

Encouraging Diversity- and Representation-Awareness in Geographically Centralized Content

Eduardo Graells-Garrido*

Telefónica I+D
Santiago, Chile

Mounia Lalmas

Yahoo Labs
London, UK

Ricardo Baeza-Yates

Yahoo Labs
Sunnyvale, USA

ABSTRACT

In centralized countries, not only population, media and economic power are concentrated, but people give more attention to central locations. While this is not inherently bad, this behavior extends to micro-blogging platforms: central locations get more attention in terms of information flow. In this paper we study the effects of an information filtering algorithm that decentralizes content in such platforms. Particularly, we found that users from non-central locations were not able to identify the geographical diversity on timelines generated by the algorithm, which were diverse by construction. To make users see the inherent diversity, we define a design rationale to approach this problem, focused on the treemap visualization technique. Then, we deployed an “in the wild” implementation of our proposed system. On one hand, we found that there are effects of centralization in exploratory user behavior. On the other hand, we found that the treemap was able to make users see the inherent geographical diversity of timelines. We measured these effects based on how users engaged with content filtered by the algorithm. With these results in mind, we propose practical actions for micro-blogging platforms to account for the differences and biased behavior induced by centralization.

ACM Classification Keywords

H.3.3 Information Storage and Retrieval: Information Search and Retrieval—*Information Filtering*; H.5.2 User Interfaces: Graphical user interfaces (GUI).

Author Keywords

Centralization; Location Bias; Information Filtering; Information Visualization.

INTRODUCTION

In his book on user experience, Bill Buxton wrote that “*in order to design a tool, we must make our best efforts to understand the larger social and physical context within which it is intended to function*” [8]. In today’s global Web, it is not clear if current social platforms consider those different

contexts when building their user interfaces or defining their content-based algorithms. This would not be a problem in an uniform, unbiased world, but our world is neither uniform [22] nor unbiased [29]. One of these biases is *centralization* [24], an organizational schema for governments that can be beneficial for the population, for development, and for the economy, because of the concentration of political, economical and media powers. However, at the same time, it can create regional inequalities within a country [25], making national growth biased toward central locations. Given that centralization is part of the social context of users, in this work we evaluate whether this bias affects the informational and exploratory behavior of users on the Web, particularly in micro-blogging platforms. We do so through the following research questions:

Does centralization affect how people perceive information, and how people behave when browsing informational content in micro-blogging platforms? If so, how can we encourage non-centralized exploration?

To answer our research questions, we analyzed the usage of Twitter by people in Chile interested in politics. This case study is perfect for our analysis for two reasons. On the one hand, Chile is a highly centralized country toward its capital region [4, 17]. Although this region is indeed near the geographical center of continental Chile, the country spans over 4,300 km (2,671 mi) from north to south, with only 175 km (108 mi) in average from east to west, having a geographical configuration that is not particularly fit for centralization. After returning to democracy, various initiatives attempted to decentralize the country, but enormous political and economical resistance allowed for only small progress towards this, if any [14]. On the other hand, Chile is one of the developing countries with the highest Internet penetration rate [10], and Twitter is actively used by politically-involved Chileans [40].

A previous study [20] determined that centralization is reflected in the way users interact in Twitter, creating regional inequalities in information flow with respect to user location. The authors developed an information filtering algorithm that generates geographically diverse timelines. By construction, the algorithm allows all locations to be present in a timeline. In this paper, we evaluate this algorithm through a user study. We find that users from central locations have a different sense of diversity, interestingness, and informativeness than users from peripheral locations.

The results from the user study imply that, even though in theory the algorithm solves the problem, in practice users might not notice it. After analyzing these differences, we hypothesize that centralization generates a *diversity- and representation-awareness* problem, and thus we define a design rationale to

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

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IUI 2016, March 7–10, 2016, Sonoma, CA, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-4137-0/16/03 ...\$15.00.

DOI: <http://dx.doi.org/10.1145/2856767.2856775>

address these issues. In compliance with our design rationale we then test a visualization of categorical news headlines by Weskamp [41] in our context, *i. e.*, in geographically diverse timelines with micro-posts instead of headlines. We hypothesize that this interface eases the identification in timelines by making geographical diversity salient through the usage of *squarified treemaps* [6], a well-known information visualization technique. We evaluate this design in an exploratory application deployed on the Web. The analysis of interaction data reveals that, even though there are differences in how people behave with respect to their geographical origin, our design facilitates diversity- and representation-awareness from a geographical point of view.

