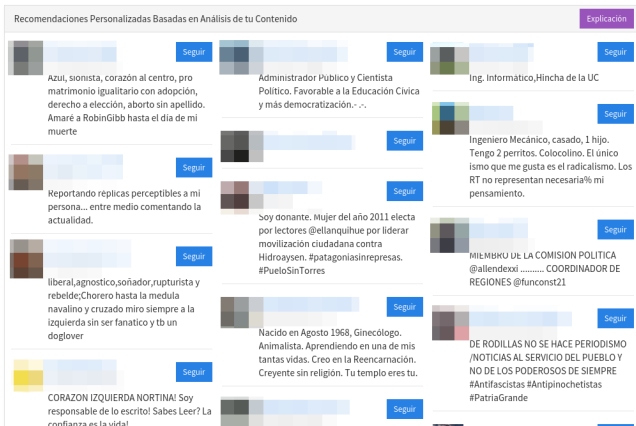


Figure 4: Baseline design of recommendations. The recommended profile identities have been pixelated to maintain anonymity.

displayed, with an engagement rate (as reported by Twitter) of 0.51%.

- Added a “Share my Profile” button to the data portrait. When clicked, the system published a tweet from the portrayed user’s account, inviting her/his followers to visit the data portrait.

The system was open to everyone. However, the user interface was available in Spanish only, and recommendations considered only Chilean users.

EVALUATION WITH INTERACTION DATA

Our hypothesis is that an indirect approach, through the mixture of data portraits and recommendations using intermediary topics, allows users to overcome the cognitive dissonance produced by exposure to potentially challenging information. To test this hypothesis our evaluation followed a *between-groups* design. For each user who signed up on the system, a random pair of conditions $\langle UI, RecSys \rangle$ was assigned. In both cases, user interface and recommender system, we considered a baseline condition in addition to our proposed one.

The *User Interface* conditions were:

- *Baseline*: The baseline recommendation UI (see Figure 4), which displays recommendations in a similar way to current mainstream user interfaces.
- *Circle Pack*: The visualization of recommendations using circle packing (see Figure 3 Top).

The *Recommender System* conditions were:

- *KLD*: recommendations generated using Kullback-Leibler Symmetric Distance only, *i. e.*, considering the most similar users according to all LDA topics.
- *IT*: recommendations generated using our proposed method based on intermediary topics (with balancing factor $\gamma = 1$).

In this way, we can compare whether using visualization influences how users explore recommendations, and whether including intermediary topics in the recommendation makes users behave in a less biased way.

For each user, we had the following independent variables:

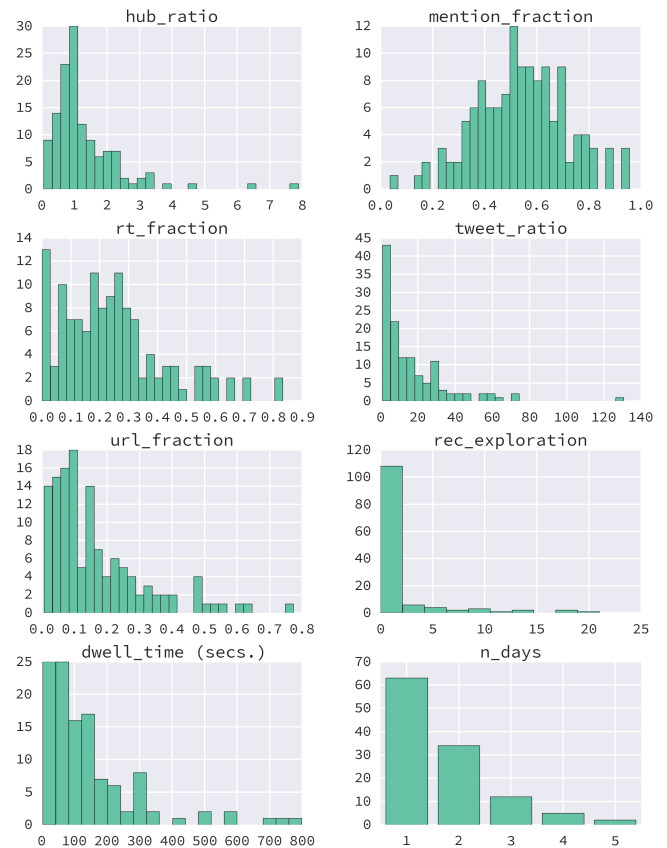


Figure 5: The first five plots display the distribution of characteristics (independent variables) of portrayed users. The last three plots display the distributions of interaction data variables under study.

