

Visualizing Targeted Audiences

Saiph Savage, Angus Forbes, Carlos Toxtli, Grant McKenzie,
Shloka Desai, and Tobias Höllerer

University of California, Santa Barbara (UCSB)
Universidad Nacional Autonoma de Mexico (UNAM)
University of Arizona (UA)
{saiph@cs., grant.mckenzie@geog., holl@cs.}ucsb.edu
{angus.forbes@sista., shlokadesai@email.}arizona.edu
{ctoxtli}@fi-b.unam.mx
<http://www.youtube.com/watch?v=vduSbDsAP30>

Abstract. Users of social networks can be passionate about sharing their political convictions, art projects, or business ventures. They often want to direct their social interactions to certain people in order to start collaborations or to raise awareness about issues they support. However, users generally have scattered, unstructured information about the characteristics of their audiences, making it difficult for them to deliver the right messages or interactions to the right people. Existing audience-targeting tools allow people to select potential candidates based on pre-defined lists, but the tools provide few insights about whether or not these people would be appropriate for a specific type of communication. We introduce an online tool, *Hax*, to explore instead the idea of using interactive data visualizations to help people dynamically identify audiences for their different sharing efforts. We provide the results of a preliminary empirical evaluation that shows the strength of the idea and points to areas for future research.

Key words: targeted audiences, targeted sharing, online audience, selective sharing, social networks, online community, Facebook

1 Introduction

Healthy and successful collaborations are fostered through meaningful online interactions [5]. Users of social networks can create favorable collaborative environments by instigating new conversations, encouraging contributions, and advertising and promoting projects [7]. Participation and action can be encouraged via postings in an online community. For example, a user can create posts that invite other community members to view interesting shared content [21]. Perhaps counterintuitively, making posts to a large group does not necessarily increase the number of people that engaged with the posted content. Communication research has found that online users receive *fewer* replies when they share content with their entire network than they would if they share it instead with a small targeted audience [4, 16]. Sociological theory on disclosures also establishes that

people are more likely to be responsive to a request when they feel as though they have been singled out based on an identification of their unique traits [17].

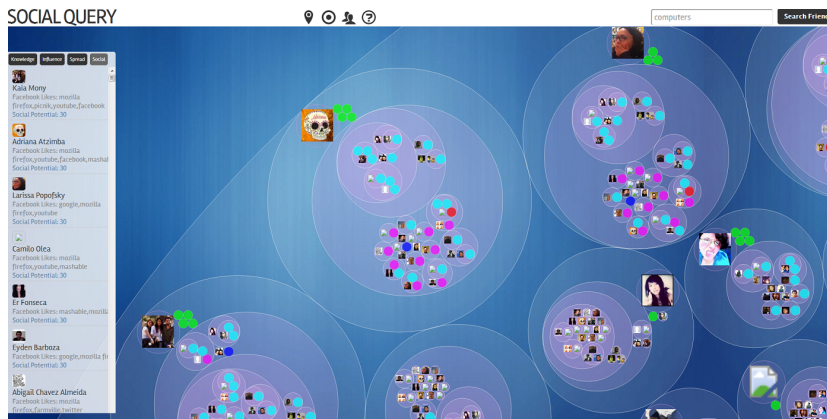


Fig. 1. Screen shots of *Hax*'s social spread interface which lets users view the social groups of their potentially interested audiences.

Many savvy users use different online sharing mechanisms to engage in *selective sharing*, directing content to specific predefined audiences [14]. These users first define collections of people with particular interests, and then post content contextualized so that it is relevant to the interests of the people in each of these collections. However maintaining up-to-date user collections can be difficult and time-consuming. This model is especially unsuitable for more dynamic collections, such as those based on the location, social affiliations, or popularity of the targeted users. For example, the administrators of an online group might want to target only the most influential users in the women's rights movement for promoting their group's cause, or the organizer of a social rally might only want to target those community members who are in town on a particular day. In these cases, predefined collections might be too coarse or include irrelevant users. Another technique involves selecting individuals to target on-the-fly and only sharing the content or message to them. This type of behavior allows for a more dynamic selective sharing experience that is context-driven. We will refer to this practice as *targeted sharing*.

Finding the right people at the right time is hard, especially in larger communities where it is difficult for a single user to keep track of every community member's specific interests and character traits. Previous work in social recommenders has explored the use of list-based interfaces in which a system recommends users with a certain expertise or skill set [6, 13]. These systems do not allow people to easily explore and compare the different characteristics of the recommended individuals. However, these characteristics can play an impor-

tant role when deciding whether or not they should be selected for a particular collaboration or interaction [27].