Our results show that centralization affects the perception and behavior of politically-involved users in Chile. They also demonstrate that awareness of social and geographical contexts when designing algorithms and user interfaces improves user perception in the presence of these systemic biases. Finally, considering that many countries are centralized [17, 24, 25], our findings and implications can be used by micro-blogging platforms to make their systems aware of important factors of social contexts.

STATE OF THE ART

This paper deals with *centralization* [24], in particular considering its geographical aspect. Geography is an important attribute to be considered in the study of social networks and Web platforms. In the context of Twitter, even though most of user ties are geographically local [35], more than a third of mentions and links are inter-countries [26]. To understand how a virtual population is distributed, each user's location must be determined, a meta-attribute not always available. When geolocating users in micro-blogging platforms, a simple approach is to query a gazetteer with the user's self-reported location [21], but there are more complex and accurate approaches that involve entity recognition [1] and language models [11]. In our context, location classifiers can become biased when the population is imbalanced [37], and care must be taken when parameterizing such algorithms.

Despite being a virtual platform, Twitter users reflect or are influenced by physical world phenomena. For instance, the number of international flights between countries is the best predictor of non-local ties [39]; tweets can predict gross community happiness by comparing aggregated word usage related to mood in communities [36]; interaction and publishing behavior on Twitter is related with cultural dimensions [19]; and international communication is still dominated by physical distance [18]. Communication and tie formation in Twitter are influenced by geography and culture—two aspects deeply connected with centralization. In Chile, the informational behavior of the population is centralized [20]. Thus, in our work, we evaluate if centralization influences user perception of geographically diverse timelines.

We focus on geographical diversity because under population imbalance and centralization, it is expected that timeline content is also imbalanced, both in volume and in popularity metrics (*e. g.*, number of retweets). Then, we propose to filter timelines using methods that increase geographical diversity

of its content. Diversification can be performed by minimizing similarity between items in a recommendation list [42], maximizing entropy of a set of content features [13], as well as context-specific diversification methods [31]. In the case of Twitter, diversified timelines (in its information entropy sense) have been found to be easier to remember and more engaging [13]. In this paper we build upon such work, by evaluating perception of timelines generated by an information filtering algorithm [20], which, in turn, is based on two previous algorithms that aim to generate diverse information sets [13, 31].

We focus on how user location influences those perceptions and how engagement differs in terms of location and user interface. Park *et al.* [33] showed that the presentation of news headlines in clusters generates more clicks on news than non-clustered displays. We extend their design guidelines by using *squarified treemaps* [6] to maintain clustered representations of micro-posts while, at the same time, making diversity visually salient and noticeable. Treemaps have been used before to visualize content from micro-blogging platforms by Archambault *et al.* [3], although their approach is different to ours: they visualize clustered keywords, whereas we visualize entire tweets using a design inspired by *Newsmap.jp* [41], a treemap visualization of news headlines. In the context of political diversity, it has been acknowledged that visualization-based interfaces help people behave differently [15].

It is known that people from different cultures behave in different ways when communicating, not only on micro-blogs [18] but also on other forms of communication, like instant messaging [23]. In our work, we propose that centralization introduces differences in how users behave, and thus, these differences should be accounted for when designing systems.

BACKGROUND

In this section we briefly describe the dataset and methods used by Graells-Garrido *et al.* [20], which provide the background of this work. Although the defined methods are general, we restrict ourselves to Twitter, a micro-blogging platform where users publish status updates called *tweets* with a maximum length of 140 characters. Users can *follow* other users, establishing directed connections between them. When user *A* follows user *B*, tweets and *re-tweets* made by *B* will show up in *A*'s timeline. A timeline is a list of tweets in reverse chronological order. Users can annotate tweets using *hashtags*, *i. e.*, keywords that start with the hash character #.