- *Political Content*: Its value is 1 if the list of the top-50 user interests has a non-empty intersection with a list of political keywords (including hashtags); the value is 0 if the intersection is empty. We considered this variable as binary, *e. g.*, a user could use few political keywords, but if used very often, then it is arguably a user interested in politics.
- *Hub Ratio*: Number of followed accounts divided by the number of followers, which measures the tendency of users to follow others based on their own popularity.
- *Mention Fraction*: Fraction of tweets that mention someone else (excluding retweets).
- *RT Fraction*: Fraction of tweets that are retweets.
- *Tweet Ratio*: Number of tweets per day, defined as total number of tweets published divided by the account age.
- *URL Fraction*: Fraction of tweets that contain a URL (excluding retweets).

The interaction data we considered are explicit actions performed on the user interface:

- *Recommendation Exploration*: Number of clicks on elements in the user interface related to recommendations (*e. g.*, click on a profile link or circle pack nodes).
- *Recommendation Acceptance*: Whether the user accepted at least one recommendation.

Since *AT* is a *Casual InfoVis System* [52], users were not expected to perform specific tasks, nor instructed to do so. Although we compare interaction data between conditions to evaluate the differences in behavior, we still require a feedback mechanism to understand how users perceive the system. We analyze user engagement using *implicit feedback* [41], where positive user engagement is used as our proxy of a positive perception of the system through the following variables:

- *Number of Days*: Number of days the user visited his/her data portrait.
- *Dwell Time*: Time (in seconds) spent interacting with or browsing the data portrait with embedded recommendations.

For our analysis, we consider the following factorial model:²

$$Y = C(ui) \times C(\text{recommendation}) \times C(\text{pol.content}) \\ + \text{tweet_ratio} + \text{hub_ratio} + \text{RT_fraction} \\ + \text{URL_fraction} + \text{mention_fraction},$$

where $C(X)$ creates dummy variables for the corresponding categories of the independent variable X , and \times represents the independent factors and the interactions between them.³ We use this model for two types of regression: Negative Binomial (NB) generalized linear models, and logistic (logit) regression. NB is used for over-dispersed count data, and logit is used to model dichotomous outcomes in terms of probabilities. In both models, if the statistical interactions between factors were found to be not significant, we analyzed the same model without interaction terms.

Participants

As our study focuses on Chilean users and Chilean politics, we discarded users whose self-reported Twitter location was not Chile, or whose IP address was not detected as Chilean by the GeoIP database. We also discarded users whose interaction data was not reliable, *e. g.*, having Javascript-blocking extensions in their browser. Lastly, we discarded users who spent less than 5 seconds on the site, and those whose tweet ratio was less than one.

As result, we have 129 valid portraits, created between 18 February and 17 March 2015. For the recommendation conditions, 59 users were assigned to *KLD*, and 70 to *IT*. With respect to the user interface, 59 users were assigned to the *Baseline*, and 70 to the *Circle Pack* condition. Finally, 69 users had political content in their portraits, and 60 did not. The means of independent variables are: *hub ratio*, 1.30; *mention fraction*, 0.54; *RT fraction*, 0.25; *tweet ratio*, 16.54; and *URL fraction*, 0.17. Figure 5 shows the distributions of these independent variables.

