University of California
Santa Barbara

Understanding and Engaging Online Audiences

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Computer Science

by

Norma Saiph Savage

Committee in charge:

Professor Tobias Höllerer, Chair
Professor Matthew Turk
Professor Xifeng Yan
Dr. Andrés Monroy-Hernández

September 2015

The Dissertation of Norma Saiph Savage is approved.

Professor Matthew Turk

Professor Xifeng Yan

Dr. Andrés Monroy-Hernández

Professor Tobias Höllerer, Committee Chair

June 2015

Understanding and Engaging Online Audiences

"If you want something from an audience, you give blood to their fantasies. It's the ultimate hustle."
—Marlon Brando

Acknowledgements

This thesis is dedicated to an indispensable audience:

- My family, you are always there for me. Thank you, I love you and appreciate you.
- My advisor Tobias Höllerer, always open to crazy ideas, and had the patience to see me grow as a researcher, and now professor. Vielen Dank für die schönen Stunden!
- My mentor and role model Andrés Monroy-Hernández who always had time to guide me to make great work in social computing, and who showed me how I could use the love for my country to make great research.
- My PhD heroes Matthew Turk, Elizabeth Belding, Michael Bernstein, Angus Forbes, Xifeng Yan, all with whom I had the wonderful opportunity of collaborating and learning their great work etiquettes and research styles.
- Stalin Munoz, Vicente Fuentes Gea, Ron S. Leder, Ana Kuri, who got me into computer science and fostered my love for research.
- All of my collaborators: Nataly Moreno, Maciej Baranski, Norma Elva Chavez, Syed Reza Ali, Rodrigo Savage, Shoji Nishimura,Carlos Toxtli, Grant McKenzie, Shloka Desai, Dulce Tania Nava, Salvador Medina, David Alvarez, Wendy Chun, Walter Angel, Anamary Leal, Jessica Cornick, Leif Singer, Derick Johnson, Bruce Henry, Berney Keane, David Cobbley, Rita Wouhaybi, Aaron Davis, Lama Nachman, Giuseppe Raffa, Saurav Sahay, Domagoj, Daniel Vaquero, Svetlin Bostandjiev, John O'Donovan, James Schaffer, Emily Fujimoto, Basak Alper, Byungkyu Jay Kang, Jungah Son, Tibor Goldschwendt, Julio Cruz, Hugo Leon, Alf Noriega, Greg Meyer, Nedda Amini, Johann ly, Luis Rodriguez. "Our whole message is that we are more powerful together" Thank you!

- The community and friends in UCSB's Four Eyes Lab and Intel.

- The wonderful people who helped make #FixIT a reality, and are transforming who creates and produces technology: Jesus Savage, Mariel Garcia, Juan Pablo Flores, Indira Cornelio, Deana Brown, Erika Smith, Alejandra Monroy, Cesar Islas, Ariadna Gomez, Laura Ballesteros, Cindy Rubio, Miray Kas, Rohan Lamba, Sarah Safir, Jenine Beekhuyzen, Carlos Gonzalez, Alvaro Norma, Paulina Bustos, Alan Bautista, Laura Lopez, Susana Hildalgo, Angel Perez, Jesus Perez, Luis Eduardo Aorve Vidal, Francisco Dorantes.

- My amazing collaborators from Crowd Research, you gave me the motivation to finish this last stretch. Thank you: Jeerel Herrejon, Veronica Jaramillo, Claudia Saviaga, Luis Daniel Palacios, Diana Everly, Will Dai, Gema Toledo, Martin Daniel Godinez, Beto Avila, Nisha KK, Dilrukshi Gamage, Mohammed Kambal, spamGirl, Ryan Compte, Adam Ginzberg, Tejas, Angela Richmond, Juechi, Shirish Goya, Rijol, Rahul, Asmintha and many more!

- My great friends! *"If you want to do really important things in life and big things in life, you can't do anything by yourself. And your best teams are your friends and your siblings."*–Deepak Chopra
  Gracias Rodrigo, Dekaru, Asaf Paris, Judith Cordero, Carly Powne, Anaid Guevara, Bety, Joshua Cardiel, Sergio Rodriguez, Miguel y Anylu Lastas, Daniel Valdez, Rojo Cordova, Veronika Stranova, Leon Chavez, Pablo Diaz, Fernanda Lugo, Ulises Matus, Ilona Gelyana, Isla Vista Collective.

- I was supported by Conacyt-UCMexus Fellowship, a Google Anita Borg Fellowship, and a generous advisor. You all make graduate school possible — thanks

Curriculum Vitæ

Norma Saiph Savage

Education

2015	Ph.D. in Computer Science University of California, Santa Barbara.
2015	M.A. in Computer Science, University of California, Santa Barbara.

Work Experience

UCSB 4 Eyes Laboratory, Santa Barbara, CA — Sept 2014 to present

Research Assistant, Imaging, Interaction, and Innovative Interfaces Group
Intel Labs, Hillsboro, OR — September 2011 to September 2013

Research Intern, Interaction and Experience Research (IXR) Group
Gifiniti, Santa Barbara, CA — January 2012 to June 2012

Co-founder & Research Scientist
Intel, Hillsboro, OR — June 2011 to September 2011

Graduate Student Technical Intern, Software Pathfinding Group

Publications

Savage, S., and Monroy-Hernandez, A., Participatory Militias: An Analysis of an Armed Movement's Online Audience. *CSCW 2015: ACM Conference on Computer-Supported Cooperative Work.* 28% acceptance rate.

Savage, S., Monroy-Hernandez, A., Hollerer, T., Tag Me Maybe: Perceptions of Public Targeted Sharing on Facebook. *HT 2015: Hypertext and Social Media.*

Savage, S., Forbes, A.G., Toxtli, C., McKenzie, G., Desai, S., and Hollerer, T., Visualizing Targeted Audiences. *COOP 2014: Conference on the Design of Cooperative Systems.* Springer.

Savage, S., and Monroy-Hernandez, A., The "Courage For" Facebook Pages: Advocacy Citizen Journalism in the Wild. *C+J 2014: Computation+Journalism Symposium.* American Journalism Review.

Chavez, N.E., Savage, R., Nava, T., and **Savage, S.**, Enchantment Under the Sea: An Intelligent Environment for User Friendly Music Mixing *IE 2012: IEEE International Conference on Intelligent Environments.*

Savage, S., Baranski, M., Chavez, N.E., and Hollerer, T., I'm Feeling LoCo: A Location Based Context Aware Recommendation System. *LBS 2011: International Symposium on Location-Based Services.* Lecture Notes in Geoinformation and Cartography. Springer.

Savage, S., Chun, W., Chavez, N.E., and Hollerer, T., Seems Familiar: An Algorithm For Automatically Inferring Spatial Familiarity. *LBS 2011: International Symposium on Location-Based Services.* Lecture Notes in Geoinformation and Cartography. Springer.

Savage, R., Nava, T., Chavez, N.E., and **Savage, S.**, Search on the Cloud File System. *PDCS 2011: IASTED International Conference on Parallel and Distributed Computing and Systems.*

Chavez, N.E., **Savage, S.**, and Savage, R., Right place at the Right Time: Novel Place Recommendation System. *CIINDET 2011: IEEE Congreso Internacional sobre Innovacin y Desarrollo Tecnolgico.*

Savage, S., Ali, R., Chavez, N.E., Savage, R., Mmmmm: A multi-modal mobile music mixer *NIME 2010: Conference on New Interfaces for Musical Expression.*

Leder, R.S., Azcarate, G., Savage, R., **Savage, S.**, Sucar, E., Reinkensmeyer, D., Toxtli, C., Roth, E., Molina, A., Nintendo Wiimote for Arm and Wrist Therapy in Stroke Survivors with Upper Extremity Hemipariesis. *IEEE Virtual Rehabilitation 2008.*

Savage, S., Nachman, L., Sahay, S., Raffa, G., Socially and Contextually Appropriate Recommendation Systems *U.S. Patent Disclosure, filed 2014.*

Savage, S., Wouhayb, R., Nachman, L., Sahay, S., Online Social Persona Management *U.S. Patent Disclosure, filed 2013.*

Savage, S., Wouhayb, R., Traversing Data Using Data Relationships, *U.S. Patent Disclosure, filed 2012, published 2013.*

Leon, H., Cruz, J., **Savage, S.,** Chavez, N.E., FiberKraft: An Interface for Designing Optical Fiber Networks. *ENC 2014: Encuentro Nacional de Ciencias de la Computacion, Posters.*

Desai, S., Rodriguez, L., **Savage, S.,** Chavez, N.E., Forbes, A.G., Crowdsourcing Volunteers. *SoCal CWIC 2014: Celebration of Women in Computing in Southern California, Posters.*

Savage, S., Chavez, N.E., Toxtli, C., Medina, S., Alvarez, D., Hollerer, T., A Social Crowd-Controlled Orchestra, *CSCW 2013: ACM Conference on Computer-Supported Cooperative Work, Posters.*

Savage, S., Forbes, A.G., Savage, R., Hollerer, T., Chavez., N.E., Directed Social Queries with Transparent User Models *UIST 2012: ACM Symposium on User Interface Software and Technology, Posters.*

Savage, S., Contextually Aware Recommender Systems. *CRA-W 2011: Committee on the Status of Women in Computing Research Cohort, Posters.*

Savage, S., Forbes, A.G., Savage, R., Hollerer, T., Chavez., N.E., Utilizing Crowdsourced Databases for Social Media Question Asking, *CSCW 2013: Workshop on Social Media Question Asking at ACM Conference on Computer-Supported Cooperative Work.*

Forbes, A.G., **Savage, S.**, Hollerer, T., Visualizing and Verifying Directed Social Queries *InfoVis 2012: IEEE Workshop on Interactive Visual Text Analytics at IEEE Transactions on Visualization and Computer Graphics.*

Savage, S., Nishimura, S., Chavez, N.E., and Yan, X., Frequent trajectory mining on GPS data., *LocWeb 2010: ACM Workshop on Location and the Web, at the Internet of Things.*

Abstract

Understanding and Engaging Online Audiences

by

Norma Saiph Savage

Social media has turned all of us into potential authors of content. This phenomenon has further facilitated the formation of new dynamic audiences – all of whom center on the data we share. Although there have been several related analyses, most research assumes that the online audience is only an observer. This has led to the design of platforms that are adaptations of traditional marketing tools wherein audiences are surveyed and categorized anonymously and content authors have minimal interaction with them. The types of collaborations produced by such tools are limited.

This research recognizes that the internet has transformed how authors and audiences operate. The thesis explores the dynamics of this emerging ecosystem, from authors, who share personal content with friends and family, to citizen reporters who collaborate with audiences to oppose drug cartels. The thesis demonstrates how to incorporate the understanding of these dynamics into the design of novel platforms. The thesis does this via individual case stories of such systems, for instance the prototype system Hax, which dynamically allows people to visualize relevant audiences for sharing and collaborating, or the tool Botivist, which dynamically recruits and assembles collective efforts with online audiences.

The thesis discusses how, together, we can create a future where platforms produce a true symbiosis between authors and audiences to facilitate collective efforts.

Contents

Chapter 1

Introduction

Traditionally, a select elite was who defined the content consumed by the majority of the population [1, 2]. These elites usually came from hierarchical organizations, and focused on generating content that followed a rigorous editorial workflow, had integrity, and was profitable. News papers, TV shows, magazines, all followed this model. Figure 1.1.A presents an overview of this traditional model.

However, the Internet, and social media in particular, have flourished a new media consumption model [2]. This new media model is much more participatory in nature, as it offers low barriers for sharing one's creations. Everyone now has the power to produce content, and have audiences forming around every piece of information they share. Audiences have also the convenience of obtaining information from diverse sources and accessing different perspectives [3]. This new medium, has created new networked communities, that both produce and consume content. Content producers, such as journalists, are no longer simply lecturers who inform others of what is taking place in the world. Journalists are assuming new roles where they can act as "forum leaders," mediators. Similarly, the audience has now a hybrid role where they act not only as consumers, but also as producers. Figure 1.1.B presents an overview of this new model.

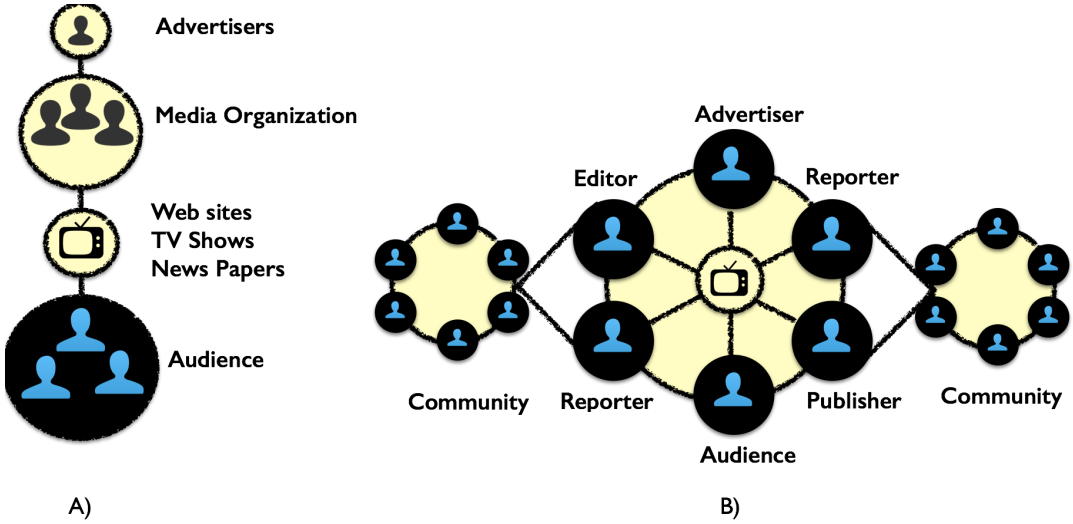

Figure 1.1: Overview of the traditional (A) and new media models (B).

This novel media model is creating new experiences, relationships, and tensions between content producers and audiences. However, we lack a detailed understanding of these emerging dynamics. This becomes especially problematic when designing tools for engaging with online audiences. The majority of these online tools have to make assumptions of how content producers and audiences interact. These tools usually just adopt the relationships that existed in the traditional model. For instance, many tools for targeting audiences assume that end-users have a clear picture of their public, and ask the person to provide the audience upfront (e.g., the traits of the audience they want to target.) While this design might be effective for advertisers who have previously conducted market studies and identified the demographics of their clients, it may not be useful for the average person, who has ever-changing interests and sharing needs. It can be difficult for the person to have a good grasp of who is her best audience for each sharing task.

Additionally, most of these tools follow a model focused on profit. Their designs mirror marketing or revenue tools. Yet, in this participatory new medium profit might not be the main priority, but rather conversation and collaboration. Our lack of knowledge

on these emerging relationships dynamics translates into tools which might not cover the true needs of content producers or online audiences.

This thesis studies how content producers and audiences interact in social media. I use the findings of my studies to create better online communication platforms, and root design principles for the next generation of online audience engagement tools.

To facilitate understanding the nature of this thesis, it is necessary to understand first the emerging field of Social Computing, of which this thesis is a part. Social computing has brought a computational approach to the study and modelling of social interactions and communications [4]. Social computing has also brought insights into the design of systems that support those interactions. From a methodological viewpoint, social computing is cross-disciplinary and uses frameworks from human-computer interaction, communication, sociology, economy, psychology, as well as organizational and behavioural theories.

Following how social computing organizes information, this thesis studies: (1) how audiences and content producers interact; and (2) the design of novel tools to engage audiences. The methodologies used in this thesis also vary, for instance I conduct qualitative analysis involving interviews and qualitative coding, usability evaluations, and quantitative analysis including user and topic modelling techniques.

To analyze how content producers and audiences interact on social media, I take Burbank's (1967) definition of interactive audience space, and consider that people engage with audiences under a continuum that ranges from friendly spaces for interacting with friends and family, to adverse spaces for planning violent oppositions against enemies.

Previous related work studied primarily online fandoms [3]. Baym researched how independent music artists and labels in Sweden interacted with their audiences. The work began to uncover the social exchange emerging among these actors. However, a deeper understanding of how content producers and audiences interact in a variety of

scenarios is still missing. For instance, comprehending the type of social exchange that exists among these actors in more adverse scenarios is lacking.

This thesis uses three different case studies to understand the interactions between content producers and audiences in different scenarios. These case studies are not exclusive, but rather help us to understand the differences and similarities across scenarios. I build on the concepts from previous work on participatory culture [2], and online fandoms [3], to provide a more detailed understanding of online audiences.

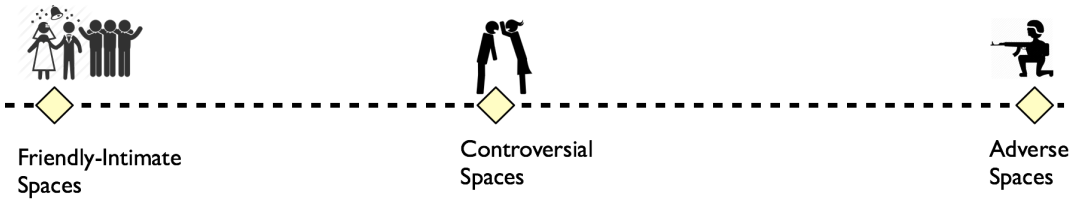

Figure 1.2: Continuum for interacting with online audiences.

I first examine the interplay between authors, audiences, and shared data in intimate settings where people share content with friends and family [5]. I focus, my analysis on one of the most important new interactions facilitated by social media: the targeting and integration of audiences in author's content. I conduct a qualitative analysis to understand people's perceptions of the interaction. Through my analysis, I uncover how authors integrate audiences in their content to gain social capital and facilitate collaborations with their audience. My study also highlights how people are reshaping their behavior to appeal to sharing algorithms.

My second study [6] focuses now on a space that is intermediate from friendly and adverse: an online community where authors discuss politics, gay-rights,feminism, racism, and other controversial topics with friends and strangers. I study another main relational aspect between authors and audiences: the self-presentation of authors. I use a mixed methods approach to understand how authors present themselves to their online audiences, the reasoning behind their selected self-presentation, and how this effects

the popularity of their content, i.e., how much audiences engage with authors. My study identifies that the most popular content usually comes from authors who adapt their self-presentation to their audience's interests. By being adaptive to their audience's interests, authors assembled more collaborations with their online audience.

In the third study [7], I analyze how authors interact with their audiences in adverse scenarios, specifically when opposing drug cartels. I inspect again some of the most important dynamics between authors and audiences: (1) the way authors garner support for their ideas; (2) the type of participation authors prompt from audiences; and (3) the type of content that triggers the most participation from the audience. My work illustrates how authors use online mobilization strategies to trigger participation from their audience in extreme and dangerous participation scenarios. It also shows how the online audience can now take part in defining the narrative of an armed conflict, something that was typically left to institutions and elites in power.

Through this analysis, I contribute an understanding of how audiences and authors engage with each other in a spectrum of online spaces. I identify key aspects of their interactions, which are present throughout the spectrum: a) authors and audiences spend time browsing social media to understand each other's traits; b) authors and audiences use their understanding of each other to probe strategies that will help them to build stronger relationships and better collaborate; c) authors and audiences are now directly engaging to produce collective efforts. These efforts can range from distributing meaningful content, enhancing one's online image, narrating collectively reality, or even engaging in offline collective action. I argue that audiences and content producers operate within a gift type economy. Wikipedia describes the gift economy as "...a mode of exchange where valuable goods and services are regularly given without any explicit agreement for immediate or future rewards. Ideally, voluntary and recurring gift exchange circulates and redistributes wealth throughout a community, and serves to build societal ties and

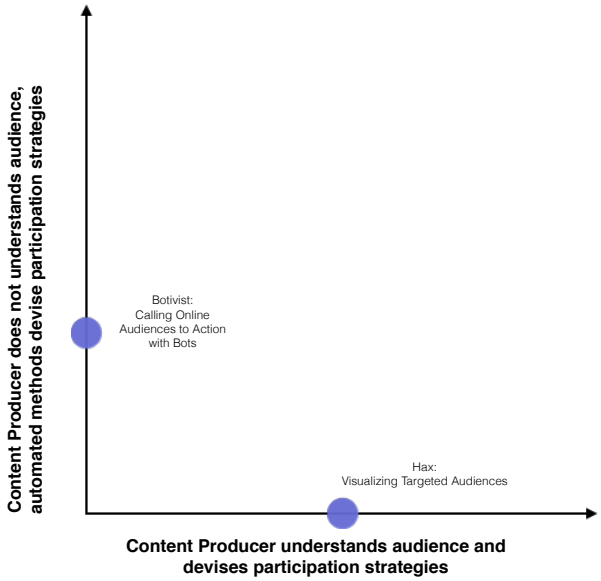

Figure 1.3: Design space of systems to engage online audiences.

obligations."

Text content in this new medium has become a gift. Content producers give the audience the stories and the information they enjoy, want and need. The audience can then reciprocate through attention, feedback, or even via actions that will help a producer's cause. Through these exchanges a network of support is forged. This support network facilitates collaboration and cooperation. The more "generous" that a content producer is to an audience, the more support she receives. Content producers and audiences have a shared sense of purpose and meaning exemplified in the gift economy. This economy is leading to several collaborations that include popularizing content, improving each other's public image, narrating events, producing offline collective actions.

I use the understanding gained from my analysis to define design principals for the next generation of tools to engage online audiences. I argue that given that these parties co-exist in a gift economy, it is important that these tools help content producers to motivate their audiences "to give back" to them. Content producers need to think how their content could realistically add value to the life of their audience. It is necessary

for content producers to think beyond their content, to how their audience could use the information, the context in which the data they share will be experienced, and especially to think of their audience's needs. Content producers thus need to think how that can create a value exchange, where they have a mutually beneficial relationship with their audience.

I postulate that to help content producers create these mutually beneficial relationships, it is important to first provide authors with tools through which they can motivate and incite action from their audience. I argue that to design tools that help authors to motivate audiences we have two axis: one where we aid authors to understand their audience in detail; so that they on their own can use the knowledge to devise how they will motivate participation from their audience; and another axis where content producers do not have to learn anything about their audience, but rather automated methods focus on probing different strategies to obtain contributions.

I evaluate my proposed design principals in two tools that focus on helping authors to better collaborate with their online audience. The first is Hax [8], an online tool that helps content producers to visualize and initiate collaborations with their desired targeted audience. I argue that multi-faceted data visualizations help content producers to better understand their audience, and engage in better "gift-giving" exchanges. The second tool is Botivist which via online bots probes different strategies to call online audiences to action. Botivist enables the next generation of audience engagement tools that not only measure the number of replies from the audience, but also facilitate collective action. I conduct an evaluation of Botivist and find that the majority of the online audiences ($>$ 80%) recruited by Botivist made contributions the helped advance a content producer's cause. My results provide evidence for the potential of this new class of systems to recruit and collaborate with online audiences .

The core idea of this thesis is to twofold. First, I provide an understanding of how

content producers and online audiences currently engage with each other on social media in a spectrum of scenarios. Secondly, I incorporate this knowledge into the design of novel tools that overcome currently limitations, and allow content producers to better collaborate with their online audience. I provide a set of design patterns for the next generation of audience engagement tools. To accomplish this goal, I (1) present a series of analysis on social media to identify the challenges and opportunities of authors and online audiences; (2) develop interactive tool that use the understanding to tackle the challenges, and use the opportunities advantageously. Both the analysis and the tools advance our knowledge on online communication and collaborations.

1.1 Contributions

The core contribution of this thesis is: (1) an analysis of how content producers and audiences interact with each other in a spectrum of online spaces; (2) design and evaluation of tools to engage online audiences. The sections below synthesize these contributions.

1.1.1 Understanding of the Online Audience Ecosystem

This thesis presents an analysis of how audiences and content producers engage with each other in a spectrum of online spaces that range from friendly to adverse. I identify key aspects of their interactions present throughout the spectrum: a) both authors and audiences spend time browsing social media to understand each other's traits; b) authors and audiences use their understanding of each other to probe strategies that will help them to build stronger relationships and better collaborate; c) authors and audiences are now directly engaging to produce collective efforts. These efforts can range from distributing meaningful content, helping each other to attain a desired online image,

narrating collectively reality, or even engaging in offline collective action. This ecosystem is the logical consequence of a gift economy.

1.1.2 Design of Online Audience Engagement Systems

This thesis provides evidence that interactive systems can target, recruit, and engage online audience's to support an author's goal. The systems I present in this dissertation for engaging with online audiences re-open areas of human-computer interaction design. For instance, Hax introduces new multi-facted data visualizations to infer the traits of an online audience. Hax demonstrates that multi-faceted data visualizations can facilitate understanding and collaborating with an online audience. An author can, e.g., find interested audiences who could physically attend an event supporting the author's ideology. Botivist introduces the concept of using autonomous agents to recruit online audiences and start to assemble collective efforts that advance an author's cause. These designs give rise to prototypes that leverage data visualizations and automated methods to facilitate collective efforts between online audiences and content producers.

1.1.3 Motivating Online Audiences

This thesis proposes two approaches to engage and obtain contributions from online audiences: (1) automated methods that probe different strategies to motivate and mobilize audiences; and (2) multi-faceted data visualizations that let authors understand an audience's interests to motivate and mobilize the audience themselves.

Chapter 2

Online Audiences in Friendly Spaces

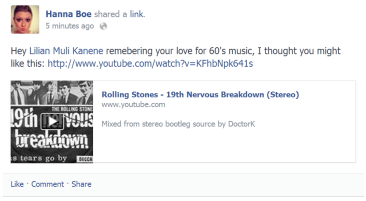

Figure 2.1: Example of public targeted sharing on Facebook.

On social networking sites (SNS), people have contacts from different life facets, e.g., college or work. This can lead to unintentionally sharing sensitive content with subsets of friends [10]. People engage in a spectrum of sharing modes to overcome this problem: from targeted sharing, where messages are shared with specific individuals, e.g., in an email [10], to public broadcasts where people share messages that are appropriate for all [10]. There has been growing interest in understanding people's perceptions of these sharing modalities. Bernstein et al. [10] studied targeted sharing in private messages and found that people saw this modality as a way to share personally relevant content.

Kairam et al. [11] used Google+ to study selective sharing, i.e., the sharing of content

This work was published in HT:26th ACM Conference on Hypertext and Social Media 2015. [9]. Special thanks to Andres Monroy-Hernandez for his valuable contributions to this research.

with specific predefined groups, e.g., "family" or "work." Their results indicate that people use selective sharing to evangelize, ask questions, or start conversations. Hogan [10] studied public broadcasting in social media. The work found that, in this sharing mode, people speak to the "lowest common denominator" to allow everyone to understand the message.

Studies have not yet fully addressed a new sharing modality of SNS: public targeted sharing. Figure 2.1 presents an example of a Facebook post with this sharing mode. Public targeted sharing combines elements of targeted sharing and broadcasting. When a person tags others, the post is shared with a specific targeted audience composed of the tagged individuals who receive a direct notice of the post. But the post's public nature brings additional audiences similar to a broadcast (e.g., the entire social graph of the content producer). This hybrid sharing mode also represents a public display of social connections, as the public posts carry the names of certain friends and links to their personal Facebook profiles [12]. Previous research has studied particular aspects of this interaction, such as the identity concerns it triggers [13] or the types of conversations that such posts elicit [14]. Little is known, however, regarding people's different views and uses of this sharing modality.