Context: Municipal Elections in Chile. In the first user study on this paper we analyze a dataset of tweets from Chile. Centralization in Chile is characterized through geography at the *regional level*, *i. e.*, the first-level administrative division of the country. Chile has 15 regions, and the capital *Región Metropolitana* (translated as *Metropolitan Region*, and *RM* hereafter) is the most central one [4, 17].

The dataset is composed of tweets crawled on October 28th, 2012, in the context of municipal elections held in Chile that day. The event had a distinctive hashtag (*#municipales2012*), which, among other related hashtags (*e.g.* *#tudecides*), keywords (*e.g.* *vote*), location and candidate names, were used as

queries for the *Twitter Streaming API*.¹ In total, this dataset contains 157,648 users who published 724,890 tweets. To explore the representativity of the sample, we estimated the Pearson product-moment correlation coefficient r between virtual and physical population, and between account rate and household Internet Access Rate in Chile [30]. Because of population imbalance, we estimated the correlation of the logarithms of population, obtaining $r = 0.95$ ($p < 0.01$). The user and household Internet access rates correlate at $r = 0.68$ ($p < 0.01$). Therefore, we consider the sample spatially representative of the physical population at the regional level.

Because its content is about local elections happening nationwide, this dataset is suitable to study the effects of centralization. The analysis made in [20] consists on the determination of potential centralization patterns in information flow. It was made by aggregating *mentions*, *replies* and *retweets* with respect to origin and destination locations of the involved user accounts. Then, an interaction graph between locations was built, in which the authors estimated *betweenness centrality* using random walks [32]. Betweenness centrality was chosen instead of other centrality metrics because it represents “the potential of a point for control of information flow in the network” [16]. The use of a random walk allowed the authors to consider the user distribution instead of binary edges in the graph, as well as to build a baseline graph based on the population distribution of the physical world. Then, by comparing both graphs, it was possible to determine whether informational behavior in Twitter, in terms of to whom people interact with, is centralized. It was found that: 1) *RM* is the most central location in Twitter discussion; 2) that centralities are significantly different between observations (interactions) and expectations (physical world population distribution); and 3) the observed centrality of *RM* is higher than the expected one, and the centralities of all other locations exhibit the opposite behavior, implying that centralization also affects informational behavior in political discussion of Chileans in Twitter.

Building Geographically Diverse Timelines. The influence of centralization makes content from peripheral locations harder to find. One way to help to reduce this bias is through generation of geographically diverse timelines where all regions are represented regardless of their centrality. This can be achieved through the usage of information filtering techniques.

In this paper, we use an information filtering algorithm previously defined Graells-Garrido *et al.* [20]. This *Proposed Method* (hereafter *PM*) mixes two other known algorithms:

1. The first algorithm [13] is greedy and based on *information entropy*. Entropy is estimated with respect to several content-based features extracted from tweets (presence of links, topical information in hashtags, author connectivity and experience, age of content, and so forth). At each iteration, the algorithm selects a tweet from a candidate pool that maximizes the current entropy of the filtered timeline. In [20] this algorithm was modified to consider popularity (number of retweets) as a feature, by preferring candidate tweets who are more popular than others.

2. The second algorithm is known as *Sidelines* [31]. Because the complexity of the information features considered in [13] can be greater than those of geography (*e. g.*, consider the number of hashtags against the number of locations), the entropy contribution of these dimensions can be higher than the entropy contribution of geography. Thus, geographical diversity is enforced by *sidelining*, *i. e.*, when a tweet is selected for inclusion by the first algorithm, its location is not considered at the following n iterations.

The mixed algorithm applies both base algorithms simultaneously: tweets are selected according to their entropy contributions, but only considering those from locations that have not been sidelined. With this schema a timeline of s tweets can be built, with ensured information and geographical diversity. Whether these timelines are perceived as more diverse, interesting and informative than baseline timelines (including the first base algorithm [13]) is what we evaluate in the next section.

USER PERCEPTION OF DIVERSE TIMELINES

We describe the user study performed to evaluate how users perceive timelines generated with the *PM* algorithm [20] in comparison with those of baseline conditions. Our hypothesis is that centralization influences user perception, given that it is part of the cultural context of users. Users are grouped according to their geographical origin: a centralized location (*RM* group) or a peripheral one (*NOT-RM* group). Then, users are asked to compare three attributes of timelines: *diversity*, *interestingness* and *informativeness*. Note that interestingness and informativeness are related to relevance and coverage, respectively. However, we used user-centered terms as our study targeted end-users who may have been more familiar with the terms we chose.