Interactive visualization tools can enable effective audience targeting by prompting a user to learn about their audience and to understand their different interests. To explore these ideas, we designed *Hax*¹. Hax is a tool that provides a query interface and multiple visualizations to support users in dynamically choosing audiences for their targeted sharing tasks. We study how users engaged with this tool in the context of sharing and connecting with an audience in a Facebook *group*. Facebook designed groups to facilitate online community-building, and we can consider each group to be an online community of its own [1]. Fig. 1 presents a screen shot of one of Hax’s visualizations for targeting audiences within a Facebook group.

The contributions of this work are:

- A novel system for discovering and visualizing the shared interests of an online group or community;
- A novel system for visualizing the spatial-temporal constraints of people;
- A novel system for visualizing the social spread of people;
- A novel system for targeting audiences on-the-fly based on a thematic task or project;
- Providing a better understanding of the way in which data visualizations transform users’ audience selection activities.

2 Motivation

One of the challenges in identifying community members to collaborate or share content with is the fact each person may have dynamic sharing and collaboration needs. For instance, a person might have just gotten a parking ticket and would like to discuss with legal experts ways of fighting the ticket; or a person could want to share a popular news piece she just read with others with like-minded political opinions. Changing events and needs affect who we want to interact with or exchange information with. As a result, social media tools need to offer dynamic mechanisms that let users easily find the people or audience that on-demand can cover their needs.

The data modeling techniques that work for content categorization and information retrieval can be adapted to mine people’s interests and retrieve audiences relevant to users’ diverse needs. But, while specialized data modeling algorithms exist that can correctly categorize data, they rarely fully capture the complex and ever-changing decision-making process for targeting an audience. We therefore opt to integrate data visualizations that incorporate a human-in-the-loop approach.

¹ *Hax* is the Mayan word for *exclusive*, referring to the idea that it is the unique characteristics of a person that are important when selecting him or her for a particular interaction or notification.

We designed different data visualizations that highlight specific traits, or *social signals*, of relevant community members in order to aid users in audience targeting tasks. Our exploration begins with the three social signals listed below. We briefly define the signal and the reasons for considering it. Note that other signals could have been contemplated, but we decided to begin with these as previous work has identified that they play an important role in targeting audiences [7, 28]:

1. **Shared interests:** This signal captures the personal thematic interests of each community member. Many researchers and practitioners view collaborations as a process that aggregates personal interests into collective choices through self-interested bargaining [28]. We believe this bargaining process can be facilitated by making users aware of the personal interests of others, and how they relate to the collaboration task they are promoting.
2. **Location:** This signal holds information about the countries, states, and cities where community members live. Collaborations supported by computers have traditionally provided users with the luxury of interacting with others without having to worry about their location [2]. However, location does play an important role when interacting and organizing events within the physical world [24] (e.g., a social rally) as others' spatial-temporal constraints can determine how much a person will engage in the activity [25].
3. **Social connectivity:** This signal holds information about the type of friends and social ties community members have. This signal is important because it can aid members in recognizing prospective newcomers who can help keep the community alive and active [7]. Additionally, the social connections of a member can also help in the spread of the community's messages and visions. Members could thus use this signal to identify the users whose social connectivity would help them the most in distributing certain content.

3 Background and Related Work

Editors have traditionally made decisions regarding the publishing and distributing of content [10], often relying on the expertise of marketing consultants for particular subjects or audiences. These consultants provided them with a clear picture of who their best audience was for a topic [10]. Via the Internet, anyone can now author, share, and distribute content. But, unlike editors, individual users typically don't have a clear image of their audience [3]. By understanding their audience and adequately targeting it, individuals could better engage their communities [20].

To overcome this lack of marketing knowledge, people rely on cues to estimate the traits of their online audiences. Unfortunately, only a few cues are available [3]. For example, a person might remember she friended her co-workers, and they are thus now in her audience. Without extensive investigation, it might be unclear exactly what these people care about [15]. In this work we explore how we can make audience cues more readily available for people. We study the impact these audience cues can have on a user's audience selection process.

Our tool, Hax, helps users of targeted sharing find a suitable audience for their content. This task is related to expert search in social networks in that the problem is finding a set of contacts that satisfy certain criteria with regard to their knowledge, traits, or social status. Perer et al. [22] present SaNDVis, a tool for visual social network analysis inside of an enterprise that also supports expertise location. In their usage study, they found that their tool helps users find authorities on certain topics, and moreover considers their location. Similarly, ContactMap [29] visualizes contacts along with their attributes and location. Work by Chen et al. [8] uses strong social links as a requirement for finding experts on a topic. Systems that support *social question asking* help users direct questions at people from their social network that are most likely to know an answer [6, 9, 18, 20].