The aim of this investigation is to organize and understand the variety of perceptions of this hybrid sharing mode on Facebook. Although Facebook may appear to have a particular design for public targeted sharing, the concepts involved are not unique. Sites such as Twitter.com or Soundcloud.com also provide hyperlinks to the profiles of the people tagged in content. Google+ links tags to people's profiles and allows people to broadcast to the social circles of those tagged. Gold users of reddit.com receive notifications when others tag them.

Our interview-based study (n=120) discovers that individuals feel public targeted sharing strengthens friendships. Individuals also believe they can use it to overcome the

algorithmic filtering powers in play, exposing each other to more surprising content and people.

2.1 Methods

To examine the perceptions of public targeted sharing, we employ a methodology similar to that used in a study about people's perceptions of Facebook disclosures [13]. Similar to Lampinen et al. [13], instead of pursuing a longitudinal study, our goal is to shed light on the variety of interpretations that individuals have regarding this sharing modality. We conducted interviews to understand the perceptions.

2.1.1 Data Collection

We recruited participants both offline and online to avoid having only highly active Facebook users. To recruit online, we posted invitations to our study in different Facebook groups and to pages randomly selected from public listings. To recruit offline, we approached people in public spaces (e.g., bus stops and parks) and invited them to take part in our study. We labeled participants recruited directly from Facebook as F# and those recruited from public spaces as P#. Our recruitment material avoided the term "tagging" to not scare off people who rarely tag.

A total of 120 individuals participated in our study, 32 recruited from Facebook, and 88 from public spaces. Participants used Facebook to varying degrees. The individuals recruited from Facebook were 18-52 years old; 44% were female, 56% were male. The participants from public spaces were 18-68 years old; 52% were female, and 48% were male. The interviews had a median duration of 11 minutes, with the longest lasting 40 minutes, and the shortest 8 minutes. The interviews covered participants' interpretations of this sharing modality from the perspectives of the different social roles involved,

i.e., content producers (individuals who tag friends in content they share on Facebook), taggees (individuals tagged in the content), and viewers (individuals who view the tagged content). We counterbalanced the order in which we asked participants about each role. Participants voiced which roles they had taken and only responded for those that they had personally assumed. All participants had been viewers at least once, 87% had been taggees at least once, and 83% had been content producers at least once.

Our study began by asking each participant to write down examples of public targeted sharing. On average, each participant provided two examples (for instance, two Facebook posts where friends were tagged). For each role, we used the examples to start a discussion. For instance, we asked: "What are your thoughts when someone tags you in this type of post?" We encouraged the participants to lead the discussion. Our role was mainly to ask them to elaborate and ejemplify. Participants mainly reported on their perceptions of public targeted sharing in posts (Facebook status updates or wall posts). Some also reported on tagging in comments, photos, or videos. However, participants did not highlight any difference among content types. The interviews also covered participants' background in SNS use (sites and devices used, type of content they or their friends share, number of Facebook friends, etc.) We did not witness any difference in interview responses across age groups or SNS usage.

2.1.2 Categorization Perspectives

We used qualitative coding to analyze interview responses. This allowed us to establish a categorization of perceptions of public targeted sharing. Our overarching goal was to identify patterns in the perceptions of participants on this sharing modality. We transcribed all audio recordings to text and separated each interview into sections based on the social role that the interviewee was reporting about. We then jointly read each

interview transcript for each social role to identify key concepts and ideas. These initial concepts were then discussed. We decided that a category would cover a general type of perception and that, when possible, it would include perspectives specific to each social role. Concepts were then aggregated and similar ones renamed.

We also quantified the number of participants who viewed the tagging of Facebook users in posts as a mechanism that involved two mixed audiences. We considered this would help to further contextualize our results. Two researchers, unfamiliar with the work, read each interview and classified whether the person saw tagging as involving two audiences. The researchers found that all interviewees referred to tagging as a mixed sharing mode. Additionally, 71% of the interviewees specifically used the term "share" to describe tagging on Facebook.

2.2 Results

The following five main categories covered participants' perceptions of public targeted sharing. One person can be associated with more than one category, as a person can have more than one perception associated with the sharing modality. We did not identify differences between participants' Facebook usage and their perceptions on sharing.

Stronger Relationships. A majority of our interviewees (79%) perceived public targeted sharing as a way to build stronger relationships. From our interviews, it appears that tagging build stronger friendships because the targeted audience feels that someone is making time for them:

> *"I think the best thing about being tagged is that there is somebody who is considering and thinking about you. I always comment on the things people tag me because they took the time to think about me." 28, male F2.*

This result matches theory on disclosure, which states that the more an interaction makes people feel singled out, the more value they attribute to it [12]. Being publicly tagged has become a social signal that helps individuals feel unique and noteworthy. Content producers appear to be aware of this dynamic. Interviewees stated that, via public targeted sharing, they secured higher-quality responses from their network:

> *"[Tagging] ensures that you are working with only loyal people that you know will help and promote your cause, not sabotage it." 30, male, F1.*

Individuals appear to use targeted sharing to jump start discussions with sympathetic crowds. Interviewees also believed that, via tagging, Facebook's ranking algorithm was more likely to share their content with others. Some thus practice public targeted sharing to revive relationships:

> *"...I basically hope that friends I haven't seen in a while will see this [a tagged post], and they will comment, and we can reconnect [...] Basically I have had other people in mind when I have tagged – sort of thinking about the fact that they are also watching." 21, male, F55.*

Surprising Content. Interviewees (61%) perceived this sharing mode as a way to expose audiences to information outside of what they normally viewed on Facebook. Interviewees had the idea that friends had similar tastes. Facebook's recommendation algorithm thus tended to show one same type of content to the members of a particular social circle. Tagging broke this by enabling content producers to reach multiple social graphs and share surprising information, outside the graphs' norms:

> *"I like to tag people that I know are interested in something and whose audience will also care. But my interest is in creating a crossover. So I involve audiences that follow people in dance, but I share with them something totally different, such as poetry..." 25, female, F10.*

This is similar to the finding of Huang et al. [10] of how people tagged on Twitter to direct content to certain social graphs. However, perhaps due to Facebook's recommendation algorithms, we observed here an emphasis on using tagging to diversify a social graph's information.

Participants also had folk theories about how tags influenced Facebook's ranking algorithm:

> *"If I tag someone and then he comments, the post becomes "active" and it'll appear at the top of the News Feed. So I'll sometimes tag someone directly in the post, or in the comments. In either case, if it makes him comment, the post will be pushed to the top." 25, female, F10.*

Participants also believed that this sharing mode facilitates the introduction of audiences to new and surprising people:

> *"Being a group of friends, I want them to meet other people who they might find interesting. The whole idea of social media is that we can share. So they might see a face, but they don't know who they are; with the tag, they can find more information about them" 26, male, P3.*

Public Image. Interviewees (60%) believed that public targeted sharing can help refine one's public image:

> *"I tag friends to represent something I want to project. So let's say that I want to project that I am a super hipster, well then I will tag friends in support of a cause, I will casually mention that my friends and I are out writing in our Moleskine journals..." 29, female F5.*

Interviewees also reported that they tagged to link their content to the reputation or image of their friends and increase their content's value. Participants even reported that

they "cherry-picked" who they tagged. They analyzed how influential or relevant the image of a person whom they planned on tagging was in reference to the content they wanted to share. Similarly, participants also believed that Facebook's ranking algorithm would give more visibility to posts linked to high-profile users. Individuals thus tagged users with influential public images in an attempt to garner more visibility for the posts they wanted to broadcast:

> *"Facebook will always rank higher content tied to verified, well known accounts. You're losing a lot of visibility if you don't link your content. You have to select well who you tag." 41, male, P83.*

In contrast to selective sharing [11], public targeted sharing is also a public display of social connections [12]. This might cause individuals to relate it to public image construction, as everyone (people and algorithms included) can view how one interacts with others.

Participants saw public targeted sharing as an opportunity to tailor not only their own public image but also that of their friends. They tagged as a way to collectively help friends find content that favored them. Interviewees also expressed that, in general, they felt honored to be tagged and be a part of their friends' posts. However, some also voiced identity concerns because others selected the content linked to them and the content might not match what they wanted to portray:

> *"[Tagging] is a very public thing. It allows anyone to broadcast something about you, something that might mislead the idea of who you are. Through tags your identity can be created for you. Your identity becomes what other people broadcast about you, their idea of you." 26, male, O3.*

Content producers were aware that others are self-conscious about the type of content that is posted about them. We found that most empathized with their friends and usually

alerted them offline that they were going to tag them:

> *"A lot of people don't want to be tagged at a party. People are self-conscious of what they want or what they don't want about them online. [...] I always ask people if they want to be tagged. I am in the car and I say: Hey guys I am making a Facebook status and I am going to tag all you guys!" 19, male O7.*

This echoes the findings of Lampinen et al. [13] regarding people negotiating self-disclosures before conflicts emerge. However, we found that, despite the negotiations taking place, individuals occasionally had to untag themselves. This occurred when individuals were excessively tagged in their friends' content. Individuals wanted to show their involvement in their friends' lives (as it helped strengthen friendships.) Yet, showing too much of that involvement, e.g., being tagged in too many posts, obscured who they were:

> *"In every wedding photo they posted, they included my name regardless of whether I was in the photo or not [...] Because I don't have any other posts on my profile, that one wedding event began dominating my page and it began characterizing me because it was the only thing that people saw on my profile. When people went into my page, they thought I was really into weddings [...] My identity was made up by someone's wedding. I felt that was too emphasized. I thought about untagging myself from the photos but decided to leave them. I felt it was sad to disassociate myself from my friend's wedding." 29, male, O8.*

The nature of this mixed sharing modality generates struggles, as individuals want to show support for friends, while also portray their desired public image.

Unwanted Content. Interviewees (38%) perceived this sharing modality as annoying, because it brought them undesired data. Individuals considered that the content should have been shared privately with the targeted audience. Some taggees also found this sharing modality annoying because they received excessive Facebook notifications. Some untagged themselves as a result. Taggees and viewers alike felt annoyed with tagged posts generated by companies, especially if they were advertisements. Individuals appeared tolerant of targeted ads created by friends. They considered that at least their friends thought about them to target them and that the posts were a chance to be updated about their friends' lives:

> *"It's usually interesting to see what my friends are selling. One girl was tagging so we could see her ceramics creations [...] it's really interesting to see what they're up to..." 25, female, P52.*

Content producers are aware that this sharing modality can be annoying. They are thus careful about how often they publicly target others and in what content they target others. Interestingly, some see this annoyance as a chance to play with friends:

> *"One friend of mine was an Obama supporter [...] I was tagging them in some anti-Obama stuff to mess with them [...] I succeeded in being annoying...hehe..." 24, male, F15.*

Despite this, taggees and viewers felt that unwanted content could also bring them serendipitous discoveries, as it could expose them to interesting strangers:

> *"...I have personally met several people through this "spam", like people with whom I've had an interest to collaborate, and people with whom I've shared interests on this and other topics. It isn't common that you get to meet people in spam. They are rare exceptions. Marvelous exceptions that bring you marvelous opportunities." 20, male, F3.*

Recollection. Interviewees (33%) perceived public targeted sharing as a type of virtual diary that helped to document who played a part in their lives. Participants reported that they enjoyed returning later to these posts to reminisce. Interviewees thought it helped them to reminisce not only individually but also collectively. Individuals surfaced posts from the past to make their audiences relive memories:

> *"...For me, they [tagged posts] are more like a path for sharing memories. They are really about making your friends relive these experiences again..." 28, female, P60.*

Content producers also have the perception that by including a person's name in a post, they can implicitly remind the taggees of events that they will conduct together. In this case, content producers admitted to mentioning certain persons so that their posts would be favored by Facebook's ranking algorithm and people would be reminded of their event:

> *"...tagging helps to build the fan base of your event, and people will be constantly reminded of the event. Their friends are also more likely to see it [i.e., the content] if you tag. That's one of the things which Facebook loves and favors: tagging!" 21, male, F202.*

2.3 Discussion

We used Facebook as a medium to investigate perceptions on an increasingly popular sharing modality: public targeted sharing. In general, individuals feel that this modality helps to reinforce relationships, as content producers publicly show that they are taking the time to consider a particular targeted audience. Tagged posts generated automatically by companies are likely to not be perceived as positively, because individuals do not see any real person making time for them. This is similar to what was observed in the

Scratch online community, where people valued credit granted automatically less than credit given by humans [15]. SNS and UI designers could contemplate how to help companies better engage with their online clients, perhaps by encouraging more humanized sharing.

Individuals appeared to use public targeted sharing to build collaborative spaces, such as a space to reminisce collectively about the past, or a type of backstage to collectively help friends craft a desired public image. Individuals seem to have adopted public targeted sharing as a building block to create the flexible online social spaces they desire. We believe that it is important to design digital structures that allow people to collectively experiment with building on the technology. It is not about just user-testing all elements of an interface but, rather, testing whether people can make use of those elements to construct jointly the dynamic spaces they want.

Our study revealed that individuals view public targeted sharing as a way to expose each other to surprising content or even strangers beyond those recommended by Facebook's algorithms. Recommendation algorithms in general have sought to filter out opinions, people, and items that are different from us, potentially limiting the diffusion of information. Individuals appear to use this sharing modality as a means to free audiences from these "algorithmic biases" and distribute information that they consider fresh and interesting.

Individuals also perceive public targeted sharing as a way to reach audiences outside their immediate social circles. Facebook does not officially present tagging as a way to reach foreign social graphs but, rather, as a way simply to let people know when they are involved in the posts that one shares. Our findings thus raise the question of whether social media would benefit from more official digital structures tailored for targeting novel audiences [16]. Audiences could be bombarded with more unwanted content in this setting. However, it might also enable more serendipitous discoveries. Future work could

further explore this trade-off.

Across categories, we uncovered how individuals feel they can use this sharing modality to try to manipulate or "game" Facebook's ranking algorithm and obtain the viewership they want. This result is consistent with the recent work of Gillespie [17], which points out, "Teens have been known to tag their status updates with unrelated brand names, in the hopes that Facebook will privilege those updates in their friends' feeds." Editors were traditionally the gatekeepers of information. Today, it is filtering algorithms that determine what is relevant. Our results raise the question of what the role of human-centered SNS should be regarding the algorithmic powers in play. Again, there is a trade-off: allowing manipulations could expose audiences to more unwanted content, yet it might also help disseminate more diverse, unpredictable, noteworthy content.

Finally, similar to [13], we call for caution in generalizing our results. Although we took care to recruit a diverse set of participants, the pool of interviewees recruited outdoors belonged to particular cultural settings. We tried to counter this issue by triangulating their responses with those participants recruited online. However, future work could focus on a quantitative analysis with a more varied population.

Chapter 3

Audiences in Controversial Online Spaces

In this section I examine how people's explicit interests–as stated in their online profiles–align with those implicitly expressed through their online posts. Using a topic modeling algorithm, I cluster individuals based on said alignment between interests in profiles and those in posts in a particular online community. I use those clusters to analyze how they relate to the popularity and number of people's posts, and investigate the motivations behind these behavior patterns through online interviews. I find that for those individuals who contribute highly popular content, as measured through number of replies, profiles interests largely do not match the topics of their online posts. Conversely, individuals whose profiles match their postings, rarely contribute popular content. Together my results suggest a complex interplay between profile statements and behavior exhibited in a community.

The work in this chapter was the basis for a US patent application [6]. Special thanks to Rita Wouhaybi, Saurav Sahay, Lama Nachman for their valuable contributions to this research.

3.1 Introduction

Online community boards thrive from posts and comments. Their content is generally contributed by strangers who voluntarily decide to post and share knowledge. However, there are several challenges to keeping community members active and contributing to the board. Computer-meditated technologies suffer from social cue deficiencies [18]. It can be difficult for users to convey through computer channels important cues about who they are, what is their intention, or current context. The fact that people interact online with a large pool of participants, each with their own backgrounds and perspectives, intensifies the problem. People may ultimately decide to stop contributing to the board if they cannot get their point across, or if they cannot find mutual understanding with others [19, 20].

In this paper, my goal is to contrast people's profile statements to their community posts, and examine how having an alignment in the two relates to the number and popularity of people's contributions.

The literature on online community participation has uncovered that some of the factors that motivate people to contribute to an online board is: having self-worth, social affiliations [21], receiving social capital [22, 23], social norms [24], or to enhance one's reputation [25].

In many scenarios, researchers and practitioners argue that one of the best formulas to cover these needs is to prompt users to have profiles that disclose information similar to what they share in a community board [26, 27, 28]. Such profiles can help give users the recognition they want, while also providing verification of the claims people make on the board.

When people have profiles that match their posts, pro-social behavior and trust should logically be promoted [27]. It lets people authenticate others' claims, and put into context

others' contributions. This can potentially encourage more replies and interactions among users. For instance, if a person posts they are against *"Gay Marriage"*, other users might consider the person a troll (purposely posting inflammatory content) and not comment on the post. But, if users can authenticate the person's beliefs and see they are genuine, they are more likely to leave replies [26]. However, it has so far been understudied how having profiles that match one's posts relates to popularity, as measured through number of replies, and number of people's posts.

I specifically study the ways people's explicit profile interests may contrast with those implicitly expressed in their posts, and study how these differences relate to the popularity and number of people's posts. Instead of a longitudinal study, I aim to shed light on the variety of behaviors, and understand people's motivations behind their behavior patterns.

I build on the work of [29] which postulates that users present themselves to others through *disclosed data* (data users share in their profiles for others to see) and *entrusted data* (information users post on pages managed by others). I study in particular how much the topics of interest users explicitly present on their profile match the topics of interest they implicitly present in a community board through their posts. For example, a tag in one's profile on *"Society for the Prevention of Cruelty to Animals"* would denote one's profile interest for *"Animal Rights",* while using the tag *"Animal Shelters"* on one's posts would show an interest for the topic in the community board.

By improving our understanding of how profile and post similarities relate to number and popularity of people's posts, I gain knowledge on how to design better self presentation and collaboration tools that prompt high participation. Additionally, my study provides insights into the relational processes occurring in online communities, especially how users showcase their interests under different social contexts and spaces. I also create awareness for the type of online disclosures (both on profiles and posts) that are important for content popularity.

I propose a work flow that, using a topic modeling algorithm, clusters users from a particular online community, based on the alignment between the interests in their profiles and those in posts. The clusters allow us to analyze how the different alignments (or matches) relate to the popularity of people's posts. Through online interviews the motivations behind the behavior patterns I observe.

3.2 Related Work

My work builds on prior research on content contributions in online communities, online interest modeling, online identity and self presentation.

Content Contributions in Online Communities

Previous work [30, 31, 32, 33] finds that some of the main motivations people have for contributing to online communities is: to build their reputation and social capital, for reciprocity, to help the community, to portray their social affiliations, and for professional self-expression.

Researchers [26, 27, 28], believe that the benefits people want from contributing to a community board can be obtained by providing accurate communication and verification of users' identity, e.g., by providing adequate profile cues.

The work of [34, 35] shows that users tend to believe others' content, not so much based on the content itself, but more by a series of heuristics, such as people's user names. Similarly, [35] finds that there is a bias between a user's name and her number of followers. Recently, [36] shows that in reputation systems, which rate participants on the quality of their contributions, users do present different behavioral patterns depending on the rankings they see in others. [37] find that people post about different topics, and depending on the gender on their profile they receive different amount of comments from their audience. We note that different profile identity signals are indeed affecting online

contributions and participation.

But, evidence is still mixed on how different profile cues can affect contributions in a community. Some of the most famous and popular Internet content has been generated from more anonymous spaces, where there are no profile cues [38]. Through our investigation, we hope to provide a better understanding of how identity signals from profiles, especially if these match a person's posts, affect content contribution and popularity.

Online Identity and Self Presentation

In real world social interactions, people strategically alter what they present about themselves to achieve their desired relational goals [39]. It can be common for people to disclose different aspects of themselves to different audiences or situation [40].

Studies on online identity and self presentation [41, 42, 12, 43, 44], find that users also fine-tune their online profiles and social interactions to highlight and present information about certain aspects of themselves, and achieve different purposes. For instance, some users strategize what they disclose on their profiles to find romantic partners [41] or what they tweet about to practice self branding [43]. The way users decide to present themselves on their profiles can ultimately effect the online connections they make [44].

The body of research on online identity and self presentation processes in online communities has focused primarily on communities formed from the interpersonal relations of its members, i.e., social networking sites [41, 42, 12, 43, 44]. Little has been studied on the self presentation processes occurring in communities constructed around a common interest [45]. It remains unclear e.g., the role that profile self presentations have in the amount of popularity a person's posts receives.

Our work extends the research on online identity and self presentation by studying the variety of ways users present their interests on their profile and in a community constructed from common interests.

Unlike previous work, which focused primarily on privacy and personality effects of

self presentations [46, 47], we discuss how profile statements relate to post popularity, and amount of posts made. Understanding the role that profiles can have in online contributions, becomes increasingly important as we evolve how users present themselves across the Internet, while trying to maintain rich discussions.

We are at a point where technically we can afford to build online spaces that augment users' profiles with what they post [48]. However, before we reach a consensus on the type of online spaces we want to create, we first need to have a deeper understanding and recognition of how profiles that match one's posts relate to content popularity.

With our investigation, we hope to contribute to scholarly conversation about the role identity plays in technology mediated contexts. For example, [25] in her study of Usenet newsgroups, emphasizes the value identity has in online contributions: *"For most participants, identity – both the establishment of their own reputation and the recognition of others – plays a vital role."*

In this work, we explore how the way users present their interests on their profile and on their posts, relates to the popularity of their content. We extend the literature by providing a better understanding of the role profile presentations play in online community contributions and participations.

Modeling Users' Interests

Data modeling has recently been extended to characterize users based on the type of content they post, share, and use [49, 50, 51, 52].

[51] uses hierarchical cluster analysis, and principal component analysis to identify from tags in users' profiles, "interest neighborhoods." The work however, only analyzes interest expression in profiles. It never designed models to contrast users' profile interests with those presented in a community; or explored whether having a match between the two influences the number of replies a person receives. Similarly, [52, 53] study how users' interests can be modeled based on their tweets, tagging, and retweeting behavior.

[54] presents two semantics for expressing a user's relationship to a topic: interest and expertise. The work compares these two semantics across social media applications. Despite these novel forms of measuring and inferring users' interests, none of these works explored the amount of alignment that exists between users' profile interests, and those they present on a community board; or studied how this alignment relates to the popularity of the content.

Work on privacy and recommender systems [55, 56, 57] has studied the discrepancy between stated preferences and actual behavior. The majority of this research however, has not explored how these discrepancies play out in online communities, or how they relate to the number of content contributions people make.

[58] began a discussion about the biases users have for deciding to posting certain content. The work found that the events users decide to discuss in a community board involve a complex dynamic sitting between important personal events, and events users believe would be of interest to their audiences. We extend these ideas by contrasting users' personal profile interests with the interests they ultimately decide to present in a community board.

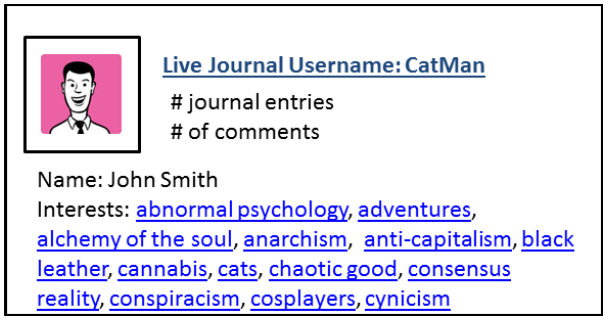

Figure 3.1: Structure of LJ profiles 3.1

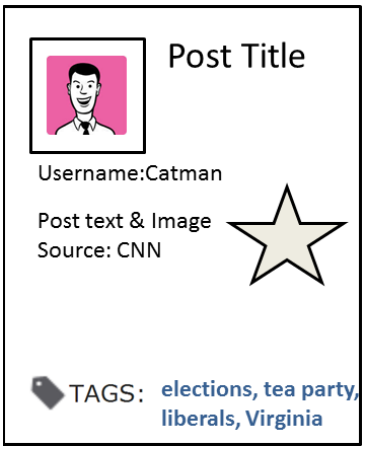

Figure 3.2: Structure of LJ posts 3.2

LiveJournal

LiveJournal (LJ) is an online blogging and social network that started in 1999 and accommodates about 40 million distinct accounts, of which about 1.8 million were considered active at the time of data capture [59].

Each LJ user picks a user name, fills a profile page, and can author a journal (blog). Users provide a list of interest tags on their profile page. A user can use free-form tags referring to concepts or names such as *"cats," "Lady Gaga," "Barack Obama," ...* We call these tags, **profile tags.** The structure of a LJ profile is shown in Figure 3.1.

Users also have the opportunity to contribute to community journals or boards. Community boards let its members create and post new entries. We analyze the data from a LiveJournal community board: ontd political (ONTDP.) Each post in an LJ board includes the contributor's user name, the posting date, title, text, a set of tags describing the post, and comments in reply to the post. Within the culture of LJ's boards, it is common practice to copy articles provided by external news sources, such as CNN [59]. Therefore, when a user posts something in a board, typically the only content the user actually creates are the tags associated with the post. The tags are free text and can

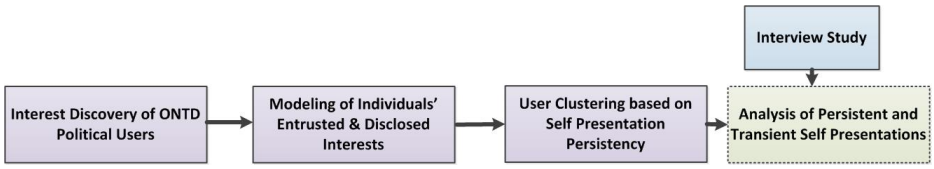

Figure 3.3: Work flow for analyzing the different ways users exhibit persistent and transient self presentations.

include anything that the user wishes to add, but often are meaningful keywords representative of the post. Figure 3.2 presents the structure of a LJ post with its different components.

3.3 Modeling Alignment in Profiles and Posts

Here we model the alignment between users' profile interests and those expressed on their online posts. For this, we propose the following work flow: (1) We discover the general interests of the community; (2) we use the discovered community interests to model each individuals profile interests, and the interests she presents on her posts (note that we model these values in terms of the general community interests to have a common ground for comparing users;) (3) we then cluster users based on the alignment between their profile and post interests.

We use the clusters to study how the alignment between profiles and posts relate to the popularity and number of people's posts, and investigate the motivations behind the behaviors we observe. We present in Figure 3.3 an overview of our work flow. In the following, we discuss each step in detail.

3.3.1 1. Discovering the Interests of the Community

We consider the general interests of a community represent the topics present in the tags of the community's posts, and in the profile tags of its members (i.e., people who contribute posts to the community board.) For instance, the tags *"Lady Gaga," "Beatles," "Country Music,"* represent an interest for the topic of *Music.*

We use a topic modeling algorithm, Latent Dirichlet Allocation (LDA,) [60] to discover the topics present in users' profile tags, and in their collection of tagged posts.