Conditions and Datasets. In addition to the *Proposed Method (PM)*, we consider the following baseline conditions: *Popularity Sampling (POP* hereafter): we select the s most popular tweets in terms of retweets; and *Diversity Filtering (DIV* hereafter): an implementation of the algorithm that maximizes information entropy of each timeline [13]. Because the different conditions require pairwise comparisons, we split the dataset from [20] in three:

1. *morning-noon*: 140,211 tweets published by 52,403 users between 10:00AM and 2:30PM.
2. *afternoon*: 180,824 tweets published by 63,388 users between 2:30PM and 9:00PM.
3. *night*: 401,029 tweets published by 106,942 users between 9:00PM and 2:00AM (next day).

This division of time matches the local culture where lunch happens between 1:00PM and 2:30PM, and dinner around 9:00PM. For each dataset we built filtered timelines with *PM*, *DIV*, and *POP*. We excluded retweets as our focus is on stand-alone, source tweets.

Participants. Participants were recruited using *snowball sampling* in Twitter using open calls to volunteer in the study, which were retweeted by participants. No compensation was offered. We recruited 125 participants. Of them, 81 were male,

¹<https://dev.twitter.com/docs/streaming-apis>

41 were female and 3 opted to not say. In terms of age, 1 was 18–19, 59 were 20–29, 54 were 30–39, 4 were 40–49, 1 was 50+ years old and 6 opted not to say. All participants were from Chile (87 in *RM* and 38 in *NOT-RM*). Experience with social networks was asked using a five-point Likert scale: those from *RM* scored 3.73 ($\sigma^2 = 0.59$), and those from *NOT-RM* scored 3.68 ($\sigma^2 = 0.58$).

Experimental Setup. For each sub-dataset we generated three timelines (with $s = 30$ tweets) using *PM*, *DIV*, and *POP*. Timelines were displayed with a format that resembled Twitter user interface.

Procedure. The study had a *within-subjects* design. First, participants were asked to fill a questionnaire about demographic information and other features such as Twitter usage. Then, in at most three steps, users performed a series of comparisons between two timelines rendered side by side, each one generated by a different condition. To avoid sequence effects, the order of pairwise comparisons (*POP/PM*, *POP/DIV*, *DIV/PM*) and the order of sub-datasets (*morning-noon*, *afternoon*, *night*) were randomized in both, the position on the screen (left or right) and the experimental step. Hence, all the participants contributed to all conditions whenever possible, as some participants were expected to not do all comparisons being an on-line, volunteered study. Additionally, not all participants were expected to do a full read of timelines, and thus we discarded comparisons where the total reading time of both timelines was less than one minute.

Task. Participants were instructed to read the two timelines side by side, and then answer the following questions:

1. Which of the two timelines is more diverse?
2. Which of the two timelines is more interesting?
3. Which of the two timelines is more informative?
4. Optional: Please explain your answers. Add examples if needed.

Questions 1 to 3 had a seven-point Likert scale from -3 to 3, where -3 (or 3) means that the timeline on the left (or right) was perceived as more *diverse*, *interesting* or *informative* than the other, and a value of 0 means that there was no perceived difference. Question 1 asked for general diversity as we did not want to prime participants into thinking primarily about geographical diversity, although we explained that it should be considered in its widest sense, including geography, by adding the following subtitle to the question: “Consider diversity in its widest sense (geographical, demographical, topical, temporal, etc).” Question 4 presented a free-text form element. After answering the four questions, a pause screen was shown for 15 seconds to allow participants to rest.