In summary, these works show interesting parallels to understanding and supporting targeted sharing. However, they focus either on user goals or audience characteristics that are distinctly different from those of targeted sharing.

3.1 Facebook Graph Search

Facebook’s *Graph Search*² offers a natural language interface for searching one’s social network; queries may consider several social variables. For instance, a typical query might be: “*TV shows liked by people who study computer science.*” A query returns a ranked list of relevant Facebook users with some of their characteristics included, such as the city where they live, the music they like, how many friends they have on the site, among others. However, it is unclear if the design of Graph Search was influenced by the requirements of targeted sharing. The attributes and interactions modes it supports are limited. The task specificity and the richer interaction modes of the tool presented in this work aim to make it more useful and accessible for targeted sharing tasks.

3.2 Facebook Advertisement Targeting Options

Facebook offers advertisers options for ensuring that their ad will reach a targeted relevant audience.³ Advertisers can target audiences based on users’ location, age, zodiacal sign, interest, education, their friends, as well as whether they have liked their particular product in the past. Facebook’s targeting options assume that the end-user has a clear image of who their desired audience is. While this design consideration can be effectively true for advertisers who have previously conducted market studies and identified the demographics of their clients, it is not necessarily valid for individual community members who engage in targeted sharing.

Bernstein et al. [3] identified that social media users consistently underestimate the audience size for their posts, guessing that their audience is just 27%

² <https://www.facebook.com/about/graphsearch>, accessed February 10, 2014

³ <https://www.facebook.com/help/www/131834970288134?rdrhc>, accessed February 10, 2014

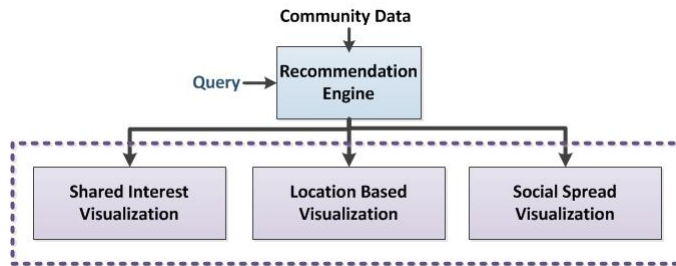


Fig. 2. The components of the *Hax* system.

of its true size. It is therefore likely, that community members also will not have an exact idea regarding the characteristics and traits of their most relevant audience for a given post. The creation of online tools could be useful in helping end-users better visualize and understand potential audiences and their different characteristics.

4 Designing Hax

Hax is a web-based tool that supports targeted sharing on Facebook via a query that indicates the topic they are interested in posting content about. Hax includes a recommendation engine that accepts and processes such queries to produce a list of relevant community members based on their *likes*. For each returned member, the recommendation engine includes their signals (e.g., their *likes*, hometown, or number of friends) and a weighting. The visualization engine provides three different visual presentations emphasizing different aspects of the recommendations. Fig. 2 presents an overview of the Hax components.

4.1 Recommendation Engine

The recommendation engine models the interests of community members based on their profile information. It then identifies those members whose interests are the most relevant to a user’s search query. We model the general interests of community members through their Facebook *likes*. A Facebook *like* typically has a name, a label, and a definition. For example, the *like* “*Everyday Feminism*” has the name “*Everyday Feminism*”, the label “*Community Organization*”, and the definition “*Everyday Feminism strives to stop the everyday violence, dominance, and silencing used against women*”. We found that the curated labeling used by Facebook to categorize interests is very general, and does not enable an easy way to explore the data further. To counter this effect, we use topic models [23] to model the community’s shared interests.

Given the nature of the data, we used a labeled Latent Dirichlet Allocation approach (labeled LDA) [23], similar to that proposed by Forbes et al. [12]. The discovered LDA topics correspond to the community’s shared interests, and labels correspond to Facebook *likes*, and each document corresponds to

a *like* with its definition. Specifically, we use a generative process to discover the interests shared by the community members. The process first detects the K number of unique labels associated to the community’s *likes*. This sets the initial number of shared interests that will be considered. For each shared interest, a unique Like and its associated data is drawn with a Dirichlet distribution α . A multinomial mixture distribution θ^d over all K shared interests is drawn for each community member with a Dirichlet prior α_ϕ . Now, using information about the labels associated with the *likes* of the user, we restrict the definition of θ^d to be defined only to the shared interest associated with the labels present in their *likes*. After this step, each community member is represented as a mixture over shared interests. An end user’s query is also modeled as a mixture over shared interests, except that, because it does not have any explicit labels, θ^d is not restricted. The community members who exhibit a shared interest mixture similar to that of the query are presented to the user via the interactive visualizations. We use the L_1 norm as our similarity metric. Our experimental experience, as well as related work in modeling micro blog conversations and users via topic models, suggest that using topic models to mine a community’s shared interests is a feasible approach [19].