Note that many profile tags can be found in isolation and can represent proper names, e.g., *Lady Gaga*. But, such tags can be difficult for data driven methods, like LDA, to interpret. We propose to incorporate a knowledge ontology to introduce the lacking cultural relationships. In particular, we use the Freebase knowledge ontology [61]. We augment profile tags with their Freebase definition when possible.

Each document fed to LDA corresponds to either the tags of a particular community post, such as *"animals", "google", "technology", "youtube"* or a unique profile tag with its Freebase definition (stop words removed,) such as: *"JavaScript," "interpreted," "computer," "programming," "language"*. We use Bayesian information criteria [62] to determine the number of desired topics to input to LDA. While the Bayesian information criteria is not a "bullet-proof" way to specify the correct number of topics to consider, we use it to obtain a notion of the different types of interests presented by a community. After this step, we have a series of topics that represent the general interests of a community.

3.3.2 2. Modeling Individuals' Interests

Here we model the interests users express on their profiles and on their posts. We use the topics discovered previously to categorize users' data (i.e., their posts and profile tags.) We then use the categorized content to compare how much the interests (topics) users present on their profiles match the interests they present in the community board.

We first use LDA to categorize users' content into the discovered topics. We calculate the topic mixtures of each profile tag and of each post (i.e., the probabilities that a particular profile tag or post is about specific topics.) We categorize the content into the topic to which it has the largest probability.

We then calculate for each discovered topic the amount of interest each user has expressed through her profile tags, and through her posts. We calculate the degree of interest a user

presents for a topic T given her profile tags as follows:

$$P_D(u,T) = \sum_{w \in \mathcal{W}} P(u)P(w|u)P(T|w), \tag{3.1}$$

where $P_D(u,T)$ is the probability user u has disclosed an interest for topic T given her profile tags; $P(u) = n^{-1}$, is the probability of selecting user u; n is the total number of users; $P(w|u) = m^{-1}$ is the probability of a particular tag w appearing in user's u profile; m is the total number of profile tags that user u has; $P(T|w) = 1$, when the profile tag w is relevant to the topic T (i.e., the profile tag is classified under that topic,) or zero otherwise.

We use a similar approach to calculate $P_E(u,T)$, the degree of interest a user presents for a topic given her posts (i.e., the interest the user has entrusted in the community board.) Instead of calculating such interest over a user's profile tags, her tagged posts are used.

Note that by calculating $P_D(u,T)$ and $P_E(u,T)$, we can contrast the interest a user presents for T on her profile and on her posts. We can analyze, e.g., whether a person who contributes the most popular content about T presents also an interest for the topic on her profile.

For each individual, we calculate for all discovered topics her $P_D(u,T)$ and $P_E(u,T)$, and form two individual vectors representing the user's profile interests, and interests expressed in the community board. The size of the individual vectors corresponds to the number of discovered topics. The elements of the individual vectors hold $P_D(u,T)$ or $P_E(u,T)$, i.e., the amount of interest the user presents for each topic either through her profile tags or posts. We then represent each user as a feature vector, composed of these two individual vectors.

We use these feature vectors to cluster users who have similar alignments between their profiles and posts. For instance, users who only present an interest for topic T in their posts, and in their profiles might be grouped together. We then analyze for each cluster the amount of popularity and number of people's posts.

Table 3.1: Statistics of our data.

# (number) of users	296
# of posts	1,200
# of comments	30,934
Total # of unique LJ profile tags	9,812
Total # of unique tags used for posts	1,622
Average # of profile tags per LJ profile	50
Average # of tags per post	8

3.3.3 3. Clustering based on Profile/Post Alignment

At this point, each person is represented as a feature vector holding information about her profile interests, and interests presented through her online posts. We cluster people based on their feature vectors. Note that the clustering gives us an overview of the variety of alignments that people exhibit through their posts and profiles. For instance, one cluster may have people who only show certain interests on the community board, while never presenting them on their profile, another cluster may have people who only post about interests on their profiles, among other possibilities. This lets us study how different combinations of profile statements and posts relate to content popularity and number of contributions.

We use mean shift algorithm [63] to group together similar feature vectors, and discover clusters of people. The mean shift algorithm discovers clusters of people with similar alignments between the interests in their profile and in their posts. We decided to use mean shift algorithm because it is based on nonparametric density estimation, and therefore we do not need to know the number of clusters beforehand (unlike K-means.) We let mean shift algorithm discover the clusters from our data.

We use the clusters to study the different profile and post alignments people have, and examine how these different alignments relate to the popularity and number of people's posts.

Table 3.2: This table provides a sample qualitative outlook on the type of posts and profile tags categorized into each topic (note the content is independent of the user)

Topic	**Sample Tags**	**Sample Posts** [Post Title:*tags*]
1	"noir", "bondage", "erotica", "roleplaying", "bdsm"	Tennessee Governor Signs Bill Defining Kissing And Holding Hands As Gateway Sexual Activity:*sex, sex ed, Tennessee, wtf*
2	"Crohns Disease", "Herbalism", "Osteology","Disability",	Doctors will resent ObamaCare because of fat people!: *"fox news," "health care", "obesity"*
3	"African Literature", "Rednecks", "Racial Inequality"	I don't feel bad for black people anymore because I think a black person stole my bike: *"fuck this guy," "race/racism," "racial profiling," "racism,"*;
4	"Amish", "Church Music", "House of Lords","Anton Lavey",	Elderly nuns turn to the Internet to find new sisters: *"catholic church", "facebook" " get thee to a nunnery/monastery", "internet", "maine"*
5	"Homoeroticism", "John Greyson", "Homosex","Bisexuality"	Queer to the Core:*gender & sexual minority rights, history, lgbtq/gender & sexual minorities*
6	"Feminist", "Jewish Feminism", "Bitch","Tristan Taormino"	In which we learn what it takes to be a Real Feminist: *batshit, feminism, women*;
7	"Terrorism", "Iraq", "Misanthropy","Genocide"	FBI Terror Plot: How the Government Is Destroying the Lives of Innocent People: *conspiracies, fbi, muslims, terrorism*;

Data

For the purpose of our analysis we crawled 4 months worth (from March 30th 2012 to July 11th 2012) of posts (tags, user name, date and time) and comments (text, user name, date and time) from the the LJ ONTP board. Additionally, we also collected the LJ profile tags of "active" board members (i.e., any user who in the time window we study contributed at least one post to the board.) Statistics of our dataset are shown in Table 4.1.

Following the work flow described in *Discovering the Interests of the Community,* we discover 11 topics describing community members' general interests at the time of data capture. To illustrate the type of topics detected, we report in Table 4.2 a set of randomly selected sample posts and profile tags that were categorized into each of the 11 topics. We warn that some vocabulary may offend.

Adhering to the work flow described in *Clustering based on Profile/Post Alignment,* we use the discovered topics to categorize people's data. We model how much interest each individual presents for each of the 11 topics, and then represent each person as a feature vector. We use mean shift algorithm to cluster people based on their feature vectors, and find the different ways users mesh their profile statements with their posts.

We discover 4 different clusters representing the different types of alignments people's posts and profile interests present.

Validation

To demonstrate the effectiveness of LDA to categorize users' data, we perform a validation exercise to determine if the topics discovered by our topic model truly capture the topics present in the community's tagged posts and LJ profile tags. For the purpose, we randomly selected 55 posts and 55 profile tags (five per discovered topic,) and sought feedback from 20 participants from Amazon's Mechanical Turk (AMT.) AMT is a system to recruit participants to do small tasks for small payments. Participants were paid $0.55 USD. Each participant coded topics to either individual posts or profile tags. Each profile tag showed the tag and its definition from Freebase. We provided Freebase definitions in case participants were unfamiliar with the tag's

meaning. Posts were presented with their title and associated set of tags. Participants were shown the most probable words of each topic, and asked to select the topic most relevant for each profile tag or post. Participants were asked to provide feedback on each of their responses. We discounted blank AMT responses (i.e., those that did not provide any written feedback on their response.) We use the remaining responses (88%) to classify each post and each profile tag into the topic selected by the majority of participants. We used Fleiss' kappa [64] to evaluate the extend to which participants had agreement. Fleiss' kappa was .70.

We then compared the "ground truth" with the classifications that LDA gave to those same 55 posts and 55 LJ profile tags. We use again Fleiss' kappa [64] to evaluate the overall agreement between the ground truth and the topic assignments from LDA. Fleiss' kappa was on average 0.64 for coding profile tags, and .69 for coding posts. Overall, we obtained a Fleiss kappa score of .66. This correlation to human-coded posts and profile tags suggests that our LDA content categorization is reliable.

3.4 Interview Study

We conduct semi-structured interviews to capture users' viewpoints on profile interests, interests expressed through board posts, and having profiles that are aligned with one's community board postings. The interviews help explain users' motivations for engaging in the behavior we observe through our data analysis. It can provide for instance, explanations about why certain users only post about interests covered in their profile.

We conduct a series of semi-structured interviews with users via LJ private messages. We sample randomly users from each of the clusters that our work flow discovers, and via LJ private messages invite these users to participate in our interview study.

We obtained participation from 3 users of each of the discovered clusters. A total of 12 people participated in the study. Participants were primarily in their early 20s and early 30s, representing a dominant age group in social networking sites [65]. Participants reported they

used LJ as well as other social networking sites. All participants had completed a LJ profile, and had contributed at least one post to the ONTDP board. As an incentive for participation, each of the participants received $5.00 USD.

We specifically questioned participants on how they selected their LJ profile tags and the content they post on the community board; and how they feel they express their interests on their profile and on a community board. We also asked participants about whether they consider it is important to have profile interests that are aligned with their online posting. We ask questions that will help us to understand the motivations behind the behaviors that our data analysis discovers. The interviews provide information that we cannot uncover with our data analysis alone. For example, assume we observe that the users of a certain cluster have many LJ profile tags about a specific topic. We would then question such users about why they selected certain thematic LJ profile tags, and what they feel motivated their behavior. We combine these qualitative interviews with our quantitative analysis to understand in detail the different interest alignments that users exhibit, and what has helped drive such behaviors.

3.5 Results

We use the data and work flow described above to uncover the variety of interest alignments that people display, and study how each interest alignment relates to the amount of posts people create, and how popular their posts are. We also discuss users' perspectives on aligning their profile interests to the interest they present in a community board.

3.5.1 Types of Interest Alignments

For this analysis we: (1) characterize users based on how much the interests they present on the LJ community board match the interests they present on their LJ profiles; (2) cluster users based on how well-aligned their profile and community board interests are; (3) quantify the amount of users that belong to each cluster, and discuss their general characteristics, such as

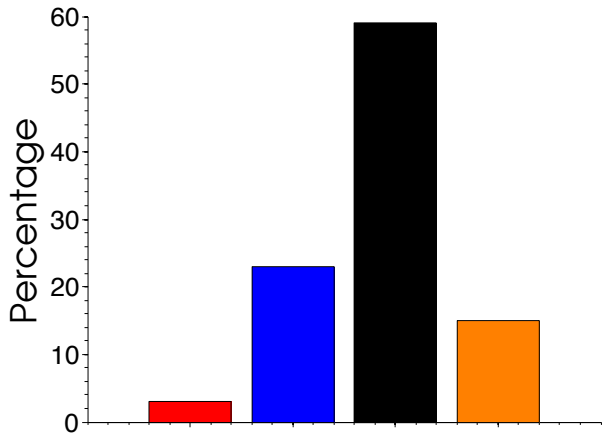

Figure 3.4: Percentage of users that belong to each cluster.

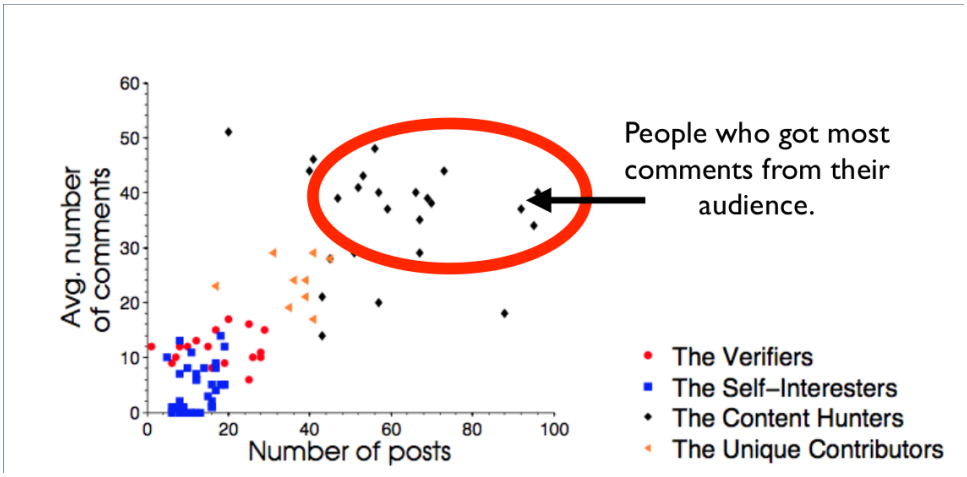

Figure 3.5: Number of posts created, and the average number of comments each of their posts receives for all users in each cluster.

amount of content they contribute, how popular their posts are, and their perspectives, among other traits.

We followed the work flow described in section *Clustering based on Profile/Post Alignment*, and identified four clusters that describe the different types of interest alignments that people exhibit. The traits of the users in each cluster are described below. We present in Figure 3.4 an overview of the percentage of users classified in each cluster. In Figure 3.5 we present the number of posts the user produced, and the average number of comments that the user's posts received (i.e., how popular the person's posts are.)

- **Cluster A *(The Verifiers)*:** These users (3% of all users) show for some interests an alignment or a match, i.e., they present the interests on both the community board and on their LJ profiles. However, on the community board they also present an interest for topics never mentioned on their LJ profile, i.e., they also show a mismatch or nonalignment between their interests. Per month they post on average about 4-5 different topics. Their content tends to receive a slightly sub-average amount of comments, around 20 comments per post. In our interviews, these users expressed they generally strive to have an alignment between their profile interests and the interests they present through their posts (they believed the cases when the two did not match was accidental, they likely forgot to include the interest in their profiles). They thought that having an alignment between profile statements and board activity, facilitates social interactions as it helps to **authenticate the opinions of others**: *"people do not have to prove credentials [i.e., have interest alignments] before engaging in a topic online but it is nice when there's some evidence that they are authentic in what they are saying [...] Determining whether a user is authentic or just a troll is difficult on the Internet." 24, Female.* This is in line with previous findings on online dating and profiles [41], which uncovered that users use different profile cues (signals) to determine the veracity of others' claims. Given this behavior of verifying the authenticity of other users, we name these users: **"The Verifiers."**

  This cluster has users who clearly see value in keeping an alignment between their profiles and their posts in a community board. However, notably, some interviewees felt they benefited from the alignment mainly in the setting of online discussions. It is unclear whether they would find such benefit in other settings: *"For discussions I participate in across the Internet, it is important [to have an alignment.] For what I surf/browse on the Internet, I'm not sure. Youtube, hulu, and imgur are time consuming" 30, Female.*

  These users appear to want an alignment between their profiles and community board activity to help their online audience put their opinions into context. They primarily

use their profiles to authenticate the interests they present in online discussions: *"The authenticity of a user is important in determining whether to take them seriously or not [...] I only presented what I think is the type of info people looking at my profile might want to know." 30, Female.*

These users on average used only one LJ profile tag to disclose a particular interest. Interviewees explained that having too many specific profile tags could be considered too personal, limit connections, and even be a security risk: *"Broader interests [profile tags] are better because people who are checking out your profile usually don't know you yet and are just trying to get a quick, general idea of what you're about. They aren't ready for your life story yet [...] Broader categories are also better to connect with a wider target audience [...] In addition, broader is better for privacy reasons." 24, Female.*

This behavior of using broad profile tags might be what is affecting the amount of comments their content receives. [44] found that users' social connections can be limited if they omit details on their online profiles.

- **Cluster B** ***(The Self-Interesters:)*** All of the interests these users (23% of all users) present on the LJ community board are the same ones they present on their profile. Their profile interests strongly match the interests they present in the community. Notably, unlike the Verifiers who present through their posts an interest in 4-5 different topics per month, these users typically only present through their posts an interest for one topic per month, and it usually is the same across months. This one interest is persistently presented. The majority of users from this cluster (98%,) presented an interest for topics 1 and 5. We find the people of this cluster tended to post almost daily about these two topics. But note from Figure 3.5 that overall they produced fewer posts than the Verifiers, and tended to receive slightly fewer comments.

  Interviewees from this cluster explained they strive to have an alignment between their profile interests and the interests they exhibit in a community board because **they view**

their profiles, and posts as part of themselves. It is thus important to be linked to content that matches who they are: *" I don't make a lot of posts to communities, but when I do make posts (or I repost or link on another site like facebook) it is almost always on a topic that would appear in my interest tags. I made a promise to myself that I would never lie about myself." 28, Female."*

Over 40% of these users use surprisingly the most profile tags among ONTDP users to display particular interests. For each of the topics presented on their LJ profile, they use on average 10 unique tags to express an interest. We observed that for expressing an interest in topic 1 this number grew to 14 profile tags per user. This is a relatively large number, considering that most other users used only one or two profile tags to denote a thematic interest.

Interviewees mentioned that they engage in this practice because they want their profiles to reflect their life story truthfully and advocate their opinions. Their different profile tags thus show how their opinions and feelings about particular interests have evolved over time: *"I promised myself that I would always be honest about who I am. It took me a long time to come to terms with the fact that I was bisexual [...] My interest tags have grown as I explored different fandoms and interests in my life. You'll find pairings that I read for different shows or books.[...] They [one particular set of tags about topic 5] specifically date to dust up over slash and gender issues that the Harry Potter fandom was dealing with several years ago [...] For me the tags are not just to define me/box me in, but to express how I feel when it comes to certain issues [...] I picked up some of my interest tags from other friends during a kerfuffle in the HP [Harry Potter] fandom a few years ago. There was a bit of history in that fandom [...] of advocacy by interest tags. As a way to voice displeasure with an outside group..." 28, Female.*

Unlike the Verifiers, none of these interviewees reported analyzing other people's profiles or activity (i.e., posts or comments) to verify truthfulness. For these users, profiles and online posts appear to be a much more personal experience. We conjecture these users

show an alignment between their interests because they only engage with the things they personally care about. The content they share has to match who they view themselves to be: *"I'm very open on the Internet (and offline if asked) about being bisexual. I tend to gravitate towards groups and people that are at least open minded when it comes to sexuality (groups like ontd_p) and through them I've learned and become interested in gender issues and identities as well. I will be more likely to read a story or a post, and more likely to comment, if it concerns an area I have particular interest in." 33, Female.*

These users also never expressed an interest with sharing content that would be of interest to the majority of the community board users. They were more concerned with sharing content that was personally true to them. Given this behavior of posting primarily about topics that were of their personal interest we name these users: **"The Self-Interesters."**

It is interesting to observe that their content tends to receive below average number of comments. Each of their posts receives on average 15 comments.

- **Cluster C *(The Content Hunters:)*** These users (59% of all users) have no alignment between their profile interests and the interests they present on the community board. However, they surprisingly do have profile tags relevant to the interests that other ONTDP users have frequently presented on the community board (mainly related to the interests expressed by the Self-Interesters.) Interviewees from this cluster explained they used their profile tags as a way to reach out to a certain type of individuals. That might be why we observe that their profile interests match the interests of others: *"my interest tags are used to link me with people I can connect with. I joined LJ to interact with people and I maxed out my list of interests towards the ends of casting a line to such people [...] If there's anything I enjoy more than being able to toss out an obscure-leaning reference, it's someone else being able to catch it." 25, Female.*

  These users use their profile self presentations to socialize. They are not displaying interests to be *"true to themselves"* or *"show their life story,"* but rather to connect

with others. Interviewees mentioned they adapt their profile information to appeal to more people: *"...I try to make myself as ambiguous as possible. People think that to have knowledge, experience, or authority, you have to be old. If I match my real profile with my writing then people write it off because they put me in a "Young white writer" bucket. That comes with the appropriate stigmas [...] I used to use my real name on LJ and I didn't receive very many comments. But then I changed my name to something ambiguous and...well..."meanigful" then people started commenting on my things because I was an idea not a person. I get around twice as many comments as "YYY" [his new user name, edited for privacy reasons.]" 23, Male.*

These users are willing to adapt what they display about themselves to socialize (connect) with a larger audience. They want to contribute content that will be well received by others. From Figure 3.5 we note that these users are actually the ones who contribute the most content on ONTDP and also the most popular. Each of their posts received on average 60 comments (almost three times the average.) The things they post on the community board is content they consider to be relevant and interesting for their audience: *"The content I post is to help a certain type of person find news information..." 23, Male.*

Interviewees thought that having profiles aligned with one's community postings could limit the way people express themselves. They also believed that on LJ, unlike certain social networking sites, people were free to exhibit interests unrelated to their profiles: *"I don't think it is too important [to have profiles that match one's board activity.] When you have personas to match your virtual self it becomes filtered, like in real life [...] There are many things I post on LJ that I can't post on facebook [...] With FB, I feel like I obviously need to try to keep my virtual face intact. Which is why I don't like it, it feels filtered [...] but on LJ you can share whatever you want (to some extent) without repercussion [...] on facebook, people care SO MUCH, because that is them..." 23, Male.*

Users feel they can show without repercussion different interests in a LJ community board and on their LJ profiles. This might explain why we see a great number exhibiting

a mismatch between their profiles and board activity. Users feel free to tailor the interests they display for each audience, and obtain high community participation in their content (i.e., large number of comments.)

The need to cater to an audience, also made some users compete with others to find and share content that the majority of the community members would appreciate: *"...I prefer to share good/cute/funny news stories and there sometimes aren't a lot of those to go around. Even when there are, I've been beaten to the draw several times by others who found and submitted the story first. " 25, Female.*

Given this behavior of searching and posting content will appeal to the rest of the community we name these users: **"The Content Hunters."**

These users appear to ascribe little value to having an alignment between the interests they display on their profile, and the ones they display on a community board. They probably benefit more from **following the community's norms,** showing interest for things the community cares about, and building connections and audiences through those norms: *"...John Smith [name changed for privacy reasons. The user used his real name here.] can't say Skrillexes Leaving EP is the best music made in the past ten years on FB because the readership of FB won't understand. So maybe it is the network I am in no so much the person I created. Maybe I created a John to fit the network [audience]. In other words, the John that fits into soundcloud is ZZZ [user's soundcloud handle, changed for privacy reasons] which is a totally different readership than the John on LJ..." 23, Male.*

The community also appears to benefit from users whose displayed interests are adapted to their audience. It is these users who create the largest community turnout.

- **Cluster D** ***(The Unique Contributors:)*** These users (15% of all users) also have no alignment between their profile interests and the interests they present on the community board. However, unlike the Content Hunters whose profile interests relate to those that others have presented on the board, these persons present profile interests that hardly

anyone posts about on ONTDP (mainly topics 9,10, and 11.)

Interviewees expressed they viewed their LJ profile tags more as a tool to give others better navigation to their personal content (e.g., personal online journal.) *"They [profile tags] are literally only category names, like ones I would use to name files. I use them [...] to provide easy navigation for visitors..." 27, Female.*

Contrastingly, on the community board they personally contribute to some of the most popular content (each user received approximately 33 comments per post, i.e. higher than average.) They tended to generate fewer content than the Content Hunters, and received also fewer comments on their posts. However, they received on average more comments than the Verifiers and the Self-Interesters.

Interviewees explained that on ONTDP they were not driven to share content about specific topics, but rather posted content when they felt they had something unique to share with the community: *"I am more likely to participate if I don't see that my thoughts have already been articulated by someone else [...] the topic of the post itself doesn't necessarily enter [...] It's more to do with if I feel like my opinion is redundant or not " 27, Female.*

"...I don't often engage, [...] because anything I could think to say would just be one more voice in a chorus" 25, Female.

Given this behavior of trying to contribute unique content we name these users: **"The Unique Contributors."** The fact these users try to contribute unique novel content that is of relevance to what the community currently is discussing might help them obtain more comments than what the Verifiers and the Self-Interesters receive.

These interviewees never mentioned that they saw value in having profiles aligned with their activity on a board. They considered that their profiles were one particular space where they could express some of their interests, and the community board was simply another space for expressing other interests: *"I care about the subjects that get posted on*

ontd_political, but that care doesn't define me as a whole or the part of me I put online any more than any other profile interest." 30, Female.

These users **use different online spaces to present different interests,** to exhibit different sides of themselves. Their online profiles are used more as a space to display interests related to their personal life. The content they contribute to the community is more attuned to what the community cares about, and whether they feel they have something unique to add. They collaborate with others' efforts to present unique perspectives.

We observe from this analysis that the Unique Contributors, and the Content Hunters contributed the most popular content, and also the most content. It appears that users who tried to have an alignment between their profiles and posts, were less in touch with their audience than users who had no match between the two. The Verifiers and Self-Interesters never mentioned that they posted content that appealed to others' interests. They were more self conscious about contributing content that showed who they really were. This might have resulted in their content receiving fewer comments, as they might have posted about things uninteresting to the rest of the community.

We believe that mismatches between one's profile statements and one's community board activity emerge not because users want to deceive or engage in disorderly behavior. But rather because users are in a "fast-paced" environment where they want to be the first to share content that will be interesting to others. This drives users to present on the board interests that they consider will appeal to their audience, even when they diverge from their profiles. These users clearly have a good grasp about what other community members care about. They thus tailor their interests to appeal to others, and obtain high participation in their content.

Our analysis shows that the users in each of the discovered clusters have signatures characterizing their communication patterns. For instance, the Self-Interesters use many profile tags to display their interests, and the posts they contribute always match their profile statements. Furthermore, we note from Figure 3.5 that the amount of content and popularity that peo-

ple's posts create also varies among clusters. Compare in Figure 3.5, e.g., the Content Hunters and Self-Interesters. Given that social roles are typically defined as a set of unique connected behaviors conceptualized by people in social situations [66], we note that by studying users' self-presentation patterns across online spaces, we may begin to uncover users' social roles.

3.6 Limitations

Our proposed work flow for contrasting users' forms of interest expression via topic modeling techniques has some limitations. First, the state-of-the-art of topic modeling oftentimes only partially captures human intuition of topics and relevance. Second, topic modeling cannot identify users' playfulness, satire, or snarkiness for presenting their interests. For our setting, we argue that while these shortcomings may add a bit of noise to our data they do not invalidate our findings because we study the relative comparison of how users present their interests for different topics under different contexts. Therefore, computational issues related to incorrectly matching a particular profile tag to a topic, or not identifying that a certain tag was satirical, will likely impact equally all topics and all tags.

We use Freebase to augment users' profile tags, and better interpret them. But, profile tags can be ambiguous. It is possible we assigned Freebase definitions that were not what the user intended, thus affecting our analysis. However, to counter this limitation we triangulated our results with semi-structured interviews with LJ users, who provided rich descriptions of how they engage with their profile tags to present their interests. Note that the semi-structured interviews also help overcome the limitations of not being able to detect users' playfulness and satire.

We also acknowledge that working with one particular online community and service does not necessarily provide generalizable results about the human behaviors we observe. However, since this area of research had not been investigated in depth before, we believe that examining the behavior of a subset of the online population a valuable starting point. It let us explore how

much users' profiles align with the content they post in a board, and study how these different alignments relate to the amount and popularity of created content.