Regression Results

The 129 portrayed users generated 1,707 interaction events. The following are the mean and max values found for each variable, as well as the results of the statistical analysis:

- *Recommendation Exploration* events (mean = 1.53, max = 21): the first NB regression did not contain significant interactions. The NB model without interactions is (log-likelihood = -169.59; deviance = 168.51; $\chi^2 = 221$).
- *Recommendation Acceptance*: 5.42% of participants accepted at least one recommendation. The first logit regression did not contain significant interactions. The logit model without interactions is (log-likelihood = -14.67, $p = 0.002$).
- *Number of Days* (mean = 1.81, max = 8): none of the performed NB regressions reported significant terms.
- *Dwell Time*: We discarded the top decile from the analysis because some users left the browser window open (the maximum dwell time observed was 9 hours), leading to $N = 116$, with mean = 147.66, and max = 798 seconds. The NB model with interactions contained significant terms (log-likelihood = -688.68, deviance = 95.89, $\chi^2 = 80.5$).

Figure 5 shows the distributions of the independent variables. Table 1 displays the regression coefficients with a p-value smaller than 0.05, and the corresponding 95% confidence intervals and p-value. We refer to each result as R_i . We discuss our results focusing first on recommendations, and then user engagement. As effect size for the logit regression coefficients, we consider the Odds-Ratio (OR) of each result. The OR is a measure of how associated a factor is to the outcome under analysis. If the OR is greater than 1, then the presence of the factor is considered to be associated with the outcome. If the OR is lesser than 1, then the opposite association holds. An OR of 1 indicates no association. The OR of each coefficient is defined as $OR = exp(\beta)$. In the case of the NB regression, the effect size of each coefficient is defined as $ES = exp(|\beta|) * sign(\beta)$. This ES means how much would the outcome increase (or decrease) with a one unit increase of the dependent variable.

Recommendation Exploration

We discuss the extent to which users explored the recommendations, and whether they accepted them (users followed the recommended users).

We found two positive effects that increase the tendency to explore recommendations. The strongest effect is the RT Fraction (R5, ES = 41.721). This can be interpreted as users who tend to retweet more, are more likely to explore recommendations because they are looking for sources to retweet. The second positive effect is the usage of Circle Pack (R2, ES = 11.750). This ES indicates that exploration recommendations increase when users are exposed to CP, if all the other factors are held constant. This effect validates our design choice of using CP.

The negative effects are the usage of Intermediary Topics (IT) and Tweet Ratio. In the case of IT, users exposed to its generated recommendations decrease exploration (R3, ES = -3.157). This effect indicates that users tend to behave homophilically, probably because the recommendees' profile information can make their political leaning explicit (*i. e.*, by using a politically-explicit avatar or self-reported description). However, its effect size is small in comparison to the positive effects found. Likewise, a one unit increase in Tweet Ratio

²Specified in R's formula syntax.

³ $A \times B = A + B + A * B$.

R#	DV	IV	β	Effect Size	95% C.I.	p-value
R1	Rec. Exploration	Intercept	-2.104	-	[-3.961, -0.247]	0.026
R2	Rec. Exploration	UI(CP)	2.464	11.750	[1.510, 3.417]	< 0.001
R3	Rec. Exploration	REC(IT)	-1.150	-3.157	[-1.937, -0.362]	0.004
R4	Rec. Exploration	Tweet Ratio	-0.030	-1.031	[-0.053, -0.007]	0.012
R5	Rec. Exploration	RT Fraction	3.731	41.721	[1.468, 5.994]	0.001
R6	P(Rec. Acceptance)	REC(IT)	-4.383	0.012	[-8.547, -0.219]	0.039
R7	P(Rec. Acceptance)	Pol. Content(True)	3.354	28.617	[0.061, 6.647]	0.046
R8	P(Rec. Acceptance)	Mention Fraction	-12.079	5.68×10^6	[-21.958, -2.200]	0.017
R9	P(Rec. Acceptance)	RT Fraction	9.872	19,380.060	[1.026, 18.717]	0.029
R10	Dwell Time	Intercept	4.825	-	[3.971, 5.680]	< 0.001
R11	Dwell Time	UI(CP), Pol. Content(True) and REC(IT)	2.187	8.908	[0.785, 3.590]	0.002

Table 1: Regression Coefficients for the dependent variables under study. Only significant terms are shown for each regression.

decreases exploration (R4, ES = -1.031). This may indicate that users who tend to publish more tweets are less likely to explore recommendations because they are generating content—instead of looking for sources; they *are* the sources.