Given a search query, the recommendation engine first identifies the community’s shared interest most relevant to the query. It then finds the community members that have Facebook *likes* most relevant to the query, weighting each of them based on their number of relevant *likes*. This list of weighted members and most relevant shared interests is then used as input for the visualization engine.

4.2 Visualization Engine

The visualization engine displays the list of recommended members with their weighted social signals. This allows users to consider these signals directly in her targeted sharing decision process. Hax provides three different interactive visualizations, each emphasizing different social signals. Following the visualization mantra [26], every visualization lets the user (a) obtain an overview of the community’s social signals; (b) zoom into particular groups of members; and (c) obtain details of a desired user’s social signals. This rich interaction is not possible with a list-based interface. List-based interfaces do not allow the user to easily obtain overviews and summaries of the data. Given that community users are often organizing things for the entire community, providing overviews of the members’ interests can help users remain relevant. Tooltips could potentially be used for offering these data summaries. However, this is not sufficient as it does not allow users to zoom in and explore particular aspects of the data.

We provide a short description of each view below. Fig. 3 presents the type of overviews each interface provides. Fig. 4 shows example screenshots of the location-based, shared interest, and social spread interfaces.

Shared Interest Interface Initially, the *shared interest interface* presents an overview of all of the discovered shared interests of the community (Fig. 3.1).

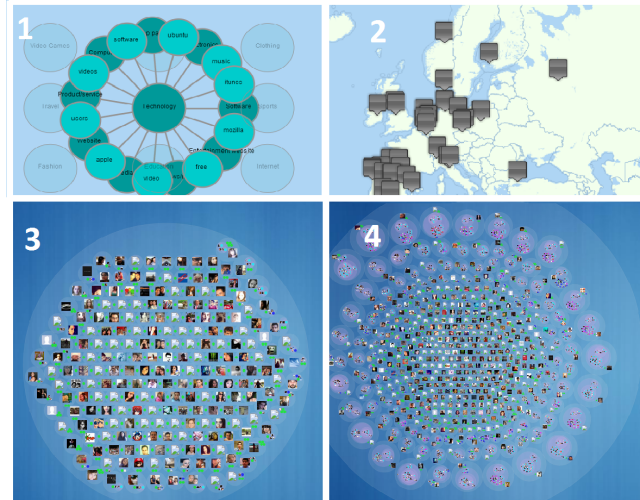


Fig. 3. Overviews given by each visualization: (1) shared interest, (2) location-based interfaces; (3, 4) social spread interface.

Shared interests are displayed as nodes on a grid. Each node has in its center the keyword most representative of the shared interest. Mousing over a shared interest displays in light green its most representative keywords, and in dark green its most representative Facebook labels. This view allows users to quickly identify the general interests of their community, as well as some of the most popular specific related interests.

When the user queries the system, a list of relevant members is displayed along with the community's shared interest topic most correlated to that query (Fig. 4, middle). Relevant members are visualized as a list of nodes on the right hand side of the interface. A large node in the center represents the most relevant shared interest topic; other shared interests are shown on the left for reference. Mousing over a member or a shared interest provides more information, e.g. the *likes* of a member that correlate to the query, the description of a *like*, or the Facebook labels associated with a shared interest.

The shared interest interface thus allows a user to quickly see the members that are likely to be interested or knowledgeable about a particular shared interest related to the query, and the user can easily investigate the connections between members, *likes*, labels, and shared interests.

Location-Based Interface The *location-based interface* lets users visualize the geographical locations of the members relevant to their search query. This information can be important when targeting members for activities that take place in the physical world, such as meetings, events, or rallies. In addition, location also provides a sense of cultural context.