Statistical Model. The aforementioned questions define three dependent variables: *diversity*, *interestingness* and *informativeness*. For each dependent variable we built the following statistical model (using *R formula syntax*):

$$Y \sim C(\text{comparison}) \times C(\text{location}),$$

where Y is an ordered response variable, $C(\text{comparison})$ is a dummy variable that encodes the specific pairwise compar-




Figure 1: Point plots of median scores given by participants in comparisons of diversity, informativeness and interestingness. For each pairwise comparison, a positive value indicates that the approach on the right of the label was perceived to be more diverse, interesting, and informative than the one on the left, and *vice-versa*.

isons performed in the study, and $C(\text{location})$ is a variable that encodes whether users are from *RM* or not. The operator \times encodes the main effects plus the interaction between factors.²

Over this model we performed a *generalized linear model* regression with a *proportional odds model* [27]. This model is also known as *ordered logistic regression*, and is used when modeling ordinal dependent variables. It extends the logistic regression model by allowing more than two categories, considering the order of the responses, and not assuming equidistant items in the Likert scale. If the statistical interaction was not found to be significant, then we performed another regression without the interaction terms (*i. e.*, using the $+$ operator instead of \times).

Results

In total, participants performed 238 comparisons: 84 for conditions *POP/DIV*, 80 for conditions *POP/PM*, and 74 for conditions *DIV/PM*. The following are the interpolated medians of pairwise comparisons given by users:

- *POP/DIV*: diversity (0), interestingness (-1), informativeness (-1). Users perceived both as equally diverse, but POP is found to be more interesting and informative.
- *POP/PM*: diversity (1), interestingness (0), informativeness (0). PM is found by users to be more diverse than POP, and both are perceived as equally interesting and informative.
- *DIV/PM*: diversity (-0.5), interestingness (1), informativeness (2). DIV is found by users to be slightly more diverse than PM, but PM is perceived as more interesting and informative.

Figure 1 showcases the medians of pairwise comparisons between conditions (questions 1, 2 and 3), considering individual differences according to location. It can be observed that statistical interactions could appear in the three attributes under consideration, specially with diversity. To formalize this intuition, we applied the regression defined previously to the user scores of comparisons. The following list describes the fit of those regressions:

- Diversity: the regression with interaction terms (log-likelihood = -447.21, AIC = 910.41, cond. H = 1,900) presented a significant interaction, and is used for analysis.
- Interestingness: the regression with interaction terms did not present significant terms according to the likelihood ratio test. The regression without interaction (log-likelihood

² $A \times B = A + B + A * B$.

Result #	Variable	Effect	β	Odds-Ratio	95% C.I.	p
R1	Diversity	Location <i>RM</i> and condition <i>POP/PM</i>	1.291	3.637	[0.317, 2.278]	0.010
R2	Interestingness	Location <i>RM</i>	0.686	1.985	[0.307, 1.069]	< 0.001
R3	Interestingness	Comparison <i>POP/DIV</i>	-1.404	0.246	[-1.885, -0.933]	< 0.001
R4	Interestingness	Comparison <i>POP/PM</i>	-0.859	0.423	[-1.353, -0.373]	< 0.001
R5	Informativeness	Location <i>RM</i>	1.364	3.910	[0.869, 1.872]	< 0.001
R6	Informativeness	Comparison <i>POP/DIV</i>	-0.926	0.396	[-1.618, -0.241]	0.008
R7	Informativeness	Location <i>RM</i> and comparison <i>POP/DIV</i>	-1.201	0.301	[-2.148, -0.261]	0.012
R8	Informativeness	Location <i>RM</i> and comparison <i>POP/PM</i>	-1.543	0.214	[-2.577, -0.521]	0.003

Table 1: Significant effects from the user study, including coefficients, effect sizes (Odds-Ratios) and confidence intervals. We only include effects with significant p-values ($p < 0.015$).

= -430.84, AIC = 873.67, cond. H = 460) is used for analysis.

- Informativeness: the regression with interaction terms (log-likelihood = -429.35, AIC = 874.70, cond. H = 200) presented a significant interaction, and is used for analysis.

The condition numbers of the respective Hessians indicate that the fits are not ill-defined. Table 1 shows the significant factors of each regression. We identify each result as R_i , to reference it later in the discussion. We mention only those that are significant according to their p-value (considering $p < 0.05$).