3.7 Discussion and Future Directions

In this paper, we investigated how the interests people explicitly listed on their profiles align with those implicitly expressed through their online posts. We studied how these different alignments relate to people's motivations, the amount of content people post on a community board, and the amount of popularity their posts garner from others.

We have shown that users who exhibit a mismatch between their profile statements and posts, generally contribute the most popular content and also the most content. They tend to understand the community, and adapt the interests they display on the board to be meaningful to others. It is interesting that some of these users reported adopting different "personas" per social network. This might hint at the possibility that we should think about designing systems that focus not so much about who the user is sharing with, but rather who she shares as. What role is the user currently playing, and how can we design systems that support users' role play?

There might be value, e.g., in designing systems that can automatically detect users' roles, and identify ways to assist users in the tasks they perform when in that role. The Content Hunters, e.g., could benefit from tools that help them find new content that will be interesting for other community members; whereas the Self-Interesters might want tools that inform them of board discussion opportunities where they could share their political convictions or opinions. Users in general, might also benefit from data visualizations that inform them of others' roles. This could provide the "authentication" that some desire when engaging in online discussions. Providing such cues could also facilitate collaborations, as users would be more aware of the motivations, and unique traits their peers bring to the table. This can ultimately help users strategize their collaborations.

We are moving toward a point of persistent online identities. Some practitioners believe

that every single piece of online content users have ever contributed, should be tied to their real identity and be reflected on their profiles [67]. Proponents postulate that this will create better, more trusting and healthier online interactions, as users will have to be accountable for their content, and think twice before contributing things that could harm others. However, our study revealed that users who worried about how their posts related back to their profile, usually did not contribute the most popular content. It is therefore possible, that if we create scenarios where users are forced to think how all their posts will tie back to them, they might become more self-conscious about the content they share, and less attuned to sharing things that could be of interest to others. This can ultimately affect the amount of participation that a community board receives, and thus, possibly shorten the community's life.

It is also interesting to reflect that some users (the Unique Contributors) thought that different online spaces (e.g., profiles and community boards,) could and should present different aspects of themselves. We believe that instead of designing more systems that augment users' profiles based on their online activity [48], there is likely more value in creating systems that support those users who would like to present different aspects of themselves to different audiences. Such systems could provide socially aware profiles, or "presentation-cards" that depending on the viewer showcased different facets of the user.

Our study shows that users had very different types of alignments between their profile interests and the interests they showed in a community board. For instance, although the Content Hunters and the Unique Contributors users both had interests that were unaligned, Content Hunters displayed profile interests that were meant to appeal to others, while Unique Contributors presented more of their own personal interests on their profiles. We believe that a key to online collaborations is creating an environment of fluid presentation, where users can decide to post, e.g., with their real names but fake interests, or with a nick name and real interests, among other options. We consider it important to give users choices of how they present themselves online. For certain contexts, users might adopt a particular role and might not want their profile information to be linked to their activities. Not because they are involved

in malicious activities, but because their personal profile information might be irrelevant and obstruct the role they want to portray, or the online contributions they want to make.

As the large Internet contenders, such as Facebook and Google, evolve the type of online presentations users maintain throughout their online interactions, it becomes increasingly important to understand how users mesh the way they present themselves in different online spaces, and how this relates to the amount of content they contribute, and to their content's popularity. We believe our study can provide an initial informed foundation to better understand this phenomenon.

Chapter 4

Online Audiences in Adverse Spaces

4.1 Introduction

Militias have historically been defined as a group of three or more concerned citizens who are non-professional fighters, but have entered a combat situation to defend their community [68]. This can include the community's people, territory, property or laws. As social media adoption spreads, militias have started to make use of these networked platforms for organizational purposes [69, 70, 71]. From militia groups in Somalia to right-wing militias in the US, social media has become part of their communication strategy. Weeber et al., for example, argued that US militias find internet technologies appealing because these technologies are user-friendly, reliable, and free [70]. Similarly, Crothers [69] found that militias take advantage of social media to publicize their ideas and even raise funds. More recently, Weimann [71] examined how militias use social media to carry out some of their main goals, such as distributing propaganda or identifying potential supporters.

Civilians, not necessarily affiliated with militias, have also been using social media to ameliorate the disruptions that warfare brings to their lives [72, 73]. For instance, citizens going

This work was published in CSCW: ACM Conference on Computer-Supported Cooperative Work 2015 [7]. Special thanks to Andres Monroy for his valuable contributions to this research.

through the perils of the Iraq War formed online networks that exhibited resilience properties such as self-organization, reconfigurability, and redundancy [72]. Similarly, in the case of the drug war, previous work has examined how residents of communities afflicted by the violence used Twitter to form alert networks and help one another identify potential danger on the streets [73].

In this paper, we focus on the relationship between militia-related groups and online audiences over social media. Previous work on militia communication practices has identified that the goals of these groups, especially when engaging in public channels, are to popularize their ideas and recruit new militants [74, 70]. With that in mind, we examine which ideas are most "popular," and which posts "recruit" participants by engaging them in conversation. We guide our exploration of militia groups on social media with the following research questions:

1. What content do these groups share online, and how does their audience react to such content?

2. What content attracts new active audience members?

3. What are common attributes of the most active audience members?

4. What are the characteristics of the most popular content?

We explore these research questions using nine months of data from a Facebook page titled *"Valor por Michoacán* SDR (VXM), which translates as "Courage for Michoacán." This page was created to inform people of "unsafe situations" or "situaciones de riesgo" (SDR) in the central state of Michoacán. Shortly after its creation, the page started to focus on the activities of the self-defense forces [75], a new militia movement that emerged in the region.

The movement of the self-defense forces began in 2011 when a group of vigilantes from the town of Cheran in Michoacán expelled a band of illegal loggers backed by the drug cartel of "La Familia" [76]. This uprising became one of the first cases of ordinary citizens joining forces to challenge the powerful drug cartels. Two years later, new armed groups known as

the *autodefensas* (self-defense forces) emerged in the same state but much stronger and better organized —and, some argued, with backing from wealthy ranchers [77]. The self-defense forces gained national and international attention after successfully taking over a number of bigger towns previously controlled by the Knights Templar cartel, a spin-off of the La Familia cartel [77, 76]. Although the self-defense forces are a decentralized movement composed of multiple smaller groups, each with their own idiosyncrasies, they share the belief that the government is unable to stop the cartels and bring safety to their towns. Thus, they decided to take action and fight the drug cartels themselves [76].

It is important to indicate that although the VXM Facebook page documents the activities of the Michoacán self-defense forces, often with access to seemingly privileged information [78], the connection between the page and the militias has not been made official. Furthermore, the administrators of the page remain anonymous. News media in the country, however, often refer to the VXM page as the militias' social media presence [75, 79]. The page is also usually one of the first to provide on-site reports about the militias' battles [78].

We characterize the content that VXM shares and its audience's participation using analytical techniques similar to those for examining discourse between politicians and their audiences [80, 81]. Our aim is to explore the context under which militia-related content gains attention and to investigate the behavioral patterns of VXM's active audience. In particular, we focus on the topics, locations, public figures, and organizations to which the posts or comments refer. We explore topics, organizations, and public figures because they provide a window into people's interests and issues they may want to make more prominent or de-emphasize [81]. For instance, perhaps the most active audience members emphasized certain social aspects of the situation in Michoacán. Additionally, geographical locations can help to further profile and contextualize people's discourse [80].

We discuss our findings by exploring the factors that have been used previously to interpret the online activity of militias and other social movements [70]. For instance, we examined the role that information and communication technologies play in mobilizing individuals and in

helping people narrate or frame their reality [82].

Through our study, we reveal how VXM and its audience appear to be constructing a participatory space where they collectively define important events and public figures. Together VXM and its audience propose solutions to the problems they recognize, and specify what each person can do to transform the situation in Michoacán. We also identify the techniques VXM is apparently using to mobilize its audience for offline collective efforts. Our case study highlights how technology is transforming who can frame the narrative of a region and facilitate the mobilization of individuals for collective efforts.

4.2 Related Work

We position our work within the area of social computing and collective action, particularly focusing on social movements. Social movements are defined as organized collective activities which take place outside political structures, but attempt to transform existing political, economic, or societal structures [82]. Militias are considered to be a type of social movement [69]. The literature [82, 83] has identified three universal themes that recur across all social movements: mobilizing structures, opportunity structures, and framing processes. We use these themes to discuss and interpret our results, while building on previous work.

Mobilizing Structures: mechanisms through which information communication technologies (ICTs) facilitate mobilizing people to promote an ideology, or political cause [82]. For instance, in the Arab Spring and Occupy movements, people used social media to mobilize crowds for street rallies or protests. Facebook also facilitated the mobilization of political activists in Palestine by letting people link their private life with political activities [84]. Additionally, digital platforms such as Turkopticon allow workers to mobilize themselves around workers' problems and hold employers accountable for their actions [85].

Opportunity Structures: characteristics of a social system that facilitate or hinder the activity of a social movement [82]. For instance, in the Israel-Palestine conflict, the Israeli

Posts	6,901
VXM Fans	158,000
Participants (active audience size)	127,374
Comments	108,967
Post Likes	1,481,008
Reshares	364,660

Table 4.1: VXM Data.

policy of not licensing Palestinian 3G providers has become an opportunity structure because it forces Palestinian activists to invest more money in providing internet via landlines, and limits how their movement operates [84]. Online donation campaigns also take advantage of the immediate emotions after a disaster or of friendships to harvest higher levels of participation and donations [86].

Framing Processes: ways to interpret reality where one's individual experiences are labeled within the experiences of society [87, 82]. In social movements, frames represent a set of beliefs and meanings that help to legitimate and motivate the actions of the group. Frames usually emerge by negotiating a shared meaning. Diamond et al. [87] found that people used social media to crowd-source framing processes, and collectively create an interpretation of street harassment.

4.3 Data collection and analysis

Using the public Facebook API, we collected nine months of posts, comments, likes, and reshares from the VXM page. The data collection took a snapshot of all activity starting from August 14, 2013 to the date of data collection, May 7, 2014. The details of the dataset are shown in Table 4.1.

Content Analysis

In order to obtain a descriptive assessment of the VXM page content, we extracted common topics, organizations, and public figures referenced in the posts. First, we familiarized ourselves with the VXM page by frequently reading its posts and exchanging notes and observations.

Topics: We used a grounded theory approach [88] to identify the main topics in VXM's posts (Table 4.2). We started by independently extracting topics from a set of 200 randomly selected posts. Then, we analyzed the emerging topics and adjusted the list accordingly. Finally, we looked at a set of 500 randomly selected posts and produced a final list of mutually exclusive topics.

We used oDesk to hire three Spanish-speaking, college- educated people to categorize the VXM posts independently. First, we asked two coders to categorize each of the 6,901 posts using the topics listed in Table 4.2. We asked the coders to pick the "most relevant" topic for each post. The two coders agreed on 6,080 posts (Cohen's kappa: 0.54). We then asked the third coder to label the 821 posts upon which the first two coders had disagreed. We then used a "majority rule" approach to determine the topic for those posts.

Public Figures and Organizations: We identified posts that referred to each of the organizations and public figures involved in the conflict, from public officials to cartel members to militia leaders. We manually created a list of full, common, and nicknames and then identified posts that mentioned any of the names. We collected the names from Wikipedia and from *Proceso* magazine. We used the Wikipedia articles, *"Caballeros Templarios Knights Templar cartel"* and *"Grupos de Autodefensa Comunitaria Self-Defense Forces,"* written both in English and Spanish, to manually identify all relevant proper names. We did the same with *Proceso* articles. We used *Proceso's* online search to find articles related to the militias using the keywords *autodefensas* and *caballeros templarios.* We went through each of the articles that appeared in the first page of the search results. We then added or merged the alternate names for each person, both shorter and nicknames. For example, *"Estanislao Beltran"* was also known as *"Papa Smurf,"* and *"Jose Manuel Mireles Valverde"* was usually referred to as just *"Mireles."*

Audience Participation Patterns We examined the behaviors exhibited by the VXM audience, particularly the commenting behavior of the most active audience members (*the outliers*). We focused on comments because they typically take more time to produce than a

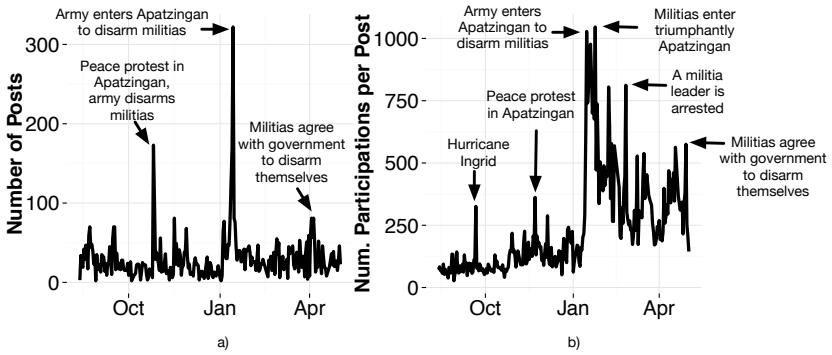

Figure 4.1: Overview of VXM's posting behavior and audience's activity. In Figure 4.1a, spikes indicate an increase in the daily number of posts; in Figure 4.1b, they show an increase in the daily number of audience likes, comments, and reshares per post. Together, VXM and its audience appear to signal when there are major events in Michoacán.

"like" and can provide clues to how dedicated a person is in a political/social movement [89].

We identified VXM's *highly active participants* by finding those individuals whose number of comments was higher than three times the standard deviation (normal procedure to find outliers.) We then profiled the behavior of these outliers by analyzing their total number of comments and the type of public figures and organizations they mentioned. For this purpose, we first grouped all of the public figures and organizations in our list based on their affiliation (i.e., type.) For instance, *"Mireles"* and *"Papa Smurf"* are grouped together because both were members of self-defense forces. We obtained the affiliations using data from the *Proceso* and Wikipedia articles mentioned earlier. We assumed that a public figure could have only one affiliation. Through this process, we identified four main types of public figures and organizations: government, militia, drug cartel, and journalist (news reporters), and classified each of the public figures in our list into one of these types.

Next, we measured how often each highly active participant mentioned each public figure in their comments. We then used a probabilistic model to implicitly calculate the degree to which they mentioned a certain type of public figure:

$$P(o,T) = \sum_{f \in \mathcal{F}} P(o)P(f|o)P(T|f), \tag{4.1}$$

where $P(o, T)$ is the probability outlier o mentions a certain type of public figures T given her comments; $P(o) = n^{-1}$ is the probability of selecting outlier o; n is the total number of outliers; $P(f|o) = m^{-1}$ is the probability of public figure f appearing in the comments of outlier o; m is the total number of words that outlier o has used in her comments; $P(T|f) = 1$, when the public figure f is of type T or zero otherwise.

For each highly active participant, we calculated their $P(o, T)$ for all types of public figures and formed a vector representing how often the participant mentioned each type. The size of each vector will correspond to the number of types— in this study, the size is four. These vectors help profile the behavior of the most active audience members in terms of the number of comments and words they produced, and the type of public figures they mentioned in their comments.

We used the vectors to cluster outliers and uncover the different behavioral patterns of the most active audience members. For instance, people who mentioned government public figures more than they mentioned militias might be clustered together. We use mean shift algorithm [63] to group together similar vectors and discover clusters of people. We opted to use mean shift algorithm because it is based on a nonparametric density estimation, and therefore we would not need to know the number of clusters beforehand (unlike K-means.) Instead, we let mean shift algorithm discover the clusters from our data. The clusters represent the different behavioral patterns that the most-active audience members exhibited.

We then looked at the characteristics of people in each cluster, particularly (1) the number of comments and words used; (2) the percentage of comments that mention certain type of public figures; and (3) the percentage of comments that reference locations. For the last, we identified comments when each person mentioned a city or town in Michoacán. We used a list from the National Institute of Statistics and Geography to identify comments that referenced such locations. We examined mentioned locations for contextualization purposes only; the clusters were primarily groupings of individuals based on their number of comments, wording, and public figures referenced.

4.4 Results

R1: The VXM Content and Audience Participation

Our first research question explored the type of content that militia-related groups share online, and how unaffiliated civilians respond to, participate in, or interact with such content. For this purpose, we examined (1) the daily activity of VXM and its audience; (2) the public figures that VXM covered; and (3) the topics VXM discussed with its audience.

1. Daily Activity: We plotted the daily number of VXM posts (Figure 4.1a) and audience participation (likes, comments, and reshares) per post (4.1b). We used this first analysis as an overview of how active the page and audience are.

Figure 4.1 shows that audience activity per post increased over time, especially in 2014, whereas VXM posting behavior stayed relatively constant. However, both audience and VXM administrators spiked in daily activity. For both, most of the peaks corresponded to times when major offline events occurred in Michoacán, especially events concerning the militias. For instance, on January 14, 2014, the day the Army entered the town of Apatzingan to disarm the self-defense forces, both the number of posts on the page and the number of times the audience participated per post spiked.

What is particularly fascinating is that the audience also exhibited spikes not related to VXM's posting behavior. For instance, on February 8, 2014 when the militia took over the Knights Templar headquarters, the audience had one of its largest activity spikes, whereas VXM did not increase its number of posts during that time. The audience might help to provide a more complete picture of when important events are transpiring in Michoacán. We question, however, why VXM only increased its number of posts during some of the militia's major events. Could VXM be posting selectively?

2. Public Figures: To inspect further the type of content VXM shares and how its audience reacts, we analyzed the top 10% of names VXM mentioned and studied how the audience responded to these references.

Figure 4.2 illustrates the total number of posts VXM created about each public figure

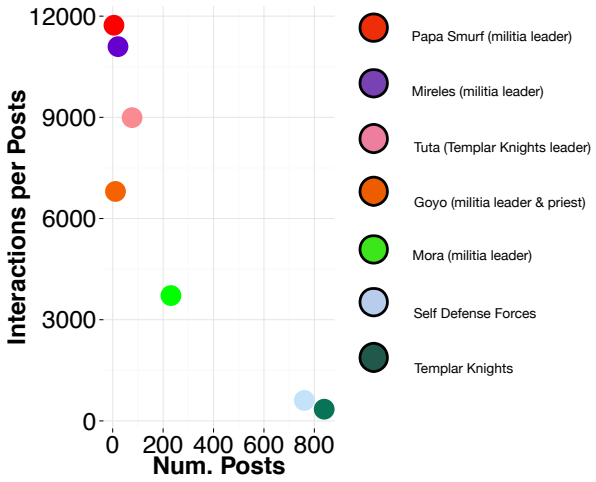

Figure 4.2: Overview of how much VXM mentions certain public figures and the public figures' popularity. The most popular public figures were not necessarily the ones most mentioned by VXM.

or organization and compares it with the subject's popularity, measured in terms of the total number of likes, comments, and reshares the content received. Each color-coded point represents a public figure or organization. The X axis represents the number of posts about the public figure or organization; and the Y axis, the public figures' popularity. A larger X value represents more mentions by VXM page administrators; a larger Y value represents more popularity with the audience.

Figure 4.2 shows that public figures and organizations that VXM mentioned most often were not actually the most popular with VXM's audience. We also observed that the audience appeared to be more participatory with posts that referenced individuals, rather than organizations. For instance, posts mentioning Knights Templar' leader *"La Tuta"* and posts mentioning *"Papa Smurf"* were much more popular than posts referencing *"Knights Templar"* or *"Self-Defense Forces."*

What is surprising about Figure 4.2 is that although *Mora* was the militia leader to whom VXM gave the most coverage, *Mora* was not the most popular. The most popular was actually the other militia leader, *Papa Smurf*, whom VXM hardly covered. *Papa Smurf's* popularity might be due to the deal he negotiated with the government in which the self-defense forces

officially became a rural police force [77].

It appears again that the audience is helping to provide a more thorough picture of who and what matters in Michoacán. Together, VXM administrators and its audience might be framing what important events are taking place in Michoacán.

Topic	Sample Post
Online activity:information on the nature of the VXM page, pointers to other reliable social media sources, and tips.	"If you want to make any complaints about things occurring in your town, try not to publish it on the Facebook wall, send us a private message instead. It is for your own SAFETY. You can also create a fake profile and send to us the complaint, although we will never share your identity."
Critiques: criticism of people or organizations, clearly targeted to attack.	"What do you guys think of this shameless governor, who is a nobody, who wants to report the brave people that told him his truths? Fausto Vallejo [governor of Michoacán] you don't serve as a governor, you are a NOBODY!"
Propaganda: Patriotic or nationalistic messages used to promote or publicize a political cause or point of view.	No more impunity! we see a Michoacán that wants justice that has a hunger for true justice. Only the people can save the people!
News reports: factual reports, and alerts. Occasionally, the messages include commentary of how the report could impact the reader, or actions the reader could take in relation to the report.	"#SDRLosReyes there is a confrontation, the Knights Templar are attacking in Los Reyes through different entrances. Please alert others who pass those zones. If you have a gun please come help the militias."
Missing people and obituaries: Reports regarding runaways, possible kidnappings, and deaths.	"#MissingPerson #Morelia, Michoacán His name is Jose Manuel Molinero Ruiz."

Table 4.2: Description and sample post of the topics covered by VXM.

3. Topics: We next examined the type of topics the VXM page presented (Table 4.2,) and how much coverage the page gave to each (Figure 4.3).

The dominate topic VXM covered was objective news reports. A distinctive trait of these reports was that they usually made requests to the public for some sort of action. For instance, the following VXM post reported a robbery in the town of Los Reyes and asked the public to help identify the criminals: *"#LosReyes there was a burglary in the jewelry shop in downtown Los Reyes today at 7 pm. Please check out the YouTube video. If anyone can identify the criminals, please report them via private message."*

These posts occasionally also included advice for keeping safe: *"#Apatzingan in the dance club Layon the organized crime [Knights Templar] is having an event. WE RECOMMEND NOT ATTENDING. IT IS FOR YOUR OWN SAFETY. The club is full of dangerous people."* VXMs news reporting appears to be factual and caring at the same time. This type of reporting

might help the page have a much deeper connection with its audience and facilitate working together towards a common goal [90].

Providing news reports aligned with the page's description in its "About" section: *"Here [at VXM] we provide and follow reports about dangerous situations occurring in Michoacán. ...We give coverage to events which take place in the state of Michoacán and are of social interest."* What is interesting is that VXM also conducted activities not reported in its description, such as sharing propaganda or providing subjective critiques. Of note, VXM incorporated the latter type of content in a subtler way. Only around 10% of all VXM posts were actually propaganda or critiques.

Communication research has shown that it is hard for a message to reach people who are not already in favor of the views the message presents [91]. Thus, for speakers who want to plant their own ideas into an audience, the most effective way is to do it subtly, so that the audience is unaware that a report or commentary is actually indoctrination [92]. Similarly, if VXM only shared propaganda or critiques, the number of people interested in their content could decrease. Subtly sharing indoctrination types of messages might be an effective technique VXM administrators use to maintain their large audience while still promoting their own ideas or opinions.

Another interesting topic the VXM page covered was online activity. These posts tried to explain the online dynamics of the page, such as what type of online content audiences should expect from the page (e.g., news reports), social media sources the page trusted, and how audiences should conduct themselves on the page. The last item focused primarily on giving tips on how people should present themselves online if they decide to participate in VXM.

This topic resembles a theme seen recently in social media content from terrorist groups [71]. Because of the potential danger of being caught, terrorists advise their audiences how to connect to their YouTube videos and present themselves online (e.g., with fake accounts), and which other internet pages are trustworthy allies. Clearly, VXM is not at all related to terrorism. However, there can be value in studying their online activity under the lens of terrorism,

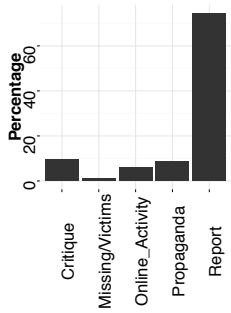

Figure 4.3: Overview of the percentage of posts VXM creates about each topic. VXM primarily provides News Reports. However, occasionally it shares Critiques and Propaganda. It also provides space for its audience by sharing their missing-persons reports, obituaries, and explaining how to keep safe online.

especially given the high risks also associated with participating in VXM. The public figures that VXM covered were notoriously dangerous (e.g., drug cartel members). Thus, VXM, just like the terrorists, needed to advise its audience on how to interact in such a potentially dangerous online space. Note that the page is an adverse space not necessarily because of its heated discussions, but because drug cartel members could easily eavesdrop on the page's content and violently target audience members whose comments they do not like. It is compelling that VXM makes an effort to explain the page dynamics to its audience and tries to keep them safe.

One topic that was hardly covered by the page (2% of all posts) but was nonetheless very fascinating was the topic of missing people and obituaries. The page not only provided reports about major events occurring in Michoacán, but also provided a space where people could report loved ones that had gone missing, run away, or recently passed away. The VXM page administrators created a space that let people mourn their losses and help each other when they suffered personal family tragedies or emergencies. Thus, VXM appeared socially engaged with its community.

R2: The VXM Newcomers

Our second research question explored the context under which people started participating in an online space linked to militias. We examined the per-day number of people who participated for the first time with VXM.

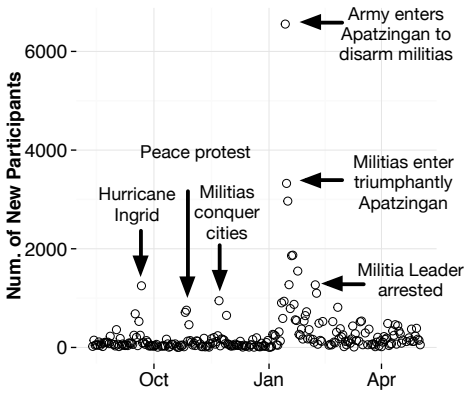

Figure 4.4: Number of daily new participants. In general, the page saw an increased number of new participants whenever there were major offline events.

Figure 4.4 shows number of new participants (newcomers) who either liked or commented on VXM page content. Most of the peaks in the number of daily newcomers correspond to major offline events that occurred in Michoacán, especially militia-related events.

For instance, around the time the militias took over the Knights Templar's headquarters on February 8, 2014, the number of newcomers increased. This increase during major offline events might explain the increase in audience participation during the same periods. New, external people suddenly participating in the page could be helping to provide a more complete picture of when there are important offline events in Michoacán.

R3: Behavioral Patterns of the Most Active Audience

This research question aimed to investigate the behavioral patterns exhibited by VXM's most active audience members. Using the procedure described in "Audiences Participation Patterns," we found that VXM's most active participants (282 individuals) exhibited three distinct behaviors (i.e., we discovered three clusters).

Cluster A *(Drug Cartel Savvy)*: All individuals in this cluster (33% of the most active) referenced primarily public figures related to drug cartels. On average, each of these individual's comments referenced drug cartels (67%), government (12%), and militias (8%), with no comments mentioning journalists. This apparent knowledge about drug cartels led us to name people in this cluster the *"Drug Cartel Savvy."* Drug Cartel Savvy contributed a median of

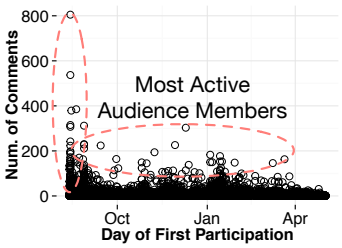

Figure 4.5: Overview of people's start date and the number of comments they generated. The most active audience members started participating either from the start or when the militias had an offline event.