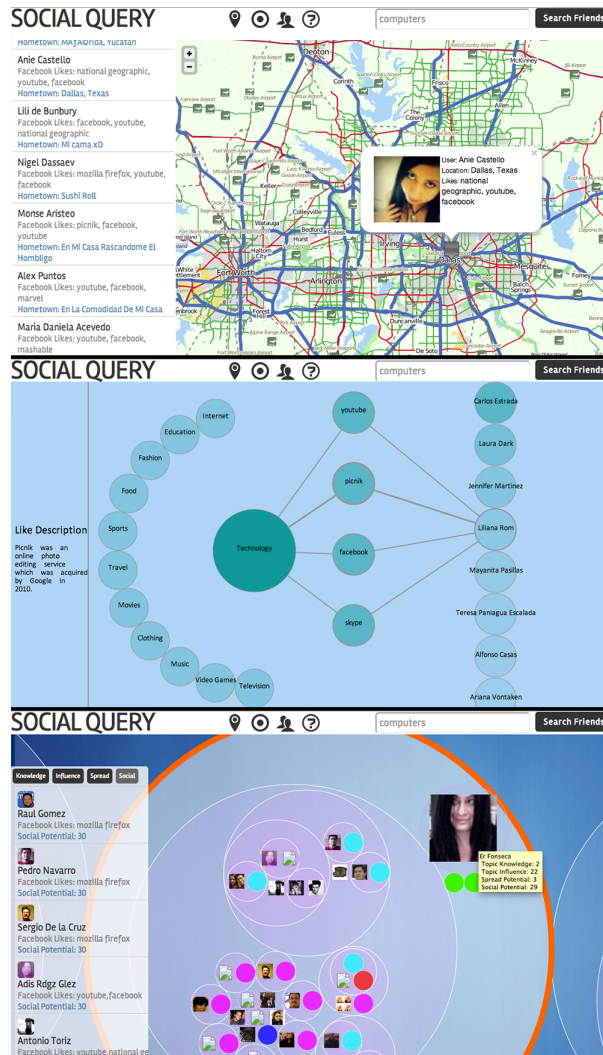


Fig. 4. Screen shots of a zoomed-in version of the different visualizations in *Hax* (top to bottom): location-based, shared interest, and social spread interfaces.

The interface shows recommended members on a geographic map, based on the city or place the member listed in their profile. At a first glance, the interface allows users to easily identify the geographical regions where the majority of the members interested in a particular topic reside (Fig. 3.2). Users can also zoom in on any member, which will show a list of their relevant Likes, their profile photo, and a more detailed map of the area (Fig. 4, top). Since not every member lists their location, this interface only includes recommended members who have shared this information.

Facebook’s targeting options for brands offers a filtering based on location. It is assumed that end-users have a good notion of the cities where their targeted audience live. However, given that users may share diverse and dynamic content with their group, it can be difficult for them to have a clear picture upfront of who their most relevant audience members are, or where they live.

We argue that location-based interfaces for targeting of audiences should allow users to obtain overviews of where their audiences are physically located, and then enable end-users to further explore the map on multiple levels. This enables users to consider community members’ different physical affordances [25] in their decision process. Knowing others’ physical affordances is important as it can influence their decisions for participating in an event [25].

Social Spread Interface The *social spread interface* helps users identify the members with interests related to their query who at the same time have the most contacts or friends with relevant interests. This interface finds members that are not just potentially interested in certain content, but rather potentially interested members whose connections help them distribute or “spread” content to large audience. These are the people who bring value to the content, not necessarily by the comments they provide to the content, but rather through the links to their social contacts.

The social spread interface receives the list of recommended community members from the recommendation engine. For each member, the recommendation engine includes a list of her Facebook *likes* relevant to the user’s query and a list of the member’s Facebook friends who also have relevant Facebook *likes*. The visualization first structures the members based on their amount of relevant social connections. Members are structured in a spiral form (cf. Fig. 3.4). The outer rings of the spiral present the members who have the most friends with the most interests related to the user’s query. The center of the spiral contains the members who have the least friends with relevant interests. When all interested members exhibit approximately the same number of interested social contacts, members are arranged in a planar circle from left to right, top to bottom, based on their amount of relevant Likes (cf. Fig. 33).

Each node in the spiral or circle represents a community member. Each member is presented with their relevant *likes*, photo, and relevant contacts. Each of these contacts is displayed with their own relevant *likes* and photo. Contacts are grouped and color-coded based on the *likes* they have in common with the community member and their relationship with the community itself. The more *likes* a community member has in common with a contact, the closer they both appear in the interface. Contacts with a light blue circle next to them are contacts that have no other connection with the community than their friendship to that particular member. Dark blue circles denote contacts that have one or more other friends who are also community members. Purple circles denote contacts that have friends who are friends with community members.

This view allows users to quickly identify the overall type of social connections that the community reveals for different topics. They can also zoom in and inspect particular members and their relevant social contacts. This enables end-

users to easily adjust their messages (and who they mention) to content that can have a larger reach and impact. It also allows users to share content with members whose social contacts could be supportive to their cause.