Discussion of Results

The experimental results support our hypothesis. Moreover, they give us two important insights. First, *popularity is more valued than diversity*. Users give more value to popularity than diversity, considering value as the mixture of informativeness and interestingness. Even *POP*, that is not geographically diverse, is perceived as more interesting than *DIV* (R3) and *PM* (R4). Likewise, *POP* is perceived as more informative than *DIV* (R6, R7 -*RM* only-) and *PM* (R8 -*RM* only-). Yet, the Odd-Ratios (OR) of these results are low, and thus, while there is an effect, this effect is subtle. Second, *perception of timeline properties depends on the geographical origin of users*. By design, we expected *PM* to be perceived as more diverse by users than *POP*, and more informative/interesting than *DIV*. However, this happened *only when users come from RM* (R1). Moreover, just being from *RM* indicates that users are more prone to find content interesting/informative, regardless of its diversity (R2, R5). The high odd-ratios of these results indicate that the effect is strong. The interaction between location and comparisons of *POP* and the diversity-driven algorithms *DIV* and *PM* (R7, R8), indicates a slight (given its ORs) accentuation of the effect.

Qualitative Feedback. Users' answers allowed us to validate the motivations and design decisions behind the proposed algorithm [20]:

"In terms of diversity, [*PM*] is partly more diverse because it is more geographically diverse. I was in doubt to choose [*POP*] as more diverse because it contained some tweets about several topics, but at the end I preferred geography" [P76, *NOT-RM*].

"I feel [*PM*] has much more information related to many locations, unlike timeline [*DIV*] which is focused on informing only about [*RM*]" [P50, *RM*].

One reason popularity was considered in the algorithm was to avoid potential noise present in *DIV*, because noisy informa-

tion is likely to increase entropy. This difference, as well as the similarities between both approaches, was noted by some users:

"It seems to me that both timelines [*DIV/PM*] are similar in the background, although [*PM*] has an informative emphasis, more than merely anecdotal as [*DIV*]. In a way, it could be said that [*DIV*] is slightly more diverse in comparison to [*PM*], but in my judgement those 'diverse' tweets were not interesting enough, instead they are more likely to be 'noise' in the main topic of both timelines" [P47, *RM*].

"I find [*POP*] more interesting as it talks about a particular topic, while [*DIV*] is about self-referential stuff that is not always of common interest" [P86, *NOT-RM*].

The above answers also surface the different concepts users have when considering diversity, as some users consider topical and geographical diversity, whereas others explicitly stated *personal opinions* as a source of diversity. However, users agree that a noticeable difference between *PM* and *DIV* relies on the fact that *PM* presents tweets from popular users:

"In general, timeline [*PM*] is based on known sources or journalists. Timeline [*DIV*] is about people from the community giving their impressions on what is going on, which is interesting but it doesn't have any backup information" [P38, *RM*].

"It could be said that timeline [*DIV*] has a diversity of users that are not, in contrast with those of timeline [*PM*], notable personalities of the 'Twitter world', however this is not enough for it to be considered more diverse (...). In terms of informativeness, timeline [*PM*] is ahead, since it is comprised of Twitter users dedicated to inform, it is expected for it to contain more complete tweets in terms of what was going on" [P47, *RM*].

If informativeness and interestingness are linked with user status, then the answers above support the quantitative results where *PM* was more interesting and informative than *DIV*. However, this feedback did not explain the differences reported in diversity by users.

Diversity and Representation. From these insights, we hypothesize that people from central locations find *PM* more diverse, and any approach more informative/interesting, because they are not used to see information from peripheral locations. It is known that the geographical span of ego-networks in Twitter is small [35], and thus, exposing those users to views from other locations expands their vision.

In contrast, people from peripheral locations do not see differences in diversity, nor find the content informative/interesting, because they are used to be exposed to views from *RM*. For peripheral users, before filtering, the timeline content focused prominently on the centralized and most populated locations; after, it contained a wider set of locations, but still not prominently their own. This concept is known as *out-group ho-*



Figure 2: Screenshot from <http://auroratwittera.cl>, the URL of our prototype implementation. The bottom bar contains the location filters. The upper bar contain navigational links to an *About* page and a feedback form.

mogeneity [38], and it can be described through the phrase “*they are alike, we are diverse*” [34]. We recall that the source dataset guarantees the presence of local content, given that it is about nationwide municipal elections happening at the same time. This implies that users did not recognize the presence of their locations in the generated timelines. Hence, in the next section we seek to enhance *representation-awareness* of users, to help users to recognize their presence in the timeline.