112 comments; each comment had an average of 36 words (compared with 2 comments, 10 words from average participants, who never mentioned any public figures). Drug Cartel Savvy contributed the largest number and the longest comments of all the behavioral clusters. Sample comments include, *"Yeah right, as if the Templar [Knights Templar] were so innocent. He [the Knights Templar leader] claims they're not stealing or blackmailing us. I think we need to tell la Tuta [Knights Templar leader] what his boys are up [to]. My mom would say that skunks don't realize how much they stink!"*

Forty percent of each of these individual's comments referenced towns in Michoacán (average users never mentioned locations). They also referenced the most locations of all the cluster types. A sample comment with locations: *"Right now it is confirmed that Tepalcatepec and Buenavista have those signs [signs that say a town was freed from the Knights Templar]. They look beautiful! I feel pride whenever I see them. God bless the brave people of my neighborhood Loma Blanca! [neighborhood in Tepalcatepec, Michoacán]."*

Cluster B *(Government Gossipers)*: All of these audience members (66% of the most active) mentioned primarily government figures and related organizations in their comments. On average, each of these individuals' comments referenced the government (28%), militias (23%), drug cartels (5%), journalists (1%), and towns or cities (4%). This apparent interest in discussing the government led us to name people in this cluster the *"Government Gossipers."* Government Gossipers contributed a median of 64 comments; each comment averaged 12 words. Government Gossipers contributed the least content of the three clusters. However, the majority

of the most active audience belonged to this cluster. Sample comments include, *"Oh yeah, but fucking Karam [Attorney General] claims he already rescued Michoacán. Dirty old liar! He and EPN [President] are a disgrace!"*

"Come on brothers, don't give up! EPN [President] will just have to learn what the people can do together!"

We also found that a few (1%) of the Government Gossiper accounts tended to share the exact same comment many times. For instance, *"We need to be united as a country, and help the people in Michoacán. The government has committed many injustices against the militias, like throwing them in jail. God isnt with the corrupt. God will not help them."*

"Even the drug cartels confirm that they are funded by high ranked US officials. Do you guys think they'll share this on TV?! It is our responsibility to distribute and share this!"

This behavior resembles how people use Twitter in war environments where there is a tendency for redundancy [72].

Cluster C *(Geographers Michoacanos)*: These audience members (1% of the most active) did not reference any of the public figures we had on our list. However, they did mention locations. Each person, on average, had some mention or reference to towns and cities in Michoacán in 20% of their comments. We thus decided to name people in this cluster *"Geographers Michoacanos."* Although they did not mention the most locations (Drug Cartel Savvy referenced more), they were the only ones who mentioned locations without apparently involving drug cartels, militias, or the government. The Geographers Michoacanos contributed a median of 74 comments; each comment averaged 17 words. A sample comment: *"Michoacán is home to brave, patriotic, and hard working people. In Nueva Valladolid [former town in Michoacán] the first university of all of Latin America was built. The priest Miguel Hidalgo [historical figure] also studied in Nueva Valladolid, today Morelia, [capital of Michoacán]."*

All of the clusters had audience members with varied start dates. Even the Geographers Michoacanos had accounts that started when the page was founded through the middle of 2014. Figure 4.5 illustrates how much each audience members commented, given their start dates.

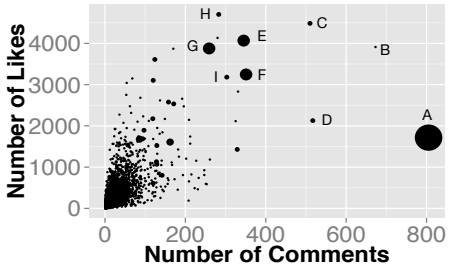

Figure 4.6: Number of likes and comments for all posts, size of a point is directly correlated with the number of reshares of a post. The most popular content was published during major offline events of the militias.

The X axis shows the date of first comment on VXM; the Y axis shows the number of comments contributed. We highlighted the audience members we considered for this analysis. Note that the majority of the most active participants appear to have started contributing either when the page was first founded or during a major offline event.

R4: Content Popularity This research question examined some traits of the content that became the most popular in militia-affiliated online spaces; that is, content that received the most number of likes, comments, and reshares. Figure 4.6 illustrates the number of likes, comments, and reshares of each VXM post. The X axis shows the number of comments the post received; the Y axis, the number of likes. The size of the circle relates to the number of reshares; the larger the circle, the more the post was reshared. We defined VXM's most popular content as the posts whose total number of likes, comments, and reshares deviated by three times the standard deviation. We labeled the most popular posts alphabetically and examined a few of their characteristics (Table 4.3).

In general, VXM's most popular content emerged during the militias' most important offline events. For instance, over half of the most popular content was created around January 15, 2014, the day the self-defense militias had their biggest confrontation with the government, who wanted to disarm them. The other half of the most popular content was created around February 8, 2014, the day the militias triumphantly entered Apatzingan, a town that had been considered to be the Knights Templar headquarters.

Label	Post Description
A	Video of a solider killing an unarmed civilian who supported the militias. (Jan 15)
B	Post announcing that the militias took over the Knights Templar headquarters. (Feb 8)
C	Post describing the type of hardships the militias and people in Michoacán have endured against the Knights Templar. (Jan 15)
D	Video of a corrupt police officer. (Feb 8)
E	Video explaining the militias' motives for fighting. (Jan 19)
F	Video of Father Goyo (priest and militia member) justifying the militias' actions. (Jan 14)
G	Video of Michoacán's youth explaining they are tired of living in the violence generated by the Knights Templar. (Feb 6)
H	Photo of Father Goyo asking civilians to take action against the drug cartels. (Feb 8)
I	Photo of Father Goyo crying because the government wants to disarm the militias. (Jan 18)

Table 4.3: Characteristics of the most popular content. The most popular content was published during major offline events of the militias.

The high number of likes, comments, and reshares exhibited during these periods might be due to the increase in newcomers (see Figure 4.4). Content published during major offline events is likely seen by a larger audience and can become popular more easily.

4.5 Discussion

We analyze our results using the universal themes that usually recur across all social movements: framing processes, mobilizing structures, and opportunity structures [82, 83].

Framing Processes. Through our case study, we observed that VXM page administrators appeared to: (1) structure and narrate meaning for particular events, public figures, and social life; (2) diagnose what things might need alteration in Michoacán; and (3) propose solutions to the diagnosed problems and, in the process, specify what people can do.

The VXM page administrators are likely framing a shared understanding of the situation in Michoacán with their audience. Such narrative has also been found in the online activity of US militias [93]. However, what is particularly fascinating here is that the audience appears to be helping to construct the narrative as well. For instance, we saw how the audience eventually took hold of the space that VXM had set up and massively participated whenever there were major offline events. The audience helped to provide a more complete picture of what was taking place in Michoacán. It was not necessary for VXM to give wide coverage to a certain event or public figure; instead, on its own, the audience adopted the space and participated extensively. The audience appeared to dictate, by itself, what was important.

This behavior has also surfaced in other online spaces. For instance, in "Hollaback!" [87], it was the online audience who defined the meaning behind street harassment. In a Facebook group set up to discuss the Coffee Party in the US, certain audience members assumed advocacy roles and extensively supported certain issues [94]. Group administrators, in fact, usually provide no input. Traditionally, it had been up to the organization's leaders (e.g., page administrators) or people in power to structure and define what matters. Technology, however, is

empowering audiences to establish for themselves what is noteworthy.

Mobilizing Structures. We witnessed that VXM was fostering solidarity with its audience by providing safety advice and even assisting them during personal family troubles (e.g., when a loved one passed away). Research on social movements has found that in order to mobilize people and produce action, it is first necessary to establish a certain degree of solidarity [70]. It is likely that the time and space that the VXM administrators are dedicating to their audience is helping them create a highly participatory environment. Note that some of these solidarity conducts resemble online strategies that politicians have adopted to encourage participation from citizens [81].

Through our study, we also saw that offline events appeared to be important for mobilizing new participants and popularizing VXM's content. This behavior resembles what has been observed on Twitter [95], where people who joined that service during an emergency event or mass convergence usually became long-term adopters of the technology. Offline events, and the numerous polemics that surround them, might help audiences concentrate on what truly matters to them. This, in turn, is likely helping to massively increase people's participation in the page. We believe that, in the end, the skillful use of offline events could play a major role in pumping new life into an online group or movement.

Opportunity Structures. Our study revealed that some of VXM's most active participants used the page to actively discuss the drug cartels and the government. There are likely several explanations for this behavior. Smelser [96] identified that social strains (i.e., impairments in the relationship between two actors) usually prompt collective action, and the type of strain generally defines the type of collective action that will emerge. In this study, it is likely that these audience members have a strain with the drug cartels or the government that led them to participate extensively and even exhibit distinct behavioral patterns in VXM. The strains might actually be an opportunity for VXM to recruit highly active participants.

It could also be worth thinking about the role these highly participatory citizens might play in Michoacán or in the country in the long run. Maruyama et al. [97] found that the people

who participated on Twitter during a debate were more likely to change their vote in a political election. Their work argued that to take part in an online discussion, people need to process the related information more elaborately; and this could lead to changes in their offline actions. Currently, the offline consequences of citizen participation in militia-affiliated online spaces are unclear. Future work could explore the offline effects that are surfacing from online spaces such as the VXM page.

4.6 Limitations

We acknowledge that working with a Facebook page linked to one particular milita group does not necessarily provide generalizable results about the human behaviors we observed. However, since this area of research has not been investigated in depth before, we believe examining the behavior of a subset of the online population to be a valuable starting point. It allowed us to explore how militia-related content is presented to an audience, and how audiences interact with and participate in such content.

4.7 Design Implications

Our results raise the question of whether showing the framings constructed by an audience could benefit social media. Summaries of an audience's constructs could certainly help historians make better sense of a social movement or politically charged event. Audience participation coupled with news reports could provide historians with a more complete picture of when important offline events transpired in a region; and also whether organizations or groups tried to frame reality in a certain way, as their reporting could be compared with the audience's .

Interfaces highlighting the narrative of an audience could come in different flavours, such as: infographics of an audience's size and number of participations through time; or visualizations

of the audience's interpretations on various matters. Some social media sites, e.g., Facebook or Twitter, already provide tools to visualize the participation of an audience, mainly to improve people's online communications [98]. Similarly, current research [99, 100] has explored visualizing the interests and social contexts of audiences to facilitate collaborations. Following these ideas, we believe that interfaces which structure and visualize audiences' framings could facilitate collective action. People could better identify collective thinking about a problem, and could work together towards possible solutions.

Many audience members appeared to have specialized knowledge about organized crime, government, and certain geographical regions. There is likely value in designing platforms that can advantageously use audiences' comprehension of the world [101]. Perhaps it is about combining citizens' knowledge with the resources from organizations to create smarter cities [102]. For example, we could design platforms that integrate the audiences' understanding of the crime in their towns or passion for their neighborhoods [103, 104] with government resources to create safer and more lively cities.

On this last point, note however, that currently individuals and organizations have to be savvy and learn on their own the best online strategies to spark collaborations or participation from audiences. As designers, it could also be worth thinking about the type of teamwork we want to facilitate. It is unclear whether it is about just building mechanisms through which organizations can donate resources for citizens to use [105], or if we need to create a more complex platform that intrinsically helps these two players build solidarity with each other.

Finally, our results found that some of the most active audience members appeared to have social strains against certain organizations. When we design collaborative tools, it could be beneficial to consider these strains (e.g., citizens might not trust the government). We likely need to think about how to create platforms that empower citizens and organizations to collaborate towards improving their cities, while taking into account possible fallouts.

4.8 Conclusions

Traditionally, it was organizations who defined what was noteworthy in a social movement. We showed that social media is empowering audiences to frame for themselves the events, and public figures that matter. Audiences appear to now collectively help narrate a social movement's story.

Our results also showed that not all major offline events translated into VXM Facebook posts. Previous work has begun an exploration between an individual's online posts and real-world activities [106, 107]. However, the relationship between online posts and real world events in more adverse scenarios has yet to be explored. An understanding of this relationship could impact research areas that range from prediction of a population's activity to crowd-sourced news [108, 109].

We found that VXM page admins conducted different mechanisms to apparently foster solidarity with their audience. Creating solidarity seems to implicitly help VXM to create a highly participatory online space. Finally, this result suggests that the correlation between the levels of participation of an audience and the solidarity created by page admins warrants deeper investigation.

Chapter 5

Discussion of Understanding Online Audiences

In these three case studies studies there were themes that recurred on how authors interact and engage with their online audience. The first theme is that authors have different strategies to obtain significant contributions from their audience. One of the most prominent strategies was to have solidarity with the audience. For instance, in friendly scenarios authors shared content they considered could help enhance their audience's online persona; in adverse scenarios, authors shared tips on how audiences could keep safe from violent drug cartels. By having solidarity with their audience, authors were able to ensure contributions that enhanced their cause.

Content producers also took time to understand their online audience and cater content for them. Sharing content that was relevant to the auduence's interests regardless of an author's own personal taste, helped authors to popularize their content and receive a larger turnout for their calls to action.

Similar to online fandoms [3], audiences and content producers appear to interact with each other within a sharing economy that is based on giving, receiving, and reciprocating. Content is here the gift. Reciprocation of these gifts can take a number of forms. Such as participating in an author's cause; enhancing an author's public image; or popularizing an author's content.

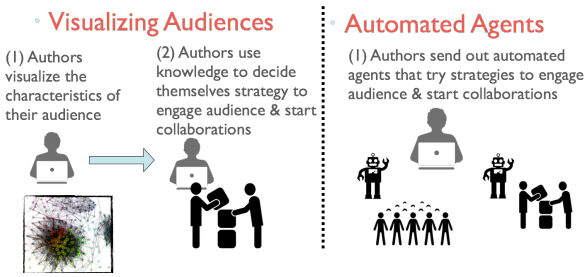

Figure 5.1: Systems design for engaging with online audiences.

Note that to strive in this gift economy, content producers must create relationships with their audience that inspire feelings of gratitude and obligation. This primarily so that the audience will want to give back, and provide support for the author.

However, current sharing tool are not specialized in helping authors to exchange contents as gifts to their audience. Authors by-themselves have to figure out how they will increment their audience's participation. Many times investing hours of their day in studying their audience and identifying how to better communicate with them.

In the next chapters I propose novel platforms that seek to overcome the limitations of current sharing tools.

I consider that to help content producers to harvest a relationship of mutual benefit with their audience, we need to provide authors with tools through which they can understand their audience's needs to then use that knowledge to motivate and incite action from their audience. I argue that such tools live within two axis: one where we aid authors to understand their audience in detail; so that they on their own can use the knowledge to devise how they will motivate participation from their audience; and another axis where content producers do not have to learn anything about their audience, but rather automated methods focus on probing different strategies to obtain contributions. Content producers in the latter case can jump in once the collaboration is initiated and take hold from their.

I evaluate my proposed design principals in two platforms, one focused on data visualizations and the other focused on automated methods.

Chapter 6

System for Visualizing Targeted Audiences

Users are passionate about sharing their political convictions, art projects, or businesses. Many times they also want to direct social interactions to certain people to start collaborations, or to raise awareness of their causes. However, users have scattered unstructured information about the characteristics of their audiences, making it difficult to deliver the right messages or interactions to the right people. Existing audience targeting tools often only allow people to select potential candidates based on long predefined lists, which provide few insights about audience candidates. We explore instead the idea of using data visualizations to help people dynamically identify audiences for their different sharing efforts. In this paper, we introduce the motivations behind visualizing online audiences, and describe a research implementation we have designed to experiment with the concept in the setting of targeting audiences in an online community. We provide the results of a preliminary empirical evaluation which shows the strength of the idea and areas for future research.

People are the main driving force behind fostering healthy and successful collaborations

This work was published in COOP: Conference on the Design of Cooperative Systems 2014 [110]. Special thanks to Angus Forbes, Carlos Toxtli, Grant Mckenzie, and Shloka Desai for their valuable contributions to *Hax*.

and interactions [111]. Through different actions, people can create favorable collaborative environments. For example, they can start conversations; encourage contributions, advertise their projects [112], among other related activities.

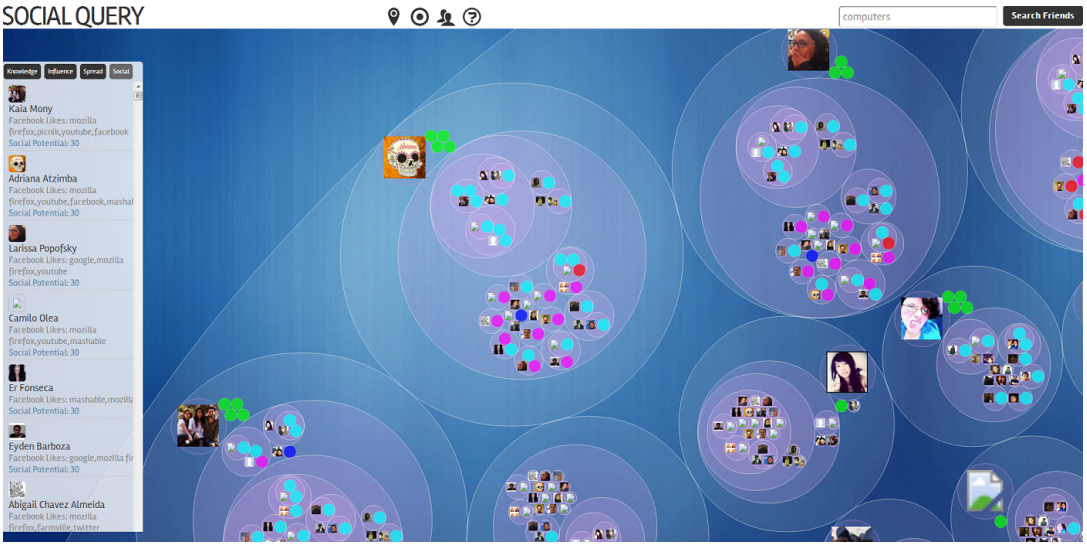

Figure 6.1: Screen shots of *Hax's* social spread interface which lets users view the social groups of their potentially interested audiences.

People get participation and action from others via postings in an online community. For example, they can create posts to request attention to interesting shared content [113]. By making posts in a large community, people potentially increment their number of responses. Yet, communication research has found that online users receive less replies when they share content with their entire network, than with a small targeted audience [10, 114]. Sociological theory on disclosures, also establishes that when a person feels signaled out due to her unique traits, the more likely the person is to be responsive [115].

Many proficient users sometimes use different online sharing mechanisms to engage in *selective sharing*, directing content to specific predefined audiences [11]. These users define collections of people, and then post content to each of these lists.

But, 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; the organizer of a social rally might only want to target those community members who are in town next weekend. In these cases, predefined collections of users might be too coarse or irrelevant.

Other techniques are selecting on the fly individuals to target, and only to them sharing the content or message. This type of behavior allows for a more dynamic selective sharing experience that is more context-driven. We will refer to this practice as *targeted sharing*.

But, finding the right people at the right time is hard. Especially in large communities where users may not have a notion of everyone's traits. Previous work in social recommenders have created list-based interfaces where the system recommended users with a certain expertise or skill set [116, 117]. These systems however, do not allow people to easily explore and compare the different characteristics of the recommended individuals. Yet, those characteristics can play an important role in people's decision for collaborating or interacting with someone [118].

We explore how data visualizations transform (or not) people's audience selection activity. Our belief is that data visualizations prompt people to learn about their audience. This empowers people to target audiences that capture their different sharing needs.

To explore these ideas, we designed *Hax* [1] — a tool that provides a query interface and many visualizations, to support users in dynamically choosing audiences for their targeted sharing tasks. Fig. 6.1 presents a screen shot of one of Hax's visualizations for targeting audiences. We set to study how users engaged with such a tool in the context of sharing and connecting with an audience in a Facebook group. Facebook designed groups to ease online community-building. We can consider each group to be an online community of its own [119].

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;

[1] Mayan for *exclusive*. This references it is a tool to select people for a task.

- 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.
- A better understanding of the way data visualizations transform users' audience selection activity.

6.1 Motivation

One of the challenges in identifying community audiences to collaborate or share content is the fact we 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 would like to share with interested folks a popular news piece she just read. These ever-changing needs affect in subtle ways who we want to exchange information with, or simply interact. 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 necessities.

We believe that the data modelling 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 data modelling algorithms are specialized at correctly categorizing data, they rarely fully capture humans' ever-changing decision process for selecting with whom to interact. We therefore opt to integrate data visualizations that help put humans in the loop, and let them make the final decision.

We designed different data visualizations highlighting specific social signals (traits) of relevant community members to aid users to select their audience. Our exploration begins with the three social signals listed below. We briefly explain 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 identified they played an important role in targeting audiences [112, 120]:

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 [120]. 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 [121]. However, location does play an important role when interacting and organizing events within the physical world [122] (e.g., a social rally,) as others' spatial-temporal constraints can determine how much a person will engage in the activity [123].

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 to recognize prospective newcomers, who can help keep the community alive and active [112]. 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.

6.2 Background and Related Work

Audience Targeting Practices

Traditionally, editors were in charge of publishing and distributing content [124]. Editors generally invested a great amount of resources in marketing consultants, who provide them with a clear upfront picture of who their best audience was [124]. The Internet has however, transformed this pattern. Anyone, can now author, share, and distribute content. However,in difference to editors, normal users typically don't have a concrete image of their audience, or even best ways of targeting them [125]. Yet targeting concrete audiences can be important for getting richer and faster responses from the public [126].

To overcome the lack of marketing knowledge, people rely on different cues to estimate the traits of their online audiences. However, few cues are available [125]. For example, a person might remember she friended her co-workers, and they are therefore now in her audience. Yet, it might still be unclear to the person exactly what these people, care about [13].

In this work, we explore how we can make audience cues more readily available for people, and we study the impact this can have on users' audience selection process.

Expert Search Tools

Hax, the tool we developed based on our qualitative study, 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. [127] 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, also taking their location into account. Similarly, ContactMap [128] visualizes contacts along with their attributes and location. In their algorithm, Chen et al. [129] add 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 [116, 130, 131, 126].

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

Facebook Graph Search

Facebook's *Graph Search* offers a natural language interface for searching one's social network. Queries may consider several social variables: an example query is *"TV shows liked by people who study computer science."* A search returns a ranked list of relevant Facebook users

with some of their characteristics — such as the city where they live, music they like, number of friends on the site, and others.

However, it is doubtful whether the design of Graph Search was influenced by the requirements of users employing targeted sharing. The attributes and interactions modes it supports — presenting a long list of matching users — are limited. The task specificity and the richer interaction modes of the tool presented in this work should make it more useful and accessible to our target group.

While these works provide means to find users to direct content to, none of them necessarily cover the needs of our target group of normal everyday users.

Facebook Advertisment Targeting Options

Facebook offers advertisers options for ensuring that their ad will reach a targeted relevant audience. Facebook allows advertisers to target audiences based on users' location, age, zodiacal sign, interest, education, as well as whether they have liked their particular product in the past, or their friends have.

Facebook's targeting options assume that the end-user has a crisp 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 community members who engage in targeted sharing.

Bernstein et al. [125] identified that social media users consistently underestimate the audience size for their posts, guessing that their audience is just 27% of its true size. It is therefore likely, that community members also will not have a clean-cut notion of the characteristics and traits of their most relevant audience for a given post. This leaves space for the creation of online tools that help end-users better visualize and understand their different audiences and their characteristics.

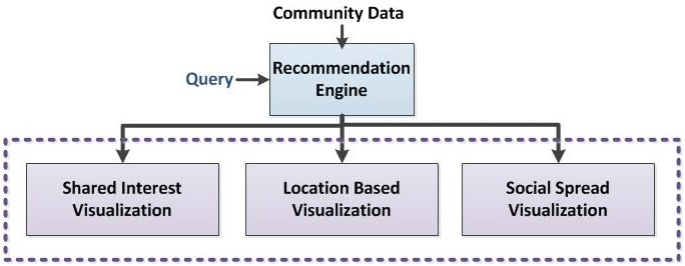

Figure 6.2: The components of the *Hax* system.

6.3 Designing Hax

Hax is a web-based tool that supports targeted sharing on Facebook. The user first enter a query that determines the topic they are interested in posting content about, such as *"women's rights"*. 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. At this point, the **visualization engine** provides three different presentations for the recommendations. Fig. 6.2 presents an overview of the Hax components. We now briefly describe these two modules.

6.3.1 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"* might have 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 Facebook's existing curated labeling to categorize interests was very general,

and did not enable an exploration of the data on various levels. To counter this effect, we use topic models [132] to model the community's shared interests.

Given the nature of the data, we used a labeled latent dirichlet allocation approach (labeled LDA) [132], similar to that proposed by Forbes et al. [133]. The discovered LDA topics correspond to the community's shared interests, labels correspond to the ones Facebook assigns to 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 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 [134].

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 interest is handed over to the visualization engine.

6.3.2 Visualization Engine

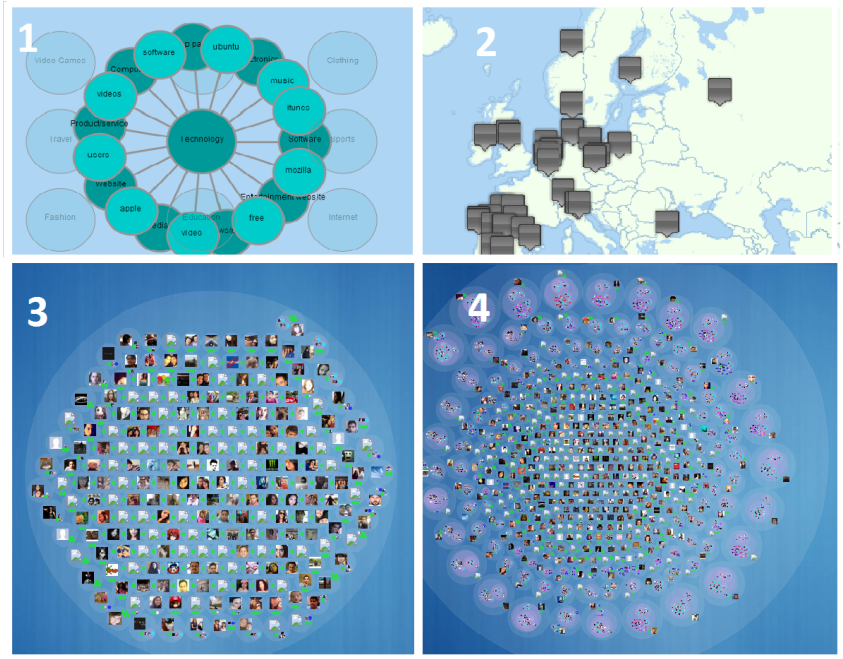

Figure 6.3: Overviews given by each visualization: (1) shared interest, (2) location-based interfaces; (3, 4) social spread interface.

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.