Recommendation Acceptance

There are two significant positive effects. The strongest one is RT Fraction (R9, OR = 19380.06). We interpret this in concordance with previous interpretations, where users were looking for sources of information to retweet from. The second positive effect is the presence of political content on a user portrait (R7, OR = 28.617). Hence, if the other variables are held constant, the odds of accepting a recommendation increase more than 28 times if the user is interested in politics. This aligns with results obtained in our first pilot study [28].

The negative effects are Intermediary Topics and Mention Fraction. In the case of IT (R6, OR = 0.012), this effect size confirms the homophilic behavior of the user population as already hinted by the previous result in recommendation exploration. The other negative effect is Mention Fraction (R8, OR = 5.68×10^6), which suggests that users with high fraction of mentions have almost negligible odds of accepting a recommendation. This suggests that users who already have a network of connections to interact with, do not need (or feel the need) to add new people in their networks.

User Engagement

We defined two variables related to user engagement: number of days that each user visited the site, and dwell time. Both variables allow to measure positive engagement with the site: someone who returns to the site in a different day may do so because s/he finds it useful, and someone who spends more time on the site, in a single session, may do so because s/he finds it interesting.

We observed that 45% of participants returned to the site at least a second time on a different day. However, what causes this return cannot be explained by our regression model, implying that in terms of visits per day, all users were engaged equally. With respect to dwell time, none of the standalone variables or main effects were found to be significant, nor the pairwise interactions. The only factor found significant is the triple interaction between Intermediary Topics, Circle Pack and Political Content (R11, ES = 8.91). In the NB model, this can be interpreted as follows: dwell time increases by

8.91 seconds when these three conditions are present, and all other factors are held constant. Note that, if all possible interactions and main effects related to user interface, recommender system, and political content would be significant, there would still be an increase in dwell time by 2.06 seconds when those three conditions are present. This could be because of increased satisfaction with the system [25], or a deeper exploration of profiles.

Overview of Results

Some results, positive and negative, support the motivation behind this work. On one hand, visualization of recommendations increased exploration. On the other hand, our proposed algorithm was not interesting for users, and was out-performed by a fully-homophilic baseline. Our results are not enough to fully support our hypothesis, although we found partial evidence in its favor: there is an effect of our mixed approach, but not for all users.

It is well known that a majority of users are *challenge-averse* [46], so it could be expected that politically-engaged users exposed to our conditions would have had a negative (*e.g.*, less engaging) experience. This was not the case – users who were exposed to our proposed conditions and are politically-involved present a comparable (and even slightly more positive) experience when using the site. In the next section we discuss why this positive experience can be linked with conscious (unbiased) behavior.

DISCUSSION

Recommendation and Individual Differences. To analyze individual differences, we focused on behavioral signals that could be extracted from user profiles, and applied a statistical model to find which ones influence user behavior with our system. Some variables were not significant: *hub ratio* (connectivity) and *URL fraction* (type of content that is shared). The significant variables were *tweet ratio* (publishing behavior), *mention fraction* (interaction with others), and *RT Fraction* (information diffusion). Knowing that some users have the tendency to explore (or not) recommendations allows identifying the users who are more likely to benefit from them, and use simpler algorithms for those users who are not.

Furthermore, the user related variable that influenced recommendation acceptance was the presence of political content.

In line with our motivation, arguably only politically-involved people are affected by selective exposure, in the sense that they look for political content, whereas non-politically involved people discard political content because of lack of interest instead of selective exposure. Not all users are interested in politics, therefore, not all users are interested in, nor need, political diversity on their timelines.

Our visualization proposal was effective: when visualizing recommendations instead of using the baseline text interface, users' exploration of recommendation was equivalent, or even greater, than when recommendations were non-diverse. We used an aesthetically attractive design based on circle packing [17], which had the property of displaying part of the underlying structure in recommendations. In that aspect, we did not find a main effect of visualization in reducing homophilic behavior of recommendation acceptance. In fact, homophilic behavior was confirmed: the non-diverse recommendation algorithm increased the likelihood of acceptance. However, due to the statistical interaction found, there is a potential usage of visualization when mixed with algorithms that recommend more diverse information.