A DIVERSITY-AWARE PLATFORM

We propose a platform design that seeks to increase awareness of representation caused by diversity in timelines generated by the *PM* algorithm. The platform is deployed “in the wild” [12], *i. e.*, we target end-users in their everyday use of Twitter.

Design Rationale: Awareness Through Identification

Since *PM* timelines are geographically diverse by definition, perhaps what people from peripheral locations need is to *become aware* of their representation in the information stream. We propose to increase this awareness by facilitating *identification* as defined by Butler [7], *i. e.*, identification relies upon differences with the others. In our context, this means that by emphasizing differences with the others, we can facilitate identification (and thus, representation-awareness). We do so by making diversity salient, allowing users to *see* the geographical diversity present in the timeline.

Making Diversity Salient. Prior work determined that presentation [33] and visualization [15] change the way users behave in the presence of diverse information. In particular, the grouping of news headlines according to agreement with political positions improved user access to diverse information, measured in clicks on those headlines [33]. Then, clustering

content in locations would make easier for users to quickly see their own locations represented. However, it must be applied with care because our context is different. In political contexts, information is usually classified into a bipartite separation of groups (*e. g.*, *democrats* and *republicans*, *conservative* and *liberal*, *etc.*), whereas in our case the number of locations is larger (*i. e.*, 15 Chilean regions), creating the need to scroll on the screen and thus inducing a positional bias, by giving more importance to those clusters already visible without scrolling.

To avoid scrolling and its associated positional bias, we consider a previous visualization of news headlines by Weskamp [41], which uses a 2D space-filling layout algorithm to partition the available screen space: the *squarified treemap layout* [6]. We visually encode locations as *internal nodes* and tweets as *leaves*. Sibling leaves appear together. The *area size* of each leaf depends on the number of retweets of the corresponding tweet, and the number of followers and friends of its author, in inverse proportions to population location. In this way, screen space is shared in a fair way between locations, as shown in Figure 2. Each leaf node is colored according to its location (*hue*) and its recency (*saturation*). Internal nodes are not displayed as they are not needed.

User Interaction. To interact with the visualization, users can *click* a leaf node to display a pop-up with detailed information about the corresponding tweet, with buttons to perform core Twitter interactions (reply, retweet, mark as favorite, follow) and a text format that resembles the typical tweet presentation. To *filter locations*, for each location we display a button that, when clicked, updates the treemap to display only tweets originated from the selected location.

Prototype

We implemented a prototype of the user interface (see Figure 2) using the *d3.js* [5] library. This interface is named “*Aurora Twittera de Chile*” (*AT* hereafter) and is available at <http://auroratwittera.cl>. Every 30 minutes a “new issue of *AT*” was generated by the filtering algorithm (having $s = 30$ as size for each timeline and $n = 5$ for turns in the sideline step in *PM*).

We used the same implementation of the filtering algorithm *PM* from the first user study, with two differences: first, for performance reasons, we did not use a location classifier for all tweets in the dataset. Instead, we considered only tweets from accounts with a known self-reported location, although we did consider non-geolocated tweets retweeted by those accounts. Second, we avoided repeated authors or tweet content in the same timeline, and we discarded tweets where almost all text was in uppercase letters to avoid *shouting*.

Input tweets were downloaded with a crawler using the *Twitter Streaming API*. As query keywords we used location names, political terms, and other terms of interest that appear constantly on the news, as well as mentions to media accounts, both at national and local levels. Each issue had a specific URL in the form <http://auroratwittera.cl/timeline/ID>, which allowed users to access it at a later time, as well as saving a permanent link. In addition, if the URL contained the code of a location (e. g., <http://auroratwittera.cl/#RM>) the interface displayed immediately the tweets related to that location in the same way as if a location filter button had been pressed.

Social Bot @todocl. Since social bots can generate social discussion and behavioral changes based on their activity [2], we created a social bot on Twitter, with username *@todocl*, to publicize *AT* and recruit users. *@todocl* presented itself as a social experiment to establish an informative community about current happenings in Chile using Twitter, and published three types of tweets:

- Whenever a timeline was generated, *@todocl* published two tweets with a link to its corresponding *issue*: one mentioning four users who authored “featured tweets”, and one mentioning four users with “featured retweets” (in both cases users were selected randomly from filtered timeline).
- After publishing those tweets, every minute *@todocl* retweeted one tweet featured in the current issue.
- Every hour past 45 minutes, *@todocl* published 15 tweets, one per location, featuring a link to the current issue with each specific location in the URL, as well as an attached image with a *word cloud* of their most representative terms obtained with *Term Frequency/Inverse Document Frequency* (TF-IDF).