Our tool provides three different interactive visualizations, each emphasizing different social signals. Following the visualization mantra [135], 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 many times 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, they do not allow users to zoom in, and explore particular

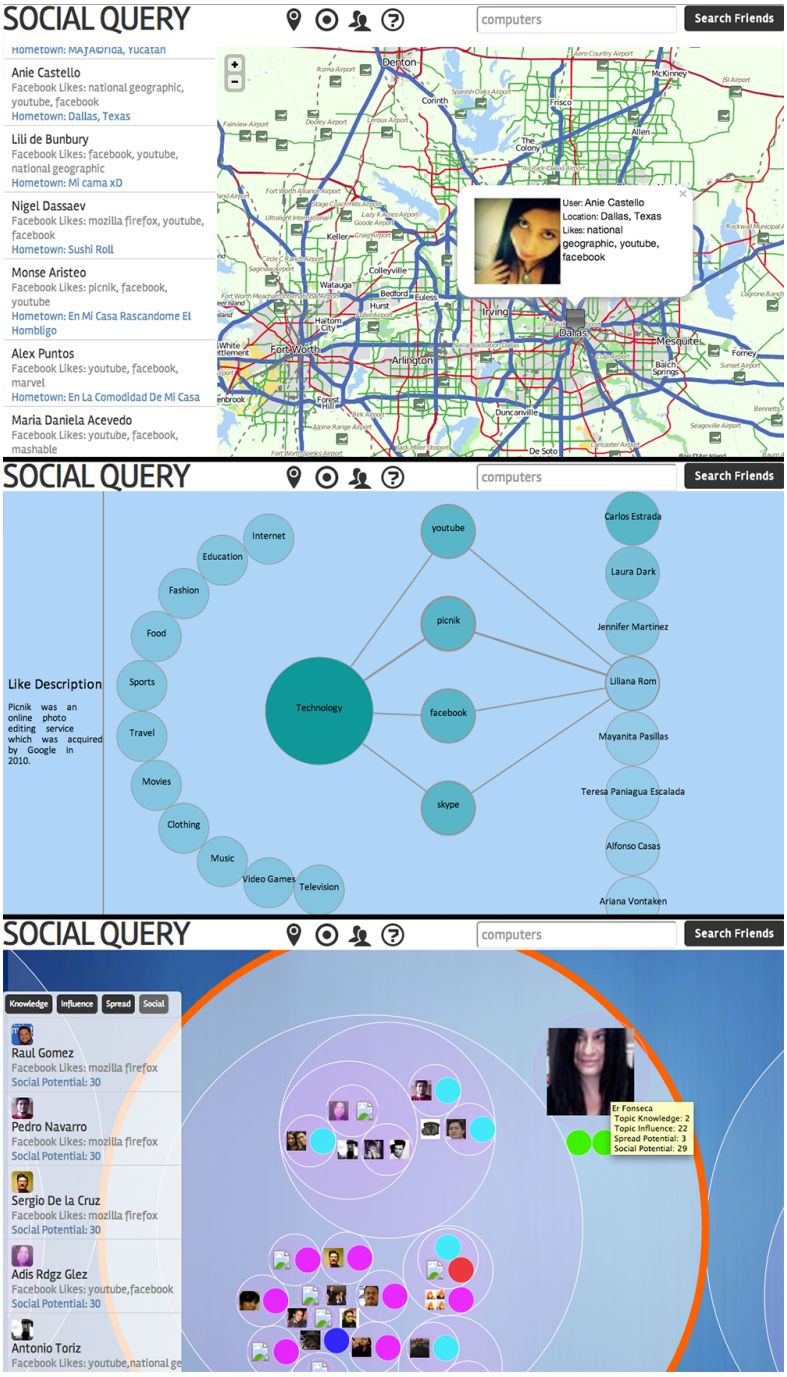

Figure 6.4: Screen shots of a zoomed-in version of the different visualizations in *Hax* (top to bottom): location-based, shared interest, and social spread interfaces.

aspects of the data.

We provide a short description of each view below. Fig. 6.3 presents the type of overviews

each interface provides. Fig. 6.4 shows a screen shot of all three visualizations with their zoomed-in view.

Shared Interest Interface Initially, the Shared Interest Interface presents an overview of all of the discovered shared interests of the community (Fig. 6.3.1). 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. 6.4, 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. By letting the user investigate the connections between members, Likes, labels, and shared interest, it moreover allows the user to explore the algorithm's rationale.

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 or rallies. In addition, location also provides a sense of cultural context.

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. 6.3.2).

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. 6.4, 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 [123] in their decision process. Knowing others' physical affordances is important, as it can influence their decisions for participating in an event [123].

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 by lending their social contacts.

From the recommendation engine, the social spread interface receives the list of recommended community members. 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. 6.3.4). 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. 6.33).

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.

From this view, users can thus quickly identify the overall type of social connections 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 design of the visualization, i.e., the structuring via a spiral, was inspired from the work of Katayoon et al. [136] In their research it was found that visualizations of hierarchical data, such as community memberssobad, ordered based on their relevant Likes and contacts, can become overcrowded. It can also be difficult to see details about specific nodes. The 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. Given the overwhelming amount of possible members of an online community, and the large amount of relevant contacts each member can have, space-saving designs become important.

6.4 Usability Inspection of Hax

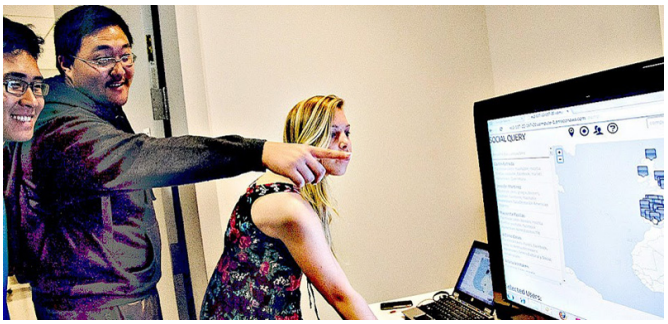

Figure 6.5: Hax at a university annual open exhibition which had hundreds of visitors.

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

Participants

Using Facebook's Group Browsers we first identified groups with large number of members, and asked the group administrators whether their group would be interested in participating. We contacted the administrators from 10 different groups, whose members (at least a certain percentage) appeared to be local to us based on information on their Facebook profile. 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 receive 10 USD for their time.

Procedure

During an hour user session each participant completed a series of targeted sharing tasks with Hax using laptops with an Internet connection we provided. We opted for participants to only conduct tasks with Hax because directly comparing with Facebook's native interface would be unfair as it 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, I observed them and 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.

To code responses, we read every questionnaire response, and identified key concepts about users' perspective 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.

6.5 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 massively 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 space 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 few spatial context.

Audience Diversity

This experience is about feeling that data visualizations bring diversity to one's targeted audience selection process. Participants reported 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 Likes' meanings, and analyse whether it made sense to include certain

candidates in their targeted audience: *"There were some brands [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 the system recommended.

Open Deployment of Hax

Hax was also installed on a large screen display for several hours in a well-attended university open exhibition to further explore how average users experience this open-ended way of selecting audience candidates (cf. Fig. 6.5.) Without prior notice or any instructions, 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 analysed and studied Hax without interacting. Average interactions times were around 1 min. Hax's visualizations appear to help people intuitively get an idea of what Hax is about and how to use it.

6.6 Outlook and Discussion

Our results show that users can target their audiences through data visualizations. Data visualizations prompted users to learn more about their peers. They also helped people find diverse audiences for their different sharing tasks. Something people felt was not facilitated with list based interfaces. This type of system design can help to have more cultural sensitivity, fostering better social interactions and collaborations. We see how data visualizations empower users to consider not only others' interests, but also their other traits, such as social, cultural, and spatial signals. This created a much more immerse and realistic 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 help users get diversity in their interactions. One threat facing society today is obtaining enough "strangeness." Technology will dictate the type of relationships we will have in the future. It is thus crucial 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 the main visualization goals.

Chapter 7

System for Assembling Action with Online Audiences

Activist groups usually have a small set of highly motivated core members who donate extensive amounts of their own time and resources to generate change in the world [138] e.g., to fight corruption. However, to achieve their goals, activist groups cannot rely solely on a small set of core members. They usually require a larger crowd of volunteers, e.g., people who can help report corruption across the globe [139]. Activist groups thus need the support of casual volunteers: a larger group of interested but, less committed individuals [140].

For years, activists went door-to-door to recruit and engage casual volunteers. Recently, activists have adopted new technologies to facilitate executing these tasks [141, 138]. Some use mailing lists to maintain continuous communication with their volunteers [142, 143]. Others are using social media [144]. Facebook has been particularly useful to recruit and coordinate volunteers, e.g., to call them to action to participate in fundraisers or demonstrations [145, 146]. Twitter has also enabled those who feel the desire to help [139]. Twitter has mobilized even those unaffected by the activists' cause or who live in distant regions [147].

This work was published in CSCW: ACM Conference on Computer-Supported Cooperative Work [137]. Special thanks to Andres Monroy-Hernandez for his valuable contributions to *Botivist*.

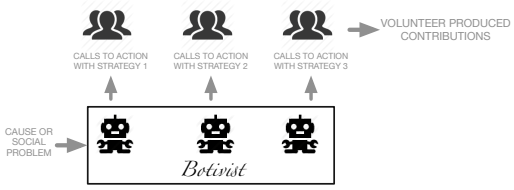

Figure 7.1: Overview of Botivist: Activists first provide the social problem for which they want a call to action. Botivist then tries different strategies to trigger contributions from volunteers.

But, despite the technological advancements, social computing technologies still do not facilitate the connecting of casual volunteers [145]. Many eager individuals receive little direction on how to help [139]. Current technologies also do not help activists mobilize people. Activist groups still have to spend time figuring out how they will present their campaigns in order to successfully trigger action [144].

Understanding how the presentation of a campaign (message) affects the engagement of volunteers and people in general has been extensively studied in both theories for civic engagement and marketing [148, 149, 150, 151, 152]. There has also been a growing interest in exploring how technology can be used to frame messages better and more persuasively [148]. Companies, organizations, and even governments have also taken their messages onto social media to reach a wider audience, and influence behavior [141]. Some have gone as far as to set up fake accounts operated by automated software (bots) in order to feign support and influence real users [153]. In Mexico's 2012 presidential election, the winning party used over 10,000 bots to swamp online discussion [154].

As we can see, several high ranking organizations have devised complex social media strategies to get their message across. However, most activist groups are still in the dark, trying to find the best strategies to mobilize people [145]. The main difficulty comes from the fact that the level of agency at which volunteers contribute to a cause online depends on how they perceive the efforts of activists [144]; but, social media has precisely transformed how people interpret activists' varying messages [142].

It can thus be very difficult for activist groups to keep up with new technology and to predict the outcomes of adopting a certain strategy. This complexity has forced many activists to limit who among their organization can use social media [145]. They usually prefer to have a "point person" in charge of the group's social media strategy. However, even when the point person finally "figures out" the best ways to trigger participation, it is not easy to transfer that knowledge to others [142]. Additionally, it can be costly or impossible for the person to present the group's campaigns differently to test what is the most effective. All of this hampers activists' success.

To help activists identify the best strategies, we present Botivist, a platform that, by leveraging online bots, allows activist groups to try different strategies for calling volunteers to action. Figure 7.1 presents an overview of Botivist. The group first presents the social issue for which they want to make a call to action, e.g., "corruption." The platform then tries different strategies to request contributions on social media, and help advance the group's plans. Activist groups thus receive help about how to prompt contributions from strangers, reducing the need to invest time in people who might never contribute.

The purpose of this paper is twofold. First, we aim to analyze the opportunities, limitations, and challenges of a system that uses online bots to call people to action in order to act upon social issues. Secondly, we analyze the online behaviours of those who are most likely to respond to activist bots. More specifically, (1) we focus on understanding how people in the wild respond to different message framing (strategies) when these come from an activist bot. Most previous work studied messages framed primarily when the information came from a company/organization/or individual and has not addressed how people respond when the messages are presented by an online bot, especially for the topic of activism. (2) In addition, an effort is made to understand the communication patterns of the people who decide to engage with Botivist. This analysis is important as these individuals could one day become core activists.

For this purpose, we designed and conducted experiments on Twitter. We designed Botivist as a Spanish speaking agent because we had more background with the language and its usage

in activism: I had direct access to activists from Mexico[1] who helped polish Botivist's messages. We focused our calls to action on impunity and corruption, which are some of the most pressing social issues fought in general by activists in Spanish speaking countries, especially those in Latin America [155, 156].

Our aim was for Botivist's calls to action to facilitate collaborations. This focus was chosen because activist groups usually need volunteers to work together towards a common goal. Botivist thus mentions three different individuals in a tweet and requests that they collaborate to fight a social problem. Note that these persons might not know each other, i.e., they can be strangers. Botivist can therefore also help people to meet and connect.

To assess and compare the reactions that Botivist triggers in volunteers, we use the evaluation methods of previous work, which studied how people react to different strategic presentations of a message [149, 150, 151, 152], as well as research which analyzed the quality of information from strangers on social media [157]. We deployed our platform publicly on Twitter, where our activist-designed and controlled bots invited people to organize around particular social issues using different strategies. 175 volunteers responded to Botivist's calls to action. These volunteers made 1,236 contributions (424 tweets, and 813 favorites and retweets). We found that Botivist's most effective strategy for prompting contributions from volunteers was to be upfront and direct. Interestingly, when Botivist used designed-to-be-persuasive techniques, known to be effective in direct human-to-human interaction, it received far fewer replies. At times these techniques even seemed to prompt individuals to discuss whether bots should help solve social problems. We found that by being less openly persuasive in its message and more direct, Botivist encouraged almost double the number of replies. Our results also indicate that those who decided to engage with Botivist, were individuals who consistently tweeted terms related to activism, politics, and also slightly marketing analytics.

Together, our results highlight the importance of understanding a community before simply

[1]We had continuous discussions with activists from Mexico's Revolutionary Action Group (GAR) who collaborate with activists' across the world to fight corruption and bring justice to workers `https://es.wikipedia.org/wiki/Grupo_de_Accion_Revolucionaria`

adopting persuasive technology, and expecting it to function. However, the benefit of Botivist is that by allowing activists to instantly probe different strategies, it might not be necessary to understand a community in detail. Our work shows the strength of this new type of civic media.

7.1 Related Work

7.1.1 Botivist: Strategically Framed Messages

There is extensive literature on how a message's framing (presentation) influences people's preferences for a product, or how much they participate in an event [149, 150]. Such research has uncovered the power that presenting something positively (benefits gained) or negatively (benefits lost) can have over people's opinions and levels of participation. e.g., people are more likely to buy a cream if its ad highlights its benefits of "helping them to look younger", than if the ad lacks such information. Related work has also identified the value of integrating particular messages into a call to action to motivate participation. Shen [158] identified that by integrating a solidarity component into the requests made to volunteers, volunteers completed their tasks more thoroughly.

"Persuasive computing" [148] has combined this literature with technology to demonstrate how computation can be used to influence people's behavior [141]. Many companies and organizations are now also using new technology, e.g., social media, and the psychological basis of how people make choices to try and shape the habits of their costumers [151], even their health choices [159]. Online platforms have also adopted persuasive computing to motivate more participation. Ling et al. [160] found that by simply reminding people of their uniqueness to an online community, people posted more. However, most persuasive computing research has paid little attention to designing technology to enhance activism. While there has been growing interest in creating systems to help activists coordinate volunteers [138, 161] most of these designs do not consider that by having different strategies for calling people to action,

individuals can be persuaded to participate more in a social cause.

7.1.2 Botivist: Online Automated Agents

Since the 60s there has been a proliferation of automated conversational agents [162], and it is now commonplace for humans to respond to them [163]. We designed Botivist considering that people respond to automated agents effortlessly and that the technology to develop such agents is now mainstream. Most previous work on online bots has focused more on developing techniques to detect these agents [153]. But the true potential of using bots to persuade and mobilize people on social media has not yet been fully understood, let alone for activism. Aiello et al. [164] began an exploration in the area by setting up a social experiment that examined how people interact with a bot in an online community. The work uncovered how a bot could acquire social relevance even when using simple canned responses and lacking profile information. Botivist builds on these findings to create online social agents that start to mobilize people to help activist groups.

7.2 Botivist

Botivist is a web application that activist groups can use to call volunteers to action to help address a social issue. We based Botivist on the well-known idea that by simply changing how a message is presented, it is possible to trigger diverse reactions [149, 150, 151, 152]. Therefore, by simply framing a call to action differently, we can prompt different amounts and kinds of contributions. However, identifying how to frame a call to action to obtain a desired turnout is not simple. Therefore Botivist probes different strategies, i.e., ways of presenting a call to action.

7.2.1 Botivist Strategies

We selected some of the most common strategies used by activists to call volunteers to action [165]. The strategies are also known to be some of the most effective to garner participation [158, 152]. We wrote the messages for each strategy, and then polished them with the help of real activists. The activists discussed the messages, and then corrected the language to produce content that they would use.

To try out each strategy, Botivist has a set of Twitter accounts to communicate with potential volunteers. All of the accounts identify themselves exactly the same. They mention they are bots (with a bio that states they are a bot; their profile picture is also of a bot.) All accounts also have at least 50 followers. The only main difference across accounts is the strategy used to communicate and call volunteers to action.

Direct Strategy

This strategy is about being upfront and direct when making petitions [166]. When in this strategy, Botivist directly calls people to action to find solutions for a social issue. For instance: *"Could we collaborate to brainstorm solutions to the problem of corruption?"*

Solidarity Strategy

This strategy posits that feelings of solidarity or empathy can drive people to respond to a call to action [165, 158]. When in this strategy, Botivist directly calls people to action, but it also shares solidarity quotes. For instance, a sample call to action is: *"Could we collaborate to brainstorm solutions to the problem of corruption?" "Remember, that: One for all, all for one!"* Note that in this strategy for each call to action or reply, two tweets are sent: one tweet for the actual reply (which mimics the replies and calls to action used under the direct strategy) and one tweet to share the solidarity quote.

Gain Strategy

This strategy is about presenting a calls to action in terms of the gain that people would receive if they participate [152]. These calls to action also mimic those used in the direct strategy, but add additional text to emphasize the gain. A sample call to action from this strategy: *"Could we collaborate to brainstorm solutions to the problem of corruption? We might improve our cities!"*

Loss Strategy

This strategy is about presenting a call to action in terms of the loss that people would receive if they do not participate [152]. A sample call to action from this strategy: *"Could we collaborate to brainstorm solutions to the problem of corruption? If not, our cities might suffer!"*

Our current design considers that activist groups would use Botivist especially at the start of their collective effort. Botivist thus focuses on prompting action in the form of discussions or plan proposals to help solve the social problem the activists are invested in (note that discussion and brainstorming are among the first steps of any collective effort.) All of Botivists' responses echo the "How Might We" questions which are known to launch brainstorming [167]. Depending on which strategy Botivist is pursuing, it will have a slightly different follow-up question. For instance, when following the gain strategy, Botivist asks for ideas, and then states that by participating in the brainstorming people could help improve their cities; whereas when following the solidarity strategy, Botivist will ask for ideas, and share a solidarity quote at the end. Table 7.1 presents the list of follow-up questions for the direct strategy. The follow-up questions of the other strategies are exactly the same, except for the added text. For instance, the gain strategy had the question: *How do we fight corruption in our cities & thus improve them?* All calls to action and followup messages for brainstorming were verified and polished by real activists. A strategy has one particular call to action associated and a list of replies to pursue brainstorming. Botivist randomizes its replies. Botivist only targets a person once, and

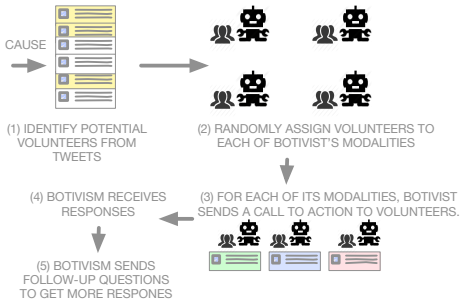

Figure 7.2: Botivist's process for calling volunteers to action and obtaining contributions.

engages further only if it receives a response.

7.2.2 Botivst Operation

An activist group first defines the social problem for which they seek volunteers, e.g., "corruption". Botivist identifies a set of potential volunteers and targets them by mentioning them in a tweet. Botivist then waits for replies, in which case it uses the opportunity to request more contributions. Botivist tries to target the same number of volunteers for each strategy. Figure 7.2 shows Botivist's work flow.

To determine who will be mentioned in a tweet, we take into account the findings from previous work on asking questions to strangers on social media. Nichols et al. [157] showed that obtaining responses to questions from strangers usually had a low response rate, even under favorable conditions. We thus designed Botivist to maximize the chances that people would respond to the call to action. Botivist targets individuals who have just publicly tweeted something related to the activists' social problem of interest (based on simple keyword matching). For instance, if the activist group wants volunteers to fight corruption, the system will target people who have just publicly shared a tweet explicitly mentioning corruption. Group managers can provide a list of relevant keyword or hashtags to the system. Note that this design follows also how activists operate online where they search keywords and make publicly addressed tweets to seek a wide range of help and build new social connections [139]. Similar

to real activists, Botivist engages people across countries, and has no prejudice against people based on their initial support or opposition to the social issue at hand [147].

Once a relevant tweet is detected, its author is randomly assigned to one of Botivist's strategies. When a strategy has three authors (persons) assigned, Botivist sends a strategic tweet that mentions and calls the individuals to action to work together on the issue. Note that Botivist mentions multiple users in a tweet to facilitate cooperation among possible volunteers. Given Twitter's space constraints, we considered that three users was the maximum which Botivist could mention in one tweet. Botivist is also constantly monitoring Twitter for responses to its calls to action. Botivist replies with canned responses to those tweets. The canned responses depend on the strategy which Botivist used to communicate with the group. The follow-up messages all focus on prompting action from the volunteers.

Q#	**Question Text**
	Direct Strategy
1	How do we fight corruption in our cities?
2	How do we fight corruption in our countries?
3	How do we use Twitter to fight corruption?
4	How do we use the people to fight corruption?
5	What should we change personally to fight corruption?
6	What should we reduce to fight corruption?
7	What should we change at home to fight corruption?

Table 7.1: Questions asked in the direct strategy.

7.3 Evaluation

This paper hypothesizes that online conversational bots can be used to call volunteers to action to work on solving a social issue. We focus our evaluation on the two main aspects of this claim. First, for a given social problem, is it possible to use bots to call volunteers to action? Can Botivist obtain responses from strangers on social media, and what solicitation strategies are most conducive for obtaining results? Second, can Botivist prompt these volunteers to produce relevant contributions? Does Botivist trigger volunteers to discuss the issue or to

propose plans? Similar to [157], our goal is to shed light on the type of contributions which autonomous agents can assemble in short bursts of times to get more immediate usable results.

To answer these questions, we used Botivist to publicly call volunteers to action regarding the social problems of impunity and corruption [155, 156].

We were able to attract 175 volunteers over the course of two days. Note that via this approach Botivist likely called to action people from different countries, similar to how real activists function [147].

7.3.1 Method

We set Botivist into operation April 26th and May 7th 2015. We used the keywords "corrupcion" and "impunidad" (Spanish for corruption and impunity) to call volunteers to action. For both social problems, Botivist prepared its calls to action and follow-up questions for each of its strategies. Botivist alternated between social problems. For instance, first it would call on people to fight corruption; and subsequently call on other individuals to fight impunity. We tried for Botivist to make its calls to action for all its strategies almost at the same time. However, we randomized the delay time between Botivist's calls to action to avoid being labeled as spam by Twitter. We also constantly monitored Botivist's Twitter accounts to ensure they were still running and not blacklisted. All tweets received and sent by our system were collected in a database for further analysis.

7.3.2 Results

Botivist called to action 376 groups of three: 94 groups per strategy (47 for corruption; 47 for impunity). In total, Botivist received reply from 175 volunteers. These volunteers made 1,236 contributions (424 tweets; 320 retweets and favorites to Botivist's content; 493 retweets and favorites to the content of other volunteers.) Table 7.2 shows a breakdown of the replies and interactions of Botivist and volunteers. Figure 7.3 shows an overview of how people responded to Botivist's calls to action. The number of participants varied across strategies, as well as their

	Botivist	Direct	Loss	Gain	Solidarity
Calls to Action	376	94	94	94	94
Followup Questions	557	158	80	79	120
Volunteers	175	94	31	27	23
Volunteer Replies	423	204	53	74	92
Reply Rate	45 %	81%	30%	43%	21%
Interactions Bot	320	90	48	57	250
Interactions Volunteers	493	274	71	85	62

Table 7.2: Summary of Botivist's and volunteers' replies and interactions. Botivist's column shows the aggregated results across strategies

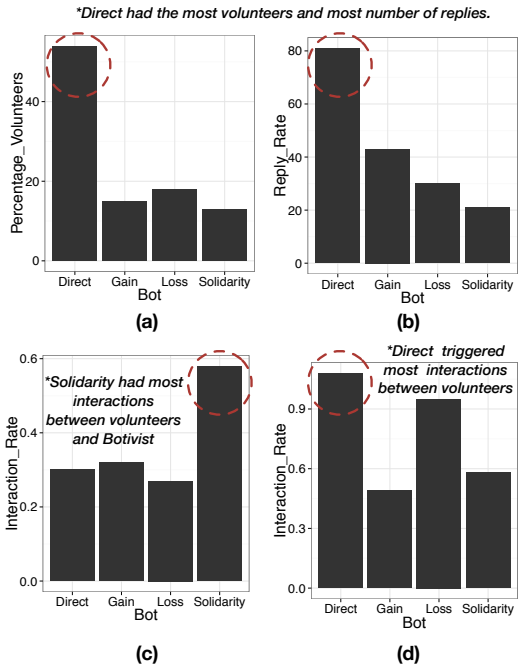

Figure 7.3: Overview of the number of volunteers and contributions which each strategy triggered. The direct strategy was the one that had the highest number of participants and triggered the most contributions.

contributions.

Figure 7.3a shows the percentage of people who responded to each strategy. There was a significant difference between number of unique contributors across strategies (a ANOVA test gave $[F(3, 174) = 38.94,\ p < 0.001]$). The direct strategy was the most effective to mobilize unique contributors. Over 50% of Botivists' participants were gathered via this strategy. By being direct and upfront, Botivist was able to obtain contributions from over 30% of the people it targeted. Compare that with a survey finding, reporting that about only 21% of social media users have contributed to groups involved in political/social issues [168]. Shepherding potential volunteers might thus actually yield more civic participation. However, it does appear to be important to know how to guide the volunteers. Botivist's Solidarity strategy was able to trigger participation from only 8% of the volunteers it originally targeted. With inadequate strategies, people's civic participation is thus likely to be inhibited.

Figure 7.3b presents the number of replies per Botivist tweet (reply rate). Note that the reply rate is a common measure to study the participation of online audiences, and has also recently been used to study how online audiences interact with advocacy organizations [169]. We found that there was a significant difference between number of replies across strategies (a ANOVA test gave $[F(3, 932) = 8.594, p-value < 0.001]$). Direct strategy had the highest reply rate (81%), receiving almost one reply per tweet sent out.

Although the direct strategy reply rate appears high, it is lower than that of crowdsourcing companies such as Innocentive,[2] which lets people pose problems for anyone to solve and gives cash awards for the best proposals. Innocentive receives approximately 20 replies (proposals) per posed problem [170]. However, a stark difference between this platform and Botivist is their monetary compensation. While Botivist relies purely on volunteers, prizes for the best proposals on Innocentive can range from $5,000 to $1 million dollars. We think there is value in exploring non-monetary ways to generate discussion and proposals from the crowd, especially for activist groups, which have limited resources.