“Indirect” Approaches and Unbiased Behavior. As an indirect approach to exposing users to people of opposing views, we proposed a data portrait [20] context for users to explore their own content. These data portraits serve to display recommendations, where users could explicitly see how the system modeled their generated content, and implicitly understand the structure of the recommendations made to them. Thus, when addressing the biased behavior of users, micro-blogging platforms may want to include other profile UIs, specially those based around data portraits [20] and *Casual InfoVis* systems [52].

Previous attempts at directly exposing users to opposing information had detrimental effects, given that a majority of users are *challenge-averse* [46]. For instance, in *Opinion Space* [23], although users did behave differently when using a visualization-based UI, their dwell times were not significantly different, and their political behavior was not less biased. Our own results showed that politically-involved users under our proposed conditions did not behave in a less biased way, although they had slightly longer sessions. This is the key difference – their extra-time browsing recommendations (from 2 to 8 seconds, with an average of 1.53 explored profiles) indicates that, whether they accepted or discarded recommendations, they took their time to explore/decide. We hypothesize that this is a *conscious* decision-making process when exploring recommendations. Note that such consciousness cannot be achieved by direct approaches because they activate cognitive dissonance [24].

Limitations and Future Work. In this work we used a proof-of-concept recommendation algorithm, based on the concept of intermediary topics. Although these topics have been shown to contain politically diverse profiles in the population under study [30], there is a need to quantify in formal terms its differences with fully homophilic algorithms. For instance, we focused on the political diversity of users provided by intermediary topics, but not on whether they actually had opposing

views in sensitive issues or not. Moreover, critics might rightly say that sentiment analysis could be used as proxy for opposed opinions instead of latent topics. However, these techniques are not language-agnostic and culture-independent, something relevant in our context (a population with Spanish as primary language), and thus this should be analyzed and tested as well. Finally, our results indicate that our algorithm has lesser acceptance than baseline algorithms. Perhaps by considering more features expected by users, like network features [14], results could be improved.

CONCLUSIONS

Behavioral and content differences influence how users perceive a *Casual InfoVis* system [52]. In our research, while trying to understand how to encourage connecting with others who think differently, we found that being politically open, and the informational and interaction behavior of users, are important features that influence interaction. This is important to be aware for the following reason. When designing systems for specific tasks, user characteristics can be assumed by visualization designers. However, in open systems like ours we cannot predict who will use the system nor their expertise level. In addition, in biased scenarios there does not seem to be a one-size-fits-all solution. Because we presented an exploratory system, our evaluation was not a task-based one focused on algorithm/visualization efficiency. Instead, we performed an “in the wild” evaluation [18], where we focused on individual differences and user interaction with the interface, as well as user engagement metrics [41]. This allowed us to obtain deep insights on user behavior and exploratory styles.

To conclude, our results show that systems that aim at unbiasing user behavior should consider an indirect approach, which can be evaluated with engagement metrics in addition to context-specific metrics (*e. g.*, recommendation acceptance). Furthermore, our work also shows that an unbiased behavior is not one that performs more actions (*e. g.*, accepting more recommendations of people who thinks differently). On the contrary, it is one that allows users to avoid the cognitive heuristics that bias the activity, enabling them to make conscious choices.

Acknowledgments. We are grateful to Daniele Quercia for inspiration and discussion. We thank Shiri Dori-Hacohen and the anonymous reviewers for valuable feedback, and Andrés Lucero for help with the user studies. This work was partially funded by Grant TIN2012-38741 (Understanding Social Media: An Integrated Data Mining Approach) of the Ministry of Economy and Competitiveness of Spain.

References

1. Al Zamal, F., Liu, W., and Ruths, D. Homophily and Latent Attribute Inference: Inferring Latent Attributes of Twitter Users from Neighbors. *International Conference on Weblogs and Social Media*, 270, 2012:
2. An, J., Quercia, D., Cha, M., Gummadi, K., and Crowcroft, J. Sharing political news: the balancing act of intimacy and socialization in selective exposure. *EPJ Data Science*, 3(1), 2014: 1–21.