The spiral structure of the visualization was inspired by the work of Katayoon et al. [11]. Their research found that complex visualizations of hierarchical data can become overcrowded and thus makes it difficult to see details about specific nodes. Their work thus proposed layouts focused on a node of interest that make use of phyllotactic patterns (spirals) via nested circles that are centered on the node of interest. This type of layout is designed to provide more space than traditional hierarchic visualizations. Space-saving designs become important given the overwhelming amount of possible members of an online community and the large amount of relevant contacts each member can have.




Fig. 5. Hax at a university annual open exhibition which had hundreds of visitors.

5 Usability Inspection of Hax

We conducted a survey study with users who utilized Hax as a tool to find relevant audiences for different content sharing tasks. We questioned participants about their experiences using Hax. We used qualitative coding to create a taxonomy of experiences that emerge from using data visualizations to target audiences. For our study, we worked closely with members of a specific Facebook group for which we were able to recruit participants.

5.1 Participants

Using the Facebook *group browser*,⁴ we first identified groups with large number of members and then asked the group administrators whether their group would be interested in participating. We contacted the administrators from 10 different groups who we determined, based on information posted on their public Facebook profile, had a large percentage of members that were local to where we planned to carry out the study.

⁴ <http://www.facebook.com/search.php?type=groups&q=%22keyword%22>, accessed February 10, 2014

One group accepted the invitation: an activist group organizing social initiatives around the world. Its 2,000+ active members are distributed world-wide. The group covers a wide range of discussions and events, ranging from the philosophy of free software to the coordination of wildlife preservation rallies. We were granted access to the public Facebook profile of all its members. From this data, our system automatically discovered the groups' interests, and produced the three different data visualizations. 15 of the group members agreed to participate in our evaluation: 2 female, 13 male, 4 long-term group members and 11 newcomers (less than one month in the group.) They ranged in age from 19 to 35. Participants came to our laboratory for the study, and received \$10 USD for their time.

5.2 Procedure

Over the course of an hour session, each participant completed a series of targeted sharing tasks using Hax on a internet-connected laptop that we provided. We decided that participants would conduct tasks with Hax only, a not in comparison to Facebook's native interface, as Facebook is not particularly designed or tailored for the specific usage of finding relevant online audiences. However, participants were asked to reflect about the benefits and drawbacks of our data visualizations and traditional list-based interfaces. We used qualitative coding based on ground theory for our analysis.

In each task, participants were told to identify 10 candidates for targeted sharing. Each participant was given 15 different tasks that we statistically varied using a Latin square design. Each task came from 5 different scenarios that represented a few of the group's audience targeting needs. Group members not taking part in the evaluation helped edit the tasks and scenarios to reflect real needs. The five scenarios were: 1) Find audiences interested in a certain thematic post; 2) Find audiences to invite to a thematic event, and who are likely to attend; 3) Find audiences to help distribute a thematic article and get others to read it; 4) Find audiences who could help spread news about a thematic event and get others involved; 5) Find audiences who could start a discussion with the group on a certain topic.

As participants performed the tasks, they were observed by one of the researchers who took notes. After participants completed all tasks, they were asked to complete a questionnaire about their experience with Hax, strategies they adopted to complete the tasks, benefits and drawbacks they saw, and a comparison between Hax and list-based interfaces. The questionnaire is available online.⁵ Two of the authors coded the responses by reading every questionnaire response and identifying key concepts about users' perspectives on using data visualizations to target audiences. Following grounded theory's coding criteria, we decided that a category would cover a general type of experience that emerges from using data visualizations to target audiences. A total of 4 main categories were identified by this process.

⁵ <http://www.surveymonkey.com/s/KN5CGF>, accessed February 10, 2014

5.3 Results

All participants were able to use Hax to complete all of the tasks assigned to them. Below we discuss each of the 4 categories that emerged from using data visualizations to target audiences. For some of the categories we provide quotes from the questionnaires to help illustrate the core of the category.

Serendipitous Discoveries This experience is about feeling that data visualizations help one make discoveries about one’s targeted audience. All participants reported that Hax prompted them to discover and learn new things about particular group members, and the group in general, something they felt was not facilitated with traditional list-based interfaces: “...*It was really neat to learn so easy and fast what everyone is into. I never experienced that with Facebook.*” Many participants mentioned out loud some of the new discoveries they made with Hax. Additionally, we observed that some started using Hax for their own personal explorations. Dynamic audience visualizations engage users and facilitate serendipitous discoveries of their social groups. This could help people share better content because they understand their audience more.