[2] `https://www.innocentive.com/`

The reply rate of Botivist's other strategies was not as high as that of the direct one. The gain and loss strategies had results similar to asking people for product reviews [157], while the reply rate for the solidarity strategy was even lower (21%).

Figure 7.3c shows for each strategy the number of retweets and favorites Botivist received per tweet it sent (interaction rate). There was a significant difference between number of retweets and favorites across strategies (a ANOVA test gave $[F(3, 932) = 3.045,\ p < 0.01]$).

The strategy of Solidarity, particularly the quotes shared under it, received the most retweets and favorites from volunteers. The rest of the strategies received almost the same amount of interactions. However, despite its high interaction rate, the Solidarity strategy received the lowest number of replies from volunteers. Interestingly, this strategy appears to prompt volunteers to retweet and favorite content; but not necessarily to make contributions in the form of replies.

Figure 7.3d shows the interaction rate per volunteer tweet, i.e., how much volunteers retweet or favorite each others' content. There was a significant difference between volunteer interaction rate across strategies (a ANOVA test gave $[F(3, 492) = 19.34,\ p < 0.001]$). Here the Direct and Loss strategies were the ones who harvested the most interactions among volunteers. We also see that in general people retweeted or favorited volunteers' content more than the content of an autonomous agent (i.e., of Botivist.) This result resembles what was observed in the Scratch online community, where people valued interactions with humans more than with computers [15].

7.3.3 Method: Analysis of Volunteer Contributions

Our goal was to understand the type of volunteer contributions that systems like Botivist can trigger. We were particularly interested in analyzing whether Botivist was able to prompt people to discuss and provide proposals for a given social problem. We used Upwork[3] to hire three Spanish-speaking, college educated individuals to independently classify Botivist's volunteers

[3]Upwork is an online platform to contract Freelancers to perform tasks online. Upwork is available at `www.upwork.com/`

into people who contributed proposals/discussions relevant to a social problem, and people who contributed to off-topic conversations. We instructed Upworkers to classify volunteers instead of single tweets because independent tweets can be harder to interpret. First, two Upworkers read through the 423 tweets of the 175 volunteers, and classified each volunteer into either on-topic volunteer or off-topic volunteer based on whether the person contributed proposals or discussion on the assigned social problem. The two coders agreed on the classification of 141 volunteers (Cohens kappa: 62: Substantial agreement). We subsequently asked a third coder to categorize the disagreement. We used a "majority rule" to categorize such volunteers.

7.3.4 Results

We found that the vast majority of volunteers (81%) responded to Botivist's call to action with discussions and proposals relevant to the social problem assigned to them. A sample of these proposals includes: *"Corruption isn't fought with street rallies! It's fought by being tough on crime, having honesty & transparency!"*

Only 19% of all of Botivist's volunteers contributed solely off-topic discussions. The word most used in these off-topic discussions was: *"bots."* A sample off-topic tweet: *"Sorry ... I can't collaborate with bots. I have a cultural bias."*

Strategy	Percentage On-Topic Volunteers
Botivist (all strategies)	81%
Loss	74%
Gain	89%
Solidarity	82%
Direct	94%

Table 7.3: Overview per stratgey of its percentage of on-topic volunteers. The Direct strategy had the most on-Topic volunteers.

To understand the type of actions promoted by each strategy, we study, per strategy, the percentage of volunteers who responded to the call to action explicitly with what was requested. Table 7.3 presents, per strategy, the percentage of volunteers who contributed proposals or discussions relevant to their assigned social problem.

The Loss strategy had the largest percentage of volunteers who responded to the call to action with off-topic replies. Over 60% of these replies used the term "bot." Upon manual inspection it appeared that the replies were mainly questioning whether autonomous agents should help solve social problems. The Direct strategy, on the other hand, had almost all of its volunteers dedicated to discussing/brainstorming solutions to their assigned social problem.

In lab settings [152], the Loss strategy were more effective than the Direct strategy. However, our results indicate the contrary. This might help to highlight the importance of understanding how messages change when adopted by technology. In this case on Twitter, the messages used in the Loss strategy appeared to have triggered people to question the participation of bots in activism. Bots with an obvious agenda might be perceived as suspect and unwanted, whereas people with an agenda would be more easily tolerated and even respected.

Category	Sample Hashtag in Tweet	Sample Twitter & Profile Description
Activism: covers hashtags fighting for a social cause, or twitter handles of people who appear to participate in activism.	*"***#WeAreAllAyotzinapa** *We continue to fight against impunity. We won't stop, let's take this to the streets!"*	**@antireforma**: *"REVO... REVOLUTION.! DAUGHTER OF THIS EARTH! LOVER OF LIFE! I FIGHT FOR TRUTH, JUSTICE AND VENEZUELA!!!"*
Politics: covers twitter handles of people who appear to officially be part of the political system or hashtags about political figures	*"Bronco & Fernando Elizondo (Mexican political candidates) will make history in Nuevo Leon* **#TweetForElBronco***"*	**@epn**:*"President of the United Mexican States."*
News: covers hashtags used to give news alerts, or twitter handles of people who appear to be news reporters.	*"Car crash in Central Street, 2 people hurt.* **#Verfollow***"*	**@epigmenioibarra**:*"I am a news reporter and producer."*
Marketing Analytics: covers hashtags used to analyze one's Twitter followers to better market oneself or Twitter handles of people that give advice on how to better market oneself.	*"Over 5% of my followes have retweeted me in the last 24 hrs and 20% in the last month* **#usefulTweet***"*	**@umspromotions**: *"Nigeria's top-class Promo-Marketing media outfit: Digital/Social Media Mgmt, Radio/TV Plugin, DJ Mixes, Strategic PR."*

Table 7.4: Description and examples of the categories that describe the groups' distinguished words.

7.3.5 Method: Communication Patterns of Volunteers

Our aim is to understand whether those who respond to Botivist's strategies communicate on social media differently than those who do not. It might be that people who discuss certain

social issues are more prone to interact with activist bots. We thus divide the people whom Botivist called to action into two groups: those who responded and those who did not. We study the tweets from the people in each group, and identify the distinguishing type of words they use.

Distinguishable Types of Words. To obtain a descriptive assessment of what distinguishes each group of people over the other, we take the 200 tweets of each group member right before Botivist called them to action (we take the data prior to Botivist's interaction as our system might have influenced people's future conversations [164].) From these tweets we then: (1) identify the words that differentiate each group; (2) categorize these words to understand their nature.

To first identify each group's distinguished words, we use a Mann-Whitney ranks test [171]. This measure highlights the words which have consistently been used more by one group than the other. Here we consider a document to be the tweets of an individual group member, and the corpus to be the collection of tweets of the entire group. For each corpus, we calculate each word's Mann-Whitney ρ measure (indicating how frequently the group has used that word, and also weighting higher the words which the group uses and the other group does not). Note that the ρ measure allows us to have corpora of different sizes. We then rank the words of each corpus based on their Mann-Whitney ρ score, and identify the 1% of the words with the highest score. We consider these represent the group's distinguishable words. In this case, we found that the most distinguishable words for both groups were hashtags and mentions to other Twitter users.

Next, we characterize each group's distinguished words (hashtags and users mentioned) via qualitative content coding. For each group, we first read a set of 200 randomly selected tweets mentioning one or more of the distinguished words. This helped us understand the context in which the hashtags were used and people referenced. We also looked up the Twitter profile of the users the groups mentioned, reading especially their profile description to understand who they were. We then began to extract categories describing the distinguished words of

both groups. We then analyzed the emerging categories and helped to adjust them. Finally, we looked at a set of 400 randomly selected tweets with distinct words from both groups and produced a final list of mutually exclusive categories. We then hired three Spanish-speaking, college educated colleagues, unrelated to the work, to categorize the distinguished words of each group. For each word, we showed the tweets in which the word was referenced, and in case the word was a twitter handle we showed also the user's profile description. We asked first the two coders to categorize each of the 155 words using the categories listed in Table 7.4. We asked the coders to pick the "most relevant" category for each word given its related information. The two coders agreed on 143 words (Cohens kappa: .79). We then asked the third coder to label the 12 words upon which the first two coders had disagreed. We then used a "majority rule" approach to categorize those remaining words.

7.3.6 Results

We found that across strategies the people who responded to Botivist tended to consistently use more hashtags related to social causes, and referenced activists in their tweets. Over 50% of the distinct words of this group were hashtags and user mentions related to activism. For instance the hashtag *#porEsoPropongo*[4] (*#that'sWhyIPropose* in English) was one of the most tweeted by the people who replied to Botivist. The hashtag has been adopted by citizens in Latin America to make proposals for fighting corruption and impunity globally and locally. Note that in addition to activism, these individuals also referenced frequently politicians and news reporters (25% of their distinct words belonged to news and 19% to politics). However, they referenced these categories to a smaller degree than the people who decided to not engage with Botivist. The latter tweeted more consistently about high profile politicians and news media (43% of all their distinct words were about politics and 40% news.) For instance, they frequently referenced the twitter handle of possible Argentine presidential candidate Jose M. de la Sota (@delasotaok) or Mexico' president (@epn). They rarely used hashtags about social

[4] `http://www.poresopropongo.mx/`

causes (only 6% of their distinct words).

What was particularly interesting is that those who did reply to Botivist also mentioned somewhat consistently hashtags and users related to marketing analytics (which is about measuring and analyzing how audiences react to one's content to maximize effectiveness when communicating online). For instance, they mentioned frequently the account @m_g_w_v, which describes itself as a guide to help others market themselves to obtain more followers, and retweeters.

Note however, that the marketing analytics hashtags accounted to only 4% of this group's distinct words (compared to less than 1% for those who did not reply.) Overall, the people who replied to Botivist mainly used hashtags related to activism. Table 7.5 presents for each group and strategy its top-3 words with highest Mann-Whitney ρ (most distinct words).

Strategy	Words People Who Reply	Words People Who Don't Reply
Loss	@antireforma (activism)	@epn (politics)
	#poresopropongo (activism)	@delasotaok (politics)
	@umspromotions (marketing)	@sanchezcastejon (politics)
Gain	#cambioclimatico (activism)	@albert_rivera (politics)
	@jaimerdznl (politics)	@alvarouribevel (politics)
	@latati2 (activism)	@eldiarioes (news)
Solidarity	#blacklivesmatter (activism)	@juanorlandoh (politics)
	#ficrea (activism)	@presidenciamx (politics)
	@maracayactiva (news)	#radarparlamentario (news)
Direct	#handicapped (activism)	@ccifuentes (politics)
	@aristeguionline (news)	@periodicovzlano (news)
	#mgwv (marketing)	@sumariumcom (news)

Table 7.5: Top 3 Mann-Whitney words found consistently in the tweets of people who replied to Botivst & those who did not reply.

7.4 Discussion

Our experiments demonstrate the potential of using online bots to call volunteers to action for a social issue. The majority of the people called to action by Botivist made relevant contributions, and even started to interact with each other to further drive the collaboration. Our

study provides insight into the deployment of platforms that use online bots to call volunteers to action for activism, as well as demonstrates their feasibility (in terms of the contributions made by volunteers).

We discuss next our findings, covering in depth the advantages and limitations of systems like Botivist. We also initiate a discussion on the design implications of future systems that wish to use automated agents to call volunteers to action.

Feasibility

Botivist's strategy of being direct yielded a reply rate of over 80%. This strategy also produced a higher percentage of relevant responses (94%). This high response rate and promising collection of relevant information showcases the feasibility of using online bots to help activist call to action new potential members. Our results also highlight how online bots enable community scaffolding around social issues.

Our results show that the way a call to action is framed makes a difference in terms of how much people participate. This matches HCI findings on how the way a computer presents its voice influences people's replies to the computer [163]. But, what was surprising was that strategies which social psycology had identified to work well when done by humans were less effective when done by Botivist. This is meaningful as social psychology findings are usually relevant to how humans respond to automated agents [163].

Integrating solidarity quotes has long been a useful technique that marketing experts use to increment preference for items [158], and which activists have also adopted to prompt people to action [165]. This strategy has also been effective online. For instance, advocacy citizen journalists in the drug war added a solidarity component to their news reports to ensure participation to their calls to action [169]. Political activists in Palestine also shared solidarity messages on Facebook to mobilize [146]. But, when adopted by Botivist, the technique resulted in the lowest number of responses. Similarly, the loss strategy is known to trigger people to produce more quality work [152]. In the 2011 Egyptian uprising activists on social media used the technique (they stated how the lack of action would result in more corruption) to drive

more effective participation [172]. But, when used by Botivist, the strategy resulted in the lowest number of relevant replies, and seemed to make people actually question the role of bots in social issues.

For bots to be effective it is necessary to understand the communities in which they are deployed. We identified that the people who responded to Botivist were those immersed in online activism and marketing. They mentioned e.g., hashtags related to social causes and marketing analytics. It is likely that people linked Botivist to online marketing schemes. Therefore, the individuals who responded to Botivist were those who already found it natural to engage with marketing agents. However, here the persuasiveness of the strategy also influenced people's contributions. People respond better to activist bots that do not appear to have a "secrete agenda" and are upfront on their requests.

The proliferation of persuasive technology in social media [153] might have changed how people react to certain strategies. It likely made people have a preference for less "manipualtive" tactics. On the other hand, the fact that by simpling changing a bot's message varied how much people contributed, without needing any interface modification or design of complex bot behavior, suggests that designing adequate dialog for online automated agents can have a strong effect on increasing volunteer participation in activism.

Generalizability

Botivist can extend beyond the social issues focused on here. There is opportunity to use Botivist to guide citizens to devise solutions that improve education, health care. City planners could also use systems like Botivist to call to action experts or minorities to design novel more inclusive urban spaces.

Bootstrapping

Previous volunteer systems have identified the difficulties in finding participants [161]. Botivist shows the power of bootstrapping onto existing networks of people to call volunteers to action using automated agents. However, if such automated social actors were to become popular, a large number of these bots could then be contesting for people's time and attention.

It could be useful for social media platform developers to consider enabling functionality for these bots via APIs. For instance, the bots could obtain knowledge of the individuals who have already been targeted by other agents to eliminate duplicating volunteering requests, and also reduce user overload. Additionally, if a strategy is known to be ineffective, the system could recommend the bot to adopt other.

Maintainability

Our results show that a reasonable fraction of people were willing to make relevant contributions to Botivist's causes for free. This is meaningful because a key problem in maintaining systems that rely on crowds is their costs.

However, the long-term success of systems like Botivist might require automated social agents that can continuously engage the volunteers they have recruited. Such systems should consider in their design that while they do need to maintain contact with volunteers they should also not ping them too frequently to become an annoyance.

Additionally, by mentioning different individuals in a tweet Botivist enabled strangers to meet and start collaborating. Future work could also analyze how to design for long-term cooperation and even maintaining lasting relationships among volunteers. In particular, future work could study designs that: follow up with volunteers about their experiences; recognize them for their work; help them to feel part of the activist group, and build a sense of community [145].

On the other hand, some activists follow "syntax rules" when writing tweets to ease coordination [139], Botivist could also motivate standardized contributions that use a certain format or hashtag. This might also help long-term collaborations.

7.5 Conclusions and Future Work

In this paper, we introduced Botivist, a system that uses online bots on social media to call volunteers to action for social issues. The system adopts the bots as an opportunity to instantly probe different strategies for calling volunteers to action. People on Twitter in

general responded to Botivist's calls to action and provided relevant contributions. Volunteers also began to interact with each other, showing the potential of using bots to scaffold collective efforts.

We found that the strategy which Botivist uses to call volunteers to action does matter. Tactics known to be effective when used by humans, were not as effective when adopted by online bots. The strategy of being direct was the most effective for getting relevant participation. We believe that the growing number of persuasive online agents [153, 154] has made people prefer interactions with less "manipualtive" approaches. Nevertheless, our experiments highlight the real-world feasibility of using online bots to call volunteers to action for activism.

Note that the insights of this work are limited by the methodology used and population studied. For example, the people targeted by Botivist were all Spanish speakers, likely with their own cultural biases towards automated accounts (especially given the recent accusations of Latin American governments using online bots to change political online discussions [141]). Our results might therefore not generalize across populations. Future work could study how people across the globe engage with bots to better understand the phenomenon. Our methods also focused on breadth instead of depth. Future work could conduct longitudinal studies and engage in in-depth interviews with volunteers.

Chapter 8

Conclusions

This thesis provides an analysis of the relationships, tensions, and dynamics emerging between content producers and audiences in this new digital media model. While previous work [3, 2] focused its analysis to more friendly scenarios (e.g., within friends or in fandoms), this thesis uses different case studies to bring a more detailed understanding in a spectrum of online spaces.

Through this analysis, I identify that across spaces, audiences and content producers interact with each other in a sharing economy, that is based on giving, receiving, and reciprocating. Within this gift economy context, producers and audiences engage to produce collective efforts that benefit both parties. To succeed, content producers must then inspire their audience to give back, and ultimately help them attain their goals.

My research identifies that current tools to engage online audiences have several limitations, which emanate from a misconceived model of how audiences and content producers interact. Such tools adopt the traditional media consumption model, and assume that audiences and content producers interact within a market exchange instead of a gift exchange. As a result, such tools focus on profit rather than enabling reciprocity or collaborations.

I use the limitations I identified to introduce systems for engaging online audiences within a gift economy. My systems focus on facilitating cooperation and relationships of mutual benefit among content producers and online audiences. My systems consider two different

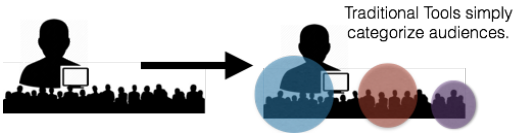

Figure 8.1: As described in chapter 1, traditional tools for engaging with online audiences usually are all adaptations of traditional marketing tools wherein audiences are surveyed and categorized anonymously and content authors have minimal interaction with them. These tools limit the type of collaborations produced.

design approaches: one in which authors can visualize relevant online audiences, and use the understanding to better engage with audiences; and autonomous methods which probe different strategies to motivate cooperation from an audience.

To conclude, this chapter will review the main contributions of the thesis and consider important avenues for future work.

8.1 Summary of Contributions

This dissertation has presented both an analysis of how audiences and content producers interact and the design and evaluation of online audience engagement systems. The knowledge contributed through this dissertation allow us to identify pain points in current tools for engaging online audiences, and to design novel systems that overcome such limitations. These systems open a design space of deployable applications that allow content producers to better cooperate with their audience. This work is an important first step in demonstrating how we can move beyond traditional tools (see Fig 8.1) that simply survey and categorize audiences, to novel collaborative systems that allow authors to reach audiences and directly start to collaborate to fulfill different goals (see Fig 8.2).

Hax (Chapter 6) opened the design space of multi-facteted data visualizations for engaging online audiences by demonstrating how authors could use such visualizations to understand their online audience, and then use the knowledge to motivate contributions from the audiences themselves. Hax enables content producers to consider and understand multiple variables, both

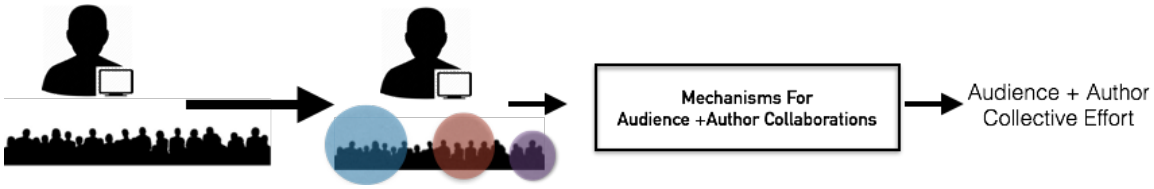

Figure 8.2: The systems proposed in this thesis add an extra element into the design space: audiences and authors can complete collective efforts which they could not perform efficiently alone. There are many under-explored areas of this design space, especially those that enable authors to drive into the intelligence of an audience and use it for more creative and complex tasks.

offline and online, to motivate and initiate collaborations with their audience.

Botivist (Chapter 7) demonstrates the power of using automated methods to recruit and engage audiences for an author's cause. Botivist opens the door for autonomous systems that can coordinate large audiences for a collective effort. Botivist shows how authors can recruit participating audiences and initiate collaborations with them with help of autonomous agents that can employ different recruitment strategies.

This thesis makes contributions in the area of social computing:

1. Analysis. Systems for cooperating with online audiences enable a new class of applications that give authors direct access to relevant and specialized online audiences with whom they can take-on multiple collective endeavours.

2. Concepts. The thesis provides an understanding of how content producers and online audiences operate within a gift economy. However current social media embodiment have several limitations that hinder reciprocity and cooperation.

3. The thesis offers an informed model of how to overcome these limitations and design the social computing systems of the future. Tools for cooperating with online audiences change the interactions and dynamics between content producers and online audiences. These interactions create novel ecosystems and power dynamics that were nonexistent before.

8.2 Outlook

Through our analysis, we observed that authors and online audiences have adopted social media not only to share relevant content, but also to harvest collaborations with each other. These collaborations range from helping each other to improve personally, to improving their community.

While existing social computing systems are powerful tools to organize collective action, their current embodiment have serious limitations. Most systems only support a particular aspect of the collective effort, e.g., tools to brainstorm a plan. There are few examples of systems that support a collective action from beginning to end, i.e., where a crowd comes together to formulate goals, brainstorm plans, and mobilize a critical mass of participants to take action.

For the future, I am interested in creating platforms that by leveraging ever-growing accumulations of big data can: 1) model people's behavior to create connections for collaborations and; 2) pipeline and dispatch large tasks to advance a collective goal. In such systems, a person could start out by presenting a neighborhood problem, e.g., her neighborhood receiving too much negative publicity in the news. The platform would then use social media traces and open data to recruit, connect, and encourage discussions among concerned citizens to help them find a solution and action plan for the problem. These systems could also incorporate data visualizations to help participants understand the problem and solutions. Perhaps, through the platform citizens would come up with the possible solution of organizing poetry events to showcase another side of the neighborhood. The platform would then pipeline the tasks needed to carry out this plan, such as identifying and recruiting promising artists, as well as designers to create advertisement materials, and audiences to attend the event. In other words, I want to explore designing end-to-end systems that use big data to connect citizens, and help them execute their plans to impact society. Such systems enable not only collaborations to improve a town, but also collaborations around healthcare, art, education, etc. I believe that the analyses, concepts, and implemented tools presented through this dissertation work demonstrate a

possible path to such a vision.

In summary, my research utilizes social media data to understand the collective action of media producers and participating audiences and to design innovative systems that hopefully can contribute to allowing societies to self-organize to their own and everyone's benefit.

Bibliography

[1] T. Harrison, *Audience interactivity and participation*, .

[2] H. Jenkins, S. Ford, and J. Green, *Spreadable media: Creating value and meaning in a networked culture*. NYU Press, 2013.

[3] N. K. Baym and R. Burnett, *Amateur experts international fan labour in swedish independent music*, *International Journal of Cultural Studies* **12** (2009), no. 5 433–449.

[4] D. Zeng, F.-Y. Wang, and K. M. Carley, *Guest editors' introduction: social computing*, *Intelligent Systems, IEEE* **22** (2007), no. 5 20–22.

[5] S. Savage, A. Monroy-Hernandez, K. Bhattacharjee, and T. Hollerer, *Tag me maybe: Perceptions of targeted sharing on facebook*, *ACM 2015 International Conference in Hypertext and Social Media* (2015).

[6] N. S. Savage, R. H. Wouhaybi, L. Nachman, and S. Sahay, *Online social persona management*, June 27, 2013. US Patent App. 14/232,706.

[7] S. Savage and A. Monroy-Hernández, *Participatory militias: An analysis of an armed movement's online audience*, in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pp. 724–733, ACM, 2015.

[8] S. Savage, A. Forbes, C. Toxtli, G. McKenzie, S. Desai, and T. Höllerer, *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.

[9] S. Savage, A. Monroy-Hernandez, K. Bhattacharjee, and T. Höllerer, *Tag me maybe: Perceptions of public targeted sharing on facebook*, in *Proceedings of the 26th ACM Conference on Hypertext & Social Media*, pp. 299–303, ACM, 2015.

[10] M. S. Bernstein, A. Marcus, D. R. Karger, and R. C. Miller, *Enhancing directed content sharing on the web*, in *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, CHI '10, pp. 971–980, ACM, 2010.

[11] S. Kairam, M. Brzozowski, D. Huffaker, and E. Chi, *Talking in circles: selective sharing in google+*, in *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, CHI '12, pp. 1065–1074, ACM, 2012.

[12] J. Donath and D. Boyd, *Public displays of connection*, *BT Technology Journal* **22** (Oct., 2004) 71–82.

[13] A. Lampinen, V. Lehtinen, A. Lehmuskallio, and S. Tamminen, *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, (New York, NY, USA), pp. 3217–3226, ACM, 2011.

[14] C. Honeycutt and S. C. Herring, *Beyond microblogging: Conversation and collaboration via twitter*, *2013 46th Hawaii International Conference on System Sciences* **0** (2009) 1–10.

[15] A. Monroy-Hernández, B. M. Hill, J. Gonzalez-Rivero, and d. boyd, *Computers can't give credit: How automatic attribution falls short in an online remixing community*, in *SIGCHI Conf. on Human Factors in Computing Systems*, 2011.

[16] S. Savage, A. Forbes, C. Toxtli, G. McKenzie, S. Desai, and T. Höllerer, *Visualizing targeted audiences*, .

[17] T. Gillespie, *The relevance of algorithms*, *Media Technologies* (2013).

[18] L. Sproull and S. Kiesler, *Reducing social context cues: Electronic mail in organizational communication*, *Management science* **32** (1986), no. 11 1492–1512.

[19] M. Burke, C. Marlow, and T. Lento, *Feed me: motivating newcomer contribution in social network sites*, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 945–954, ACM, 2009.

[20] L. Chidambaram and L. L. Tung, *Is out of sight, out of mind? an empirical study of social loafing in technology-supported groups*, *Info. Sys. Research* **16** (June, 2005) 149–168.

[21] J. Brown, A. J. Broderick, and N. Lee, *Word of mouth communication within online communities: Conceptualizing the online social network*, *Journal of interactive marketing* **21** (2007), no. 3 2–20.

[22] C. Lampe, J. Vitak, and N. B. Ellison, *Users and nonusers: interactions between levels of adoption and social capital.*, in *CSCW* (A. Bruckman, S. Counts, C. Lampe, and L. G. Terveen, eds.), pp. 809–820, ACM, 2013.

[23] D. Ganley and C. Lampe, *The ties that bind: Social network principles in online communities*, *Decision Support Systems* **47** (2009), no. 3 266–274.

[24] T. Zhou, *Understanding online community user participation: a social influence perspective*, *Internet Research* **21** (2011), no. 1 67–81.

[25] J. Donath, *Identity and Deception in the Virtual Community*. Routledge, London, 1999.

[26] M. Ma and R. Agarwal, *Through a Glass Darkly: Information Technology Design, Identity Verification, and Knowledge Contribution in Online Communities*, .

[27] D. R. Millen and J. F. Patterson, *Identity disclosure and the creation of social capital*, in *CHI '03 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '03, (New York, NY, USA), pp. 720–721, ACM, 2003.

[28] S. A. Rains, *The impact of anonymity on perceptions of source credibility and influence in computer-mediated group communication: A test of two competing hypotheses*, *Communication Research* **34** (2007), no. 1 100–125.