3. Archambault, D., Greene, D., Cunningham, P., and Hurley, N. ThemeCrowds: multiresolution summaries of Twitter usage. *Proceedings of the 3rd international workshop on Search and mining user-generated contents*. 2011, 77–84.
4. Asch, S. E. Forming impressions of personality. *The Journal of Abnormal and Social Psychology*, 41(3), 1946: 258.
5. Assogba, Y., and Donath, J. Mycrocosm: visual microblogging. *42nd Hawaii International Conference on System Sciences*. 2009, 1–10.
6. Baeza-Yates, R., and Ribeiro-Neto, B. Modern Information retrieval: the concepts and technology behind search, 2nd. Edition. Addison-Wesley, Pearson, 2011.
7. Barberá, P. Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data. *Political Analysis*, 23(1), 2015: 76–91.
8. Bigi, B. Using Kullback-Leibler distance for text categorization. In: *Advances in Information Retrieval*. Springer, 2003, 305–319.
9. Blei, D. M., Ng, A. Y., and Jordan, M. I. Latent Dirichlet allocation. *The Journal of Machine Learning Research*, 3, 2003: 993–1022.
10. Bostock, M., Ogievetsky, V., and Heer, J. D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12), 2011: 2301–2309.
11. Bruls, M., Huizing, K., and Van Wijk, J. J. Squarified treemaps. Springer, 2000.
12. Brzozowski, M. J., and Romero, D. M. Who should I follow? Recommending people in directed social networks. *International Conference on Weblogs and Social Media*. 2011.
13. Center, P. R. Internet seen as positive influence on education but negative on morality in emerging and developing nations. <http://www.pewglobal.org/2015/03/19/internet-seen-as-positive-influence-on-education-but-negative-influence-on-morality-in-emerging-and-developing-nations/>. 2015.
14. Chen, J., Geyer, W., Dugan, C., Muller, M., and Guy, I. Make new friends, but keep the old: recommending people on social networking sites. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2009, 201–210.
15. Chen, K., Chen, T., Zheng, G., Jin, O., Yao, E., and Yu, Y. Collaborative personalized tweet recommendation. *Proceedings of the 35th international ACM SIGIR conference on Research and Development in Information Retrieval*. 2012, 661–670.
16. Chhabra, S., and Resnick, P. Does clustered presentation lead readers to diverse selections? *CHI'13 Extended Abstracts on Human Factors in Computing Systems*. 2013, 1689–1694.
17. Collins, C. R., and Stephenson, K. A circle packing algorithm. *Computational Geometry*, 25(3), 2003: 233–256.
18. Crabtree, A., Chamberlain, A., Grinter, R. E., Jones, M., Rodden, T., and Rogers, Y. Introduction to the special issue of “The Turn to The Wild”. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 20(3), 2013: 13.
19. Diakopoulos, N., Naaman, M., and Kivran-Swaine, F. Diamonds in the rough: social media visual analytics for journalistic inquiry. *IEEE Symposium on Visual Analytics Science and Technology (VAST)*. 2010, 115–122.
20. Donath, J., Dragulescu, A., Zinman, A., Viégas, F., and Xiong, R. Data portraits. *ACM SIGGRAPH 2010 Art Gallery*. 2010, 375–383.
21. Dörk, M., Gruen, D., Williamson, C., and Carpendale, S. A visual backchannel for large-scale events. *Transactions on Visualization and Computer Graphics*, 2010: 1129–1138.
22. Dragulescu, A. Lexigraphs: Twitter data portrait: jakedfw. <http://vimeo.com/2404119>. 2009.
23. Faridani, S., Bitton, E., Ryokai, K., and Goldberg, K. Opinion space: a scalable tool for browsing online comments. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2010, 1175–1184.
24. Festinger, L. A theory of Cognitive Dissonance. Stanford University Press, 1962.
25. Fox, S., Karnawat, K., Mydland, M., Dumais, S., and White, T. Evaluating implicit measures to improve web search. *ACM Transactions on Information Systems (TOIS)*, 23(2), 2005: 147–168.
26. Goel, A., Sharma, A., Wang, D., and Yin, Z. Discovering similar users on Twitter. *11th Workshop on Mining and Learning with Graphs*. 2013.
27. Gou, L., You, F., Guo, J., Wu, L., and Zhang, X. L. Sfviz: interest-based friends exploration and recommendation in social networks. *Proceedings of the 2011 Visual Information Communication-International Symposium*. 2011, 15.
28. Graells-Garrido, E., Lalmas, M., and Quercia, D. Data portraits: connecting people of opposing views. [arXiv:1311.4658 \[CS.LG\]](https://arxiv.org/abs/1311.4658).
29. Graells-Garrido, E., Lalmas, M., and Quercia, D. People of opposing views can share common interests. *Proceedings of the companion publication of the 23rd international conference on World Wide Web (poster)*. 2014, 281–282.
30. Graells-Garrido, E., Lalmas, M., and Baeza-Yates, R. Finding intermediary topics between people of opposing views: a case study. *International Workshop on Social Personalisation & Search co-located with the 38th Annual ACM SIGIR Conference*. 2015.
31. Graells-Garrido, E., Lalmas, M., and Baeza-Yates, R. Encouraging diversity- and representation-awareness in geographically centralized content. *Proceedings of the 21st International Conference on Intelligent User Interfaces (to appear)*. 2016.