Visualizing Diffusion and Participation This experience is about considering data visualizations to be helpful in finding large pools of people likely to take action in regard to a message, e.g., comment, or attend an invitation. 70% of participants found Hax useful for distributing content to audiences who would be engaged with the content. Participants felt list-based targeting tools did not provide such perspective. Participants believed the location-based visualization facilitated finding audiences from big cities who could easily spread messages to large pools of actionable people, e.g., by making announcements on the streets about an event people could walk to. Participants also felt that by visualizing social connections and interests they could distribute content to mass audiences likely of participating in collaborative action afterwards, such as a discussion. Additionally, the location-based interface helped participants make a connection between the virtual event on Facebook and participation in the physical world, especially selecting an audience who could travel and attend: “*The map really made me think about the actual event, and like really including the person.*” It is interesting to observe how just having a map helped people integrate location in their audience decision process. Our results hint there is value in designing systems that enable users to visualize and explore others’ spatial affordances. This signal could provide the perspective needed to make online interactions more realistic, especially compared to list-based interfaces that provide little spatial context.

Audience Diversity This experience is about feeling that data visualizations bring diversity to one’s targeted audience selection process. Participants reported that the shared interest visualization helped them find relevant candidates who had different perspectives. Participants also mentioned that the location-based

interface let them have more diverse selections: *“I tried to have diversity in who I selected. People who like the same things or are from the same town will have same interests and maybe not that much new to add.”*

Audience Verification This experience is about using data visualizations to verify the recommended audiences. 10% of all participants reported this experience. Participants especially used the shared interest interface to figure out the meaning of the *likes* and to analyze whether it made sense to include certain candidates in their targeted audience: *“There were some brands [i.e., likes] that I didn’t know, but the knowledge interface [i.e., shared interest visualization] helped me know what they were about.”* Participants particularly enjoyed not having to leave the tool to comprehend the audience that our system recommended.

5.4 Open Deployment of Hax

Additionally, Hax has been installed on a large screen display for several hours in a well-attended university exhibition to further explore how average users experience this open-ended way of selecting audience candidates (cf. Fig. 5.) Even without prior notice or instruction visitors to the exhibition were able to approach the display and begin interacting with Hax. During the deployment, approximately 150 visitors approached Hax: around 70 visitors interacted with Hax while the rest analyzed and studied Hax without interacting. Average interactions times were around 1 min.

6 Outlook and Discussion

Our results show that users can target their audiences through interactive data visualizations. Our data visualizations prompted users to learn more about their peers. They also helped people find diverse audiences for their different sharing tasks, something not facilitated by list based interfaces. This type of system design can help users to have more cultural sensitivity and to foster better social interactions and collaborations. Hax empowers users to consider not only others’ interests, but also other traits, such as social, cultural, and spatial signals. This creates a more compelling sharing experience. We believe there is value in designing systems focused on the visualization of people’s traits. Such systems could facilitate serendipitous discoveries and encourage diverse interactions. It is important to think about creating digital opportunities where strangers with different opinions can find each other and connect. Social media data mixed with data mining and visualization techniques provide a unique opportunity for giving users diversity. Our results encourage future studies that address audience understanding as a main visualization goal.

Acknowledgements: This work was partially supported by a UC MEXUS-CONACYT fellowship, by the U.S. Army Research Laboratory under Cooperative Agreement No. W911NF-09-2-0053 and by NSF grant IIS-1058132. Special thanks to our participants

and the anonymous reviewers whose thoughtful feedback helped improve the presentation of this work.

References

1. Askanase, D. Returning to facebook groups, May 2013. <http://communityorganizer20.com/2012/05/31/returning-to-facebook-groups/>.
2. Bentley, R., Hughes, J. A., Randall, D., Rodden, T., Sawyer, P., Shapiro, D., and Sommerville, I. Ethnographically-informed systems design for air traffic control. In *Proceedings of the 1992 ACM conference on Computer-supported cooperative work, CSCW '92*, ACM (New York, NY, USA, 1992), 123–129.
3. Bernstein, M. S., Bakshy, E., Burke, M., and Karrer, B. Quantifying the invisible audience in social networks. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems, CHI '13*, ACM (2013), 21–30.
4. Bernstein, M. S., Marcus, A., Karger, D. R., and Miller, R. C. Enhancing directed content sharing on the web. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems, CHI '10*, ACM (2010), 971–980.
5. Brennan, K., Monroy-Hernandez, A., and Resnick, M. Making projects, making friends: Online community as catalyst for interactive media creation. *New Directions for Youth Development 2010*, 128 (2010), 75–83.
6. Brzozowski, M. J., and Romero, D. M. Who should i follow? recommending people in directed social networks. In *ICWSM* (2011).
7. Butler, B., Sproull, L., Kiesler, S., and Kraut, R. Community effort in online groups: Who does the work and why?, 2001.
8. Chen, D., Tang, J., Li, J., and Zhou, L. Discovering the staring people from social networks. In *Proc. of the 18th Int. Conf. on the World Wide Web*, ACM (2009), 1219–1220.
9. Chen, J., Geyer, W., Dugan, C., Muller, M., and Guy, I. Make new friends, but keep the old: recommending people on social networking sites. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems, CHI '09*, ACM (2009), 201–210.
10. Cheyne, T. L., and Ritter, F. E. Targeting audiences on the internet. *Communications of the ACM* 44, 4 (2001), 94–98.
11. Etemad, K., and Carpendale, S. Symmetry and node focused visualization of large trees. In *GRAND NCE Annual Conf.* (2010).
12. Forbes, A. G., Savage, S., and Höllerer, T. Visualizing and verifying directed social queries. In *IEEE Workshop on Interactive Visual Text Analytics* (October 2012).
13. Hu, Y., Farnham, S. D., and Monroy-Hernández, A. Who.ly: facilitating information seeking for hyperlocal communities using social media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '13*, ACM (New York, NY, USA, 2013), 3481–3490.
14. Kairam, S., Brzozowski, M., Huffaker, D., and Chi, E. Talking in circles: selective sharing in google+. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems, CHI '12*, ACM (2012), 1065–1074.
15. Lampinen, A., Lehtinen, V., Lehmuskallio, A., and Tamminen, S. We're in it together: Interpersonal management of disclosure in social network services. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11*, ACM (New York, NY, USA, 2011), 3217–3226.