[29] B. Schneier, *A taxonomy of social networking data*, *IEEE Security and Privacy* **8** (July, 2010) 88–88.

[30] L. I. Labrecque, E. Markos, and G. R. Milne, *Online personal branding: processes, challenges, and implications*, *Journal of Interactive Marketing* **25** (2011), no. 1 37–50.

[31] S. Rafaeli and Y. Ariel, *11 online motivational factors: Incentives for participation and contribution in wikipedia*, .

[32] M. M. Wasko and S. Faraj, *Why should i share? examining social capital and knowledge contribution in electronic networks of practice*, *MIS quarterly* (2005) 35–57.

[33] N. B. Peddibhotla and M. R. Subramani, *Contributing to public document repositories: A critical mass theory perspective*, *Organization Studies* **28** (2007), no. 3 327–346.

[34] M. R. Morris, S. Counts, A. Roseway, A. Hoff, and J. Schwarz, *Tweeting is believing?: Understanding microblog credibility perceptions*, in *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, CSCW '12, (New York, NY, USA), pp. 441–450, ACM, 2012.

[35] A. Pal and S. Counts, *What's in a@ name? how name value biases judgment of microblog authors.*, in *ICWSM*, 2011.

[36] A. Pal, S. Chang, and J. A. Konstan, *Evolution of experts in question answering communities*, in *ICWSM*, 2012.

[37] Y.-C. Wang, M. Burke, and R. E. Kraut, *Gender, topic, and audience response: An analysis of user-generated content on facebook*, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, (New York, NY, USA), pp. 31–34, ACM, 2013.

[38] M. S. Bernstein, A. Monroy-Hernández, D. Harry, P. André, K. Panovich, and G. Vargas, *4chan and/b: An analysis of anonymity and ephemerality in a large online community*, .

[39] E. Goffman, *The Presentation of Self in Everyday Life*. Anchor, 1 ed., June, 1959.

[40] M. Sunnafrank, *Predicted outcome value during initial interactions a reformulation of uncertainty reduction theory*, *Human Communication Research* **13** (1986), no. 1 3–33.

[41] N. Ellison, R. Heino, and J. Gibbs, *Managing impressions online: Self-presentation processes in the online dating environment*, *Journal of Computer-Mediated Communication* **11** (2006), no. 2 415–441.

[42] Z. Tufekci, *Can you see me now? audience and disclosure regulation in online social network sites*, 2008.

[43] A. Marwick and d. boyd, *To see and be seen: Celebrity practice on twitter*, *Convergence* (2011).

[44] C. Lampe, N. Ellison, and C. Steinfield, *A familiar face(book): profile elements as signals in an online social network*, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, (New York, NY, USA), pp. 435–444, ACM, 2007.

[45] E. Schwammlein and K. Wodzicki, *What to tell about me? self-presentation in online communities*, *J. Comp.-Med. Commun.* **17** (July, 2012) 387–407.

[46] V. Barash, N. Ducheneaut, E. Isaacs, and V. Bellotti, *Faceplant: Impression (mis)management in facebook status updates.*, in *ICWSM* (W. W. Cohen and S. Gosling, eds.), The AAAI Press, 2010.

[47] J. Vitak, *The impact of context collapse and privacy on social network site disclosures*, *Journal of Broadcasting & Electronic Media* **56** (2012), no. 4 451–470.

[48] J. Subercaze, C. Gravier, and Laforest, *Towards an expressive and scalable twitter profile hash for users recommendation*, *ACM International Conference on Web Intelligence* (2013).

[49] N. Agarwal, H. Liu, L. Tang, and P. S. Yu, *Identifying the influential bloggers in a community*, in *Proceedings of the 2008 International Conference on Web Search and Data Mining*, WSDM '08, (New York, NY, USA), pp. 207–218, ACM, 2008.

[50] M. Pennacchiotti and A.-M. Popescu, *Democrats, republicans and starbucks afficionados: user classification in twitter*, in *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '11, (New York, NY, USA), pp. 430–438, ACM, 2011.

[51] J. C. Paolillo, S. Mercure, and E. Wright, *The social semantics of livejournal foaf: Structure and change from 2004 to 2005*, in *Proceedings of the ISWC 2005 Workshop on Semantic Network Analysis*, vol. 72, pp. 69–80, Citeseer, 2005.

[52] M. Michelson and S. A. Macskassy, *Discovering users' topics of interest on twitter: a first look*, in *Proceedings of the fourth workshop on Analytics for noisy unstructured text data*, AND '10, (New York, NY, USA), pp. 73–80, ACM, 2010.

[53] N. Nori, D. Bollegala, and M. Ishizuka, *Exploiting user interest on social media for aggregating diverse data and predicting interest*, 2011.

[54] I. Guy, U. Avraham, D. Carmel, S. Ur, M. Jacovi, and I. Ronen, *Mining expertise and interests from social media*, in *Proceedings of the 22nd international conference on World Wide Web*, WWW '13, (Republic and Canton of Geneva, Switzerland), pp. 515–526, International World Wide Web Conferences Steering Committee, 2013.

[55] X. Su and T. M. Khoshgoftaar, *A survey of collaborative filtering techniques*, *Advances in artificial intelligence* **2009** (2009) 4.

[56] B. Berendt, O. Günther, and S. Spiekermann, *Privacy in e-commerce: Stated preferences vs. actual behavior*, *Commun. ACM* **48** (Apr., 2005) 101–106.

[57] A. N. Joinson, C. Paine, T. Buchanan, and U.-D. Reips, *Measuring self-disclosure online: Blurring and non-response to sensitive items in web-based surveys*, *Comput. Hum. Behav.* **24** (Sept., 2008) 2158–2171.

[58] E. Kiciman, *OMG, I Have to Tweet that! A Study of Factors that Influence Tweet Rates*, in *ICWSM*, 2012.

[59] Wikipedia, *Livejournal*, Jan., 2013. http://en.wikipedia.org/wiki/LiveJournal.

[60] D. Blei, A. Ng, and M. Jordan, *Latent dirichlet allocation*, *the Journal of machine Learning research* **3** (2003) 993–1022.

[61] K. Bollacker, C. Evans, P. Paritosh, T. Sturge, and J. Taylor, *Freebase: a collaboratively created graph database for structuring human knowledge*, in *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, SIGMOD '08, (New York, NY, USA), pp. 1247–1250, ACM, 2008.

[62] G. Schwarz, *Estimating the Dimension of a Model*, *The Annals of Statistics* **6** (1978), no. 2 461–464.

[63] Y. Cheng, *Mean shift, mode seeking, and clustering*, *IEEE Trans. Pattern Anal. Mach. Intell.* **17** (Aug., 1995) 790–799.

[64] J. L. Fleiss, *Measuring nominal scale agreement among many raters.*, *Psychological bulletin* **76** (1971), no. 5 378.

[65] A. Lenhart, *Adults and social network websites*, .

[66] R. K. Merton, *The Role-Set: Problems in Sociological Theory*, *The British Journal of Sociology* **8** (1957), no. 2.

[67] C. Willyard, *Separating work, friends, and family on facebook isn't easy*, May, 2011. http://spectrum.ieee.org/telecom/internet/separating-work-friends-and-family-on-facebook-isnt-easy.

[68] W. H. Sumner, *An Inquiry Into the Importance of the Militia to a Free Commonwealth: In a Letter from William H. Sumner... to John Adams, Late President of the United States; with His Answer*, vol. 40. Cummings and Hilliard, 1823.

[69] L. Crothers, *Rage on the right: The American militia movement from Ruby Ridge to homeland security*. Rowman & Littlefield Publishers, 2004.

[70] S. Weeber and D. Rodeheaver, *Militias at the millennium*, *The Sociological Quarterly* **44** (2003), no. 2 181–204.

[71] G. Weimann, *New Terrorism and New Media*. Wilson Center Common Labs, 2014.

[72] G. Mark, B. Al-Ani, and B. Semaan, *Resilience through technology adoption: Merging the old and the new in iraq*, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, (New York, NY, USA), pp. 689–698, ACM, 2009.

[73] A. Monroy-Hernández, d. boyd, E. Kiciman, D. Choudhury, and S. Counts, *The new war correspondents: The rise of civic media curation in urban warfare*, in *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, CSCW '13, (New York, NY, USA), pp. 1443–1452, ACM, 2013.

[74] S. Gerwehr and S. Daly, *Al-Qaida: Terrorist Selection and Recruitment*, research report 5, pp. 73–89. McGraw-Hill, 2005.

[75] I. García, *Redes sociales, el otro campo de accin de las autodefensas*. ADN Politico, 2014. http://www.adnpolitico.com/gobierno/2014/01/14/redessocialeseł-otrocampodebatalladelasautodefensas.

[76] K. Zabludovsky, *Reclaiming the Forests and the Right to Feel Safe*. New York Times, 2012. http://www.nytimes.com/2012/08/03/world/americas/inmexico-reclaimingtheforestsandtherighttofeelsafe.html.

[77] O. Garcia-Ponce and A. Lajous, *How does a drug cartel become a lime cartel?* Washington Post, 2014. http://www.washingtonpost.com/blogs/monkey-cage/wp/2014/05/20/how-doesadrugcartelbecomealimecartel/.

[78] Processo, *Se enfrentan militares y autodefensas en Michoacan; reportan muerte de civiles*. 2014. http://www.proceso.com.mx/?p=362352.

[79] R. A. Noticias, *Autodefensas en Facebook relegan al gobierno a la hora de difundir informacion*. Aristegui Noticias, 2014. http://aristeguinoticias.com/2603/mexico/autodefensasenfacebookrelegana-lgobiernoalahoradedifundirinformacion/.

[80] V. Dijk, *Discourse as social interaction*, vol. 2. Sage, 1997.

[81] S. Robertson, R. Vatrapu, and R. Medina, *Off the wall political discourse: Facebook use in the 2008 u.s. presidential election*, *Info. Pol.* **15** (Apr., 2010) 11–31.

[82] K. Garrett, *Protest in an information society: A review of literature on social movements and new icts*, *Information, communication & society* **9** (2006), no. 02 202–224.

[83] D. McAdam, J. D. McCarthy, and M. N. Zald, *Comparative perspectives on social movements: Political opportunities, mobilizing structures, and cultural framings*. Cambridge University Press, 1996.

[84] V. Wulf, K. Aal, I. Abu Kteish, M. Atam, K. Schubert, M. Rohde, G. P. Yerousis, and D. Randall, *Fighting against the wall: Social media use by political activists in a palestinian village*, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, (New York, NY, USA), pp. 1979–1988, ACM, 2013.

[85] L. C. Irani and M. S. Silberman, *Turkopticon: Interrupting Worker Invisibility in Amazon Mechanical Turk*, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, (New York, NY, USA), pp. 611–620, ACM, 2013.

[86] Y. Mejova, V. R. K. Garimella, I. Weber, and M. C. Dougal, *Giving is caring: understanding donation behavior through email*, in *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, pp. 1297–1307, ACM, 2014.

[87] J. Dimond, M. Dye, D. Larose, and A. Bruckman, *Hollaback!: the role of storytelling online in a social movement organization*, in *Proceedings of the 2013 conference on Computer supported cooperative work*, pp. 477–490, ACM, 2013.

[88] A. Strauss and C. Juliet, *Basics of Qualitative Research : Techniques and Procedures for Developing Grounded Theory*. SAGE Publications, Sept., 1998.

[89] J. Feezell, M. Conroy, and M. Guerrero, *Facebook is... fostering political engagement: A study of online social networking groups and offline participation*, in *APSA 2009 Toronto Meeting Paper*, 2009.

[90] E. Jones, *Ingratiation*, *The Journal of Social Psychology* (1964).

[91] E. Cooper and M. Jahoda, *The evasion of propaganda: How prejudiced people respond to anti-prejudice propaganda*, *The Journal of psychology* **23** (1947), no. 1 15–25.

[92] E. King, *Exposing the age of lies: The propaganda menace as portrayed in american magazines in the aftermath of world war i*, *Journal of American Culture* **12** (1989), no. 1 35–40.

[93] R. Benford and D. Snow, *Framing processes and social movements: An overview and assessment*, *Annual review of sociology* (2000) 611–639.

[94] C. Mascaro and S. Goggins, *Brewing up citizen engagement: the Coffee Party on Facebook*, in *Proceedings of the 5th International Conference on Communities and Technologies*, pp. 11–20, ACM, 2011.

[95] A. Hughes and L. Palen, *Twitter adoption and use in mass convergence and emergency events*, *International Journal of Emergency Management* **6** (2009), no. 3 248–260.

[96] N. Smelser, *Theory of collective behavior*, *Cambridge New York, Cambridge University Press, Coll.* (1962).

[97] M. T. Maruyama, S. P. Robertson, S. K. Douglas, B. C. Semaan, and H. A. Faucett, *Hybrid media consumption: How tweeting during a televised political debate influences the vote decision*, in *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '14, (New York, NY, USA), pp. 1422–1432, ACM, 2014.

[98] M. S. Bernstein, E. Bakshy, M. Burke, and B. Karrer, *Quantifying the invisible audience in social networks*, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 21–30, ACM, 2013.

[99] S. Savage, A. Forbes, C. Toxtli, G. McKenzie, S. Desai, and T. Höllerer, *Visualizing targeted audiences*, in *COOP*, pp. 17–34, 2014.

[100] A. G. Forbes, S. Savage, and T. Höllerer, *Visualizing and verifying directed social queries*, in *IEEE Workshop on Interactive Visual Text Analytics. Seattle, WA*, Citeseer, 2012.

[101] R. Zambrano, *New technologies and democratic governance: Towards a more inclusive participation?*, *IN THE CONSTRUCTION OF DEMOCRACY* 139.

[102] A. Monroy-Hernández, S. Farnham, E. Kiciman, S. Counts, and D. Choudhury, *Smart societies: from citizens as sensors to collective action*, *interactions* **20** (2013), no. 4 16–19.

[103] J. Cranshaw, R. Schwartz, J. I. Hong, and N. M. Sadeh, *The livehoods project: Utilizing social media to understand the dynamics of a city.*, in *ICWSM*, 2012.

[104]

[105] M. Maruyama, S. Douglas, and S. Robertson, *Design teams as change agents: Diplomatic design in the open data movement*, in *System Sciences (HICSS), 2013 46th Hawaii International Conference on*, pp. 1860–1869, IEEE, 2013.

[106] E. Kıcıman, *OMG, I Have to Tweet that! A Study of Factors that Influence Tweet Rates*, in *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*, 2012.

[107] G. McKenzie, K. Deutsch, M. Raubal, P. Mooney, and K. Rehrl, *What, when and where: The real-world activities that contribute to online social networking posts*, in *International Workshop on Action and Interaction in Volunteered Geographic Information (ACTIVITY), Leuven, Belgium, edited by Peter Mooney, and Karl Rehrl*, pp. 25–30, 2013.

[108] E. Agapie and A. Monroy-Hernandez, *Eventful: Crowdsourcing local news reporting*, .

[109] D. Dailey and K. Starbird, *Journalists as crowdsourcerers: Responding to crisis by reporting with a crowd*, *Computer Supported Cooperative Work (CSCW)* **23** (2014), no. 4-6 445–481.

[110] S. Savage, A. Forbes, C. Toxtli, G. McKenzie, S. Desai, and T. Höllerer, *Visualizing targeted audiences*, in *COOP*, pp. 17–34, 2014.

[111] K. Brennan, A. Monroy-Hernndez, and M. Resnick, *Making projects, making friends: Online community as catalyst for interactive media creation*, *New Directions for Youth Development* **2010** (2010), no. 128 75–83.

[112] B. Butler, L. Sproull, S. Kiesler, and R. Kraut, *Community effort in online groups: Who does the work and why?*, 2001.

[113] P. Otto and M. Simon, *Dynamic perspectives on social characteristics and sustainability in online community networks*, *System Dynamics Review* **24** (2008), no. 3 321–347.

[114] B. Latané and J. Darley, *Group Inhibition of Bystander Intervention in Emergencies*, *Journal of Personality and Social Psychology* **10** (1968), no. 3 215–221.

[115] J. C. McCroskey and T. A. McCain, *The measurement of interpersonal attraction*, *Speech Monographs* **41** (1974), no. 3 261–266.

[116] M. J. Brzozowski and D. M. Romero, *Who should i follow? recommending people in directed social networks*, in *ICWSM*, 2011.

[117] Y. Hu, S. D. Farnham, and A. Monroy-Hernández, *Whoo.ly: Facilitating information seeking for hyperlocal communities using social media*, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, (New York, NY, USA), pp. 3481–3490, ACM, 2013.

[118] L. Singer, F. M. F. Filho, B. Cleary, C. Treude, M.-A. D. Storey, and K. Schneider, *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*, pp. 103–116, 2013.

[119] D. Askanase, *Returning to facebook groups*, May, 2013. http://communityorganizer20.com/2012/05/31/returning-to-facebook-groups/.

[120] A. M. Thomson and J. L. Perry, *Collaboration processes: Inside the black box*, *Public Administration Review* **66** (2006) 20–32.

[121] R. Bentley, J. A. Hughes, D. Randall, T. Rodden, P. Sawyer, D. Shapiro, and I. Sommerville, *Ethnographically-informed systems design for air traffic control*, in *Proceedings of the 1992 ACM conference on Computer-supported cooperative work*, CSCW '92, (New York, NY, USA), pp. 123–129, ACM, 1992.

[122] M. Raubal, H. Miller, and S. Bridwell, *User Centered Time Geography For Location-Based Services*, *Geografiska Annaler B* (2004).

[123] S. Savage, M. Baranski, N. E. Chavez, and T. Hollerer, *I'm feeling loco: A location based context aware recommendation system*, in *Advances in Location-Based Services: 8th International Symposium on Location-Based Services, Vienna 2011*, Lecture Notes in Geoinformation and Cartography. Springer, 2011.

[124] T. L. Cheyne and F. E. Ritter, *Targeting audiences on the internet*, *Communications of the ACM* **44** (2001), no. 4 94–98.

[125] M. S. Bernstein, E. Bakshy, M. Burke, and B. Karrer, *Quantifying the invisible audience in social networks*, in *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, CHI '13, pp. 21–30, ACM, 2013.

[126] J. Nichols, M. Zhou, H. Yang, J.-H. Kang, and X. H. Sun, *Analyzing the quality of information solicited from targeted strangers on social media*, in *Proc. of the 2013 Conf. on Computer-Supported Cooperative Work*, pp. 967–976, ACM, 2013.

[127] A. Perer, I. Guy, E. Uziel, I. Ronen, and M. Jacovi, *Visual social network analytics for relationship discovery in the enterprise*, in *2011 IEEE Conf. on Visual Analytics Science and Technology (VAST)*, pp. 71–79, IEEE, 2011.

[128] S. Whittaker, Q. Jones, B. Nardi, M. Creech, L. Terveen, E. Isaacs, and J. Hainsworth, *Contactmap: Organizing communication in a social desktop*, *ACM Transactions on Computer-Human Interaction (TOCHI)* **11** (2004), no. 4 445–471.

[129] D. Chen, J. Tang, J. Li, and L. Zhou, *Discovering the staring people from social networks*, in *Proc. of the 18th Int. Conf. on the World Wide Web*, pp. 1219–1220, ACM, 2009.

[130] J. Chen, W. Geyer, C. Dugan, M. Muller, and I. Guy, *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, pp. 201–210, ACM, 2009.

[131] D. W. McDonald and M. S. Ackerman, *Expertise recommender: a flexible recommendation system and architecture*, in *Proc. of the 2000 ACM Conf. on Computer-Supported Cooperative Work*, pp. 231–240, ACM, 2000.

[132] D. Ramage, D. Hall, R. Nallapati, and C. D. Manning, *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*, pp. 248–256, Association for Computational Linguistics, 2009.

[133] A. G. Forbes, S. Savage, and T. Höllerer, *Visualizing and verifying directed social queries*, in *IEEE Workshop on Interactive Visual Text Analytics*, October, 2012.

[134] G. McKenzie, B. Adams, and K. Janowicz, *A thematic approach to user similarity built on geosocial check-ins*, in *Geographic Information Science at the Heart of Europe*, pp. 39–53. Springer, 2013.

[135] B. Shneiderman, *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.

[136] K. Etemad and S. Carpendale, *Symmetry and node focused visualization of large trees*, in *GRAND NCE Annual Conf.*, 2010.

[137] M.-H. A. H. T. Savage, S., *Botivist: Calling volunteers to action using online bots*, in *Proceedings of the 19th ACM conference on Computer supported cooperative work & social computing*.

[138] E. Massung, D. Coyle, K. F. Cater, M. Jay, and C. Preist, *Using crowdsourcing to support pro-environmental community activism*, in *SIGCHI Conf. on Human Factors in Computing Systems*, ACM, 2013.

[139] K. Starbird and L. Palen, *"voluntweeters": Self-organizing by digital volunteers in times of crisis*, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, (New York, NY, USA), pp. 1071–1080, ACM, 2011.

[140] D. H. Smith, *Altruism, volunteers, and volunteerism*, *Journal of Voluntary Action Research* (1981).

[141] N. Byrnes, *Persuasive technologies surround us, and theyre growing smarter. how do these technologies work? and why?*, in *Technology and Persuasion*, MIT Technology Review, 2015.

[142] S. Saeed, M. Rohde, and V. Wulf, *Analyzing political activists organization practices: findings from a long term case study of the european social forum*, *Computer Supported Cooperative Work (CSCW)* **20** (2011), no. 4-5 265–304.

[143] E. Sommerfeldt, *Activist online resource mobilization: Relationship building features that fulfill resource dependencies*, *Public Relations Review* **37** (2011), no. 4 429–431.

[144] A. Velasquez and R. LaRose, *Youth collective activism through social media: The role of collective efficacy*, *New Media & Society* (2014) 1461444813518391.

[145] A. Voida, E. Harmon, and B. Al Ani, *Bridging between organizations and the public: volunteer coordinators' uneasy relationship with social computing*, in *SIGCHI Conference on Human Factors in Computing Systems*, ACM, 2012.

[146] V. Wulf, K. Aal, I. Abu Kteish, M. Atam, K. Schubert, M. Rohde, G. P. Yerousis, and D. Randall, *Fighting against the wall: Social media use by political activists in a palestinian village*, in *SIGCHI Conf. on Human Factors in Computing Systems*, pp. 1979–1988, ACM, 2013.

[147] K. Starbird and L. Palen, *(how) will the revolution be retweeted?: information diffusion and the 2011 egyptian uprising*, in *Proceedings of the acm 2012 conference on computer supported cooperative work*, pp. 7–16, ACM, 2012.

[148] B. J. Fogg, *Persuasive technology: using computers to change what we think and do*, *Ubiquity* **2002** (2002), no. December 5.

[149] I. P. Levin, S. K. Schnittjer, and S. L. Thee, *Information framing effects in social and personal decisions*, *Journal of Experimental Social Psychology* **24** (1988), no. 6 520–529.

[150] D. Maheswaran and J. Meyers-Levy, *The influence of message framing and issue involvement*, *Journal of Marketing research* (1990) 361–367.

[151] R. Rettie, U. Grandcolas, and B. Deakins, *Text message advertising: Response rates and branding effects*, *Journal of Targeting, Measurement and Analysis for Marketing* **13** (2005), no. 4 304–312.

[152] A. Tversky and D. Kahneman, *Rational choice and the framing of decisions*, *Journal of business* (1986) S251–S278.

[153] X. Yang, Q. Yang, and C. Wilson, *Penny for your thoughts: Searching for the 50 cent party on sina weibo*, in *Ninth International AAAI Conference on Web and Social Media*, 2015.

[154] M. Orcutt, *Twitter mischief plagues mexico's election*, *MIT Technology Review*.

[155] P. R. Center, *Crime and corruption top problems in emerging and developing countries*, *Pew Research Center*. http://www.pewglobal.org/2014/11/06/crime-andcorruptiontopproblemsinemerginganddevelopingcountries/.

[156] P. McSherry and M. Raul, *Introduction to shadows of state terrorism: Impunity in latin america*, *Social Justice* **26** (1999), no. 4. https://www.socialjusticejournal.org/SJEdits/78Edit.html.

[157] J. Nichols, M. Zhou, H. Yang, J.-H. Kang, and X. H. Sun, *Analyzing the quality of information solicited from targeted strangers on social media*, in *Proceedings of the 2013 conference on Computer supported cooperative work*, pp. 967–976, ACM, 2013.

[158] L. Shen, *Mitigating psychological reactance: The role of message-induced empathy in persuasion*, *Human Communication Research* **36** (2010), no. 3 397–422.

[159] E. P. Baumer, S. J. Katz, J. E. Freeman, P. Adams, A. L. Gonzales, J. Pollak, D. Retelny, J. Niederdeppe, C. M. Olson, and G. K. Gay, *Prescriptive persuasion and open-ended social awareness: Expanding the design space of mobile health*, in *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*, CSCW '12, pp. 475–484, 2012.

[160] K. Ling, G. Beenen, P. Ludford, X. Wang, K. Chang, X. Li, D. Cosley, D. Frankowski, L. Terveen, A. M. Rashid, *et. al.*, *Using social psychology to motivate contributions to online communities*, *Journal of Computer-Mediated Communication* **10** (2005), no. 4 00–00.

[161] E. Brady, M. R. Morris, and J. P. Bigham, *Gauging receptiveness to social microvolunteering*, in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 1055–1064, ACM, 2015.

[162] J. Weizenbaum, *Eliza: a computer program for the study of natural language communication between man and machine*, *Commun. ACM* **9** (Jan., 1966) 36–45.

[163] C. Nass, J. Steuer, and E. R. Tauber, *Computers are social actors*, in *SIGCHI Conf. on Human Factors in Computing Systems*.

[164] L. M. Aiello, M. Deplano, R. Schifanella, and G. Ruffo, *People are strange when youre a stranger: Impact and influence of bots on social networks*, *Links* **697** (2012), no. 483,151 1–566.

[165] J. A. Fleishman, *Collective action as helping behavior: Effects of responsibility diffusion on contributions to a public good.*, *Journal of Personality and Social Psychology* **38** (1980), no. 4 629.

[166] T. Aleahmad, A. D. Balakrishnan, J. Wong, S. R. Fussell, and S. Kiesler, *Fishing for sustainability: The effects of indirect and direct persuasion*, in *Extended Abstracts on Human Factors in Computing Systems*.

[167] M. Carroll, S. Goldman, L. Britos, J. Koh, A. Royalty, and M. Hornstein, *Destination, imagination and the fires within: Design thinking in a middle school classroom*, *International Journal of Art & Design Education* **29** (2010), no. 1 37–53.

[168] A. Smith, *Civic engagement in the digital age*, *Pew Internet* (2013).

[169] S. Savage and A. Monroy-Hernández, *Participatory militias: An analysis of an armed movement's online audience*, in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '15, pp. 724–733, 2015.

[170] Wikipedia, *Innocentive*, . http://en.wikipedia.org/wiki/InnoCentive.

[171] A. Kilgarriff, *Comparing corpora*, *International journal of corpus linguistics* **6** (2001), no. 1 97–133.

[172] N. Hamdy and E. H. Gomaa, *Framing the egyptian uprising in arabic language newspapers and social media*, *Journal of Communication* **62** (2012), no. 2 195–211.