32. Gretarsson, B., O'Donovan, J., Bostandjiev, S., Hall, C., and Höllerer, T. Smallworlds: visualizing social recommendations. *Computer Graphics Forum*. 2010, 833–842.
33. Gretarsson, B., O'donovan, J., Bostandjiev, S., Höllerer, T., Asuncion, A., Newman, D., and Smyth, P. Topicnets: visual analysis of large text corpora with topic modeling. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3(2), 2012: 23.
34. Gupta, P., Goel, A., Lin, J., Sharma, A., Wang, D., and Zadeh, R. WTF: the who to follow service at Twitter. *Proceedings of the 22nd international conference on World Wide Web*. 2013, 505–514.
35. Hannon, J., Bennett, M., and Smyth, B. Recommending Twitter users to follow using content and collaborative filtering approaches. *Proceedings of the fourth ACM Conference on Recommender Systems*. 2010, 199–206.
36. Harrower, M., and Brewer, C. A. ColorBrewer. org: an online tool for selecting colour schemes for maps. *The Cartographic Journal*, 40(1), 2003: 27–37.
37. Hart, W., Albarracín, D., Eagly, A. H., Brechan, I., Lindberg, M. J., and Merrill, L. Feeling validated versus being correct: a meta-analysis of selective exposure to information. *Psychological bulletin*, 135(4), 2009: 555.
38. Heer, J., and boyd, danah. Vizster: visualizing online social networks. *IEEE Information Visualization (InfoVis)*. 2005, 32–39.
39. Heer, J., Bostock, M., and Ogievetsky, V. A tour through the visualization zoo. *Commun. Acm*, 53(6), 2010: 59–67.
40. Herlocker, J. L., Konstan, J. A., and Riedl, J. Explaining collaborative filtering recommendations. *Proceedings of the 2000 ACM conference on Computer supported cooperative work*. 2000, 241–250.
41. Lalmas, M., O'Brien, H., and Yom-Tov, E. Measuring user engagement. Morgan & Claypool Publishers, 2014.
42. Liao, Q. V., and Fu, W.-T. Can you hear me now? Mitigating the echo chamber effect by source position indicators. *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*. 2014, 184–196.
43. Marcus, A., Bernstein, M. S., Badar, O., Karger, D. R., Madden, S., and Miller, R. C. Twitinfo: aggregating and visualizing microblogs for event exploration. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2011, 227–236.
44. McPherson, M., Smith-Lovin, L., and Cook, J. M. Birds of a feather: homophily in social networks. *Annual review of sociology*, 2001: 415–444.
45. Michelson, M., and Macskassy, S. A. Discovering users' topics of interest on Twitter: a first look. *Proceedings of the fourth workshop on Analytics for noisy unstructured text data*. 2010, 73–80.
46. Munson, S. A., and Resnick, P. Presenting diverse political opinions: how and how much. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2010, 1457–1466.
47. Myers, D. G., and Lamm, H. The polarizing effect of group discussion: the discovery that discussion tends to enhance the average prediscussion tendency has stimulated new insights about the nature of group influence. *American Scientist*, 1975: 297–303.
48. Pariser, E. The filter bubble: what the Internet is hiding from you. Penguin UK, 2011.
49. Park, S., Kang, S., Chung, S., and Song, J. NewsCube: delivering multiple aspects of news to mitigate media bias. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2009, 443–452.
50. Parra, D., and Brusilovsky, P. User-controllable personalization: a case study with SetFusion. *International Journal of Human-Computer Studies*, 2015:
51. Pennacchiotti, M., and Popescu, A.-M. Democrats, republicans and Starbucks aficionados: user classification in Twitter. *Proceedings of the 17th ACM SIGKDD international conference on Knowledge Discovery and Data Mining*. 2011, 430–438.
52. Pousman, Z., Stasko, J. T., and Mateas, M. Casual information visualization: Depictions of data in everyday life. *IEEE Transactions on Visualization and Computer Graphics*, 13(6), 2007: 1145–1152.
53. Ramage, D., Dumais, S., and Liebling, D. Characterizing microblogs with topic models. *International Conference on Weblogs and Social Media*. 2010, 130–137.
54. Řehůřek, R., and Sojka, P. Software framework for topic modelling with large corpora. *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. May 2010, 45–50.
55. Rello, L., and Baeza-Yates, R. Good fonts for dyslexia. *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*. 2013, 14.
56. Rout, D., Bontcheva, K., Preotjiuc-Pietro, D., and Cohn, T. Where's @Wally?: a classification approach to geolocating users based on their social ties. *Proceedings of the 24th ACM Conference on Hypertext and Social Media*. 2013, 11–20.
57. Savage, S., Forbes, A., Toxtli, C., McKenzie, G., Desai, S., and Höllerer, T. Visualizing targeted audiences. *COOP 2014-Proceedings of the 11th International Conference on the Design of Cooperative Systems, 27-30 May 2014, Nice (France)*. 2014, 17–34.
58. Sunstein, C. R. The law of group polarization. *Journal of Political Philosophy*, 10(2), 2002: 175–195.
59. Sunstein, C. R. Going to extremes: how like minds unite and divide. Oxford University Press, 2009.