16. Latané, B., and Darley, J. Group Inhibition of Bystander Intervention in Emergencies. *Journal of Personality and Social Psychology* 10, 3 (1968), 215–221.
17. McCroskey, J. C., and McCain, T. A. The measurement of interpersonal attraction. *Speech Monographs* 41, 3 (1974), 261–266.
18. McDonald, D. W., and Ackerman, M. S. Expertise recommender: a flexible recommendation system and architecture. In *Proc. of the 2000 ACM Conf. on Computer-Supported Cooperative Work*, ACM (2000), 231–240.
19. McKenzie, G., Adams, B., and Janowicz, K. A thematic approach to user similarity built on geosocial check-ins. In *Association of Geographic Information Laboratories in Europe*. Springer Int. Publishing, 2013, 39–53.
20. Nichols, J., Zhou, M., Yang, H., Kang, J.-H., and Sun, X. H. Analyzing the quality of information solicited from targeted strangers on social media. In *Proc. of the 2013 Conf. on Computer-Supported Cooperative Work*, ACM (2013), 967–976.
21. Otto, P., and Simon, M. Dynamic perspectives on social characteristics and sustainability in online community networks. *System Dynamics Review* 24, 3 (2008), 321–347.
22. Perer, A., Guy, I., Uziel, E., Ronen, I., and Jacovi, M. Visual social network analytics for relationship discovery in the enterprise. In *2011 IEEE Conf. on Visual Analytics Science and Technology (VAST)*, IEEE (2011), 71–79.
23. Ramage, D., Hall, D., Nallapati, R., and Manning, C. D. Labeled lda: a supervised topic model for credit attribution in multi-labeled corpora. In *Proc. of the 2009 Conf. on Empirical Methods in Natural Language Processing*, vol. 1 of *EMNLP '09*, Association for Computational Linguistics (2009), 248–256.
24. Raubal, M., Miller, H., and Bridwell, S. User Centered Time Geography For Location-Based Services. *Geografiska Annaler B* (2004).
25. Savage, N. S., Baranski, M., Chavez, N. E., and Höllerer, T. I'm feeling loco: A location based context aware recommendation system. *Advances in Location-based Services: 8th Int. Symposium on Location-based Services, Vienna 2011* (2012), 37.
26. Shneiderman, B. The eyes have it: A task by data type taxonomy for information visualizations. In *Proc. of the 1996 IEEE Symposium on Visual Languages, VL '96*, IEEE Computer Society (1996).
27. Singer, L., Filho, F. M. F., Cleary, B., Treude, C., Storey, M.-A. D., and Schneider, K. Mutual assessment in the social programmer ecosystem: an empirical investigation of developer profile aggregators. In *Proc. of the 2013 Conf. on Computer-Supported Cooperative Work* (2013), 103–116.
28. Thomson, A. M., and Perry, J. L. Collaboration processes: Inside the black box. *Public Administration Review* 66 (2006), 20–32.
29. Whittaker, S., Jones, Q., Nardi, B., Creech, M., Terveen, L., Isaacs, E., and Hainsworth, J. Contactmap: Organizing communication in a social desktop. *ACM Transactions on Computer-Human Interaction (TOCHI)* 11, 4 (2004), 445–471.