60. Valenzuela, S., Arriagada, A., and Scherman, A. Facebook, Twitter, and youth engagement: A Quasi-experimental study of social media use and protest behavior using propensity score matching. *International Journal of Communication*, 8, 2014: 25.
61. Van Ham, F., Wattenberg, M., and Viégas, F. B. Mapping text with phrase nets. *IEEE Transactions on Visualization and Computer Graphics*, 15(6), 2009: 1169–1176.
62. Viégas, F. B., Golder, S., and Donath, J. Visualizing email content: portraying relationships from conversational histories. *Proceedings of the SIGCHI conference on Human Factors in computing systems*. 2006, 979–988.
63. Viégas, F. B., and Wattenberg, M. Tag clouds and the case for vernacular visualization. *Interactions*, 15(4), 2008: 49–52.
64. Viegas, F. B., Wattenberg, M., and Feinberg, J. Participatory visualization with Wordle. *IEEE Transactions on Visualization and Computer Graphics*, 15(6), 2009: 1137–1144.
65. Wattenberg, M., and Viégas, F. B. The word tree, an interactive visual concordance. *IEEE Transactions on Visualization and Computer Graphics*, 14(6), 2008: 1221–1228.
66. Xiong, R., and Donath, J. PeopleGarden: creating data portraits for users. *Proceedings of the 12th annual ACM symposium on User interface software and technology*. 1999, 37–44.
67. Yi, J. S., Kang, Y.-a., Stasko, J. T., and Jacko, J. A. Understanding and characterizing insights: how do people gain insights using information visualization? *Proceedings of the 2008 Workshop on BEyond time and errors: novel evaluation methods for Information Visualization*. 2008, 4.