This implementation, comprised of filtering algorithm, user interface and social bot, creates a platform where users can access geographically diverse information, in the form of an external application to Twitter, as well as injected into the platform itself. We evaluate this application next.

EVALUATION WITH INTERACTION DATA

In this second study, we evaluate user behavior based on the interaction data obtained from *AT*, in particular at differences with respect to centralization. Our hypothesis is that the presentation using treemaps will help users overcome the awareness problems, due to the reasons specified in our design rationale. Note that *AT* is not a task-based system, *i. e.*, we do not expect users to visit the site to perform a specific task. Instead, the system is designed as an exploratory interface to geographically diverse timelines. By being exploratory, common evaluation metrics such as task accuracy or task performance cannot be applied. Instead, we evaluate user interaction with our system from the point of view of three dimensions: diversity-awareness, representation-awareness, and interestingness. From the analysis of those dimensions, we elaborate plausible explanations of potential differences in behavior as indicated by statistical analysis.

Experimental Setup

In this experiment we analyze interaction data obtained from user interaction with our implementation in *AT*. Each interaction is considered an event. We logged implicit events on the server, such as *session created/restored*, and from the client browser using Javascript, such as *timeline and UI loaded*, and *pings*. As explicit events we logged clicks on every element of the user interface, allowing us to identify whether users filtered locations, clicked on links, asked for a more detailed view of a tweet, as well as clicking on the Twitter native buttons to reply, mark as favorite, or follow a tweet’s author.

Using the logged events we analyze user behavior from three dimensions derived from interaction data and related to the problem at hand:

- **Diversity-Awareness:** defined as the number of different locations that generated the content the user interacted with. Every time the user interacted with content, we logged the location of the author that published it. Thus, we define a variable named *distinct locations* that encodes the number of different locations the user has interacted with. Note that we assume that users who clicked elements from a tweet are aware of the content they are clicking.
- **Representation-Awareness:** defined as the likelihood of selecting specific locations for browsing by users. We define a variable named *filter likelihood* that is 0 if the user has not selected specific locations using the location filters from the UI, and 1 otherwise.
- **Interestingness:** defined as the number of interactions with content by the user per day. We define a variable named *content events* that sums the number of clicks in links, the number of tweets seen with more detail, the number of tweets marked as favorite or replied, as well as the number of times other users have been followed. Given that some users visited the site more than once, we normalize this variable according to the number of different days each user has accessed the site.

Participants. We gathered interaction data from October 1st, 2014 until January 20th, 2015. The server logged each user request and was able to identify sessions based on cookies

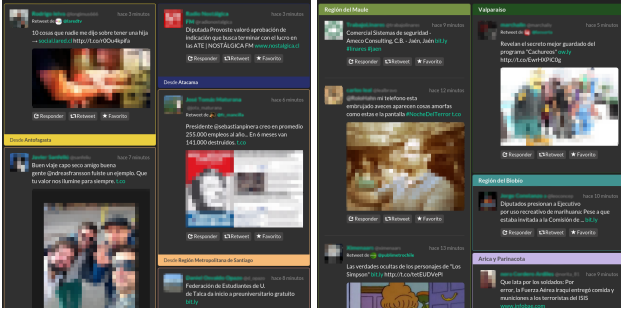


Figure 3: Design baselines implemented for the study. Left: standalone tweets (each in its own box). Right: clustered tweets by location.

placed on user browsers. IP addresses were used to identify each user’s location, using the GeoIP Legacy Database.³ The User Agent information was used to determine if the user was browsing from a mobile device; those users were served with a minimal version of the site but were not considered in the study because of platform heterogeneity (in terms of interaction capabilities, screen sizes, etc.). From 16,969 valid interaction events, we have 321 users, of which 193 are in *RM* and 128 are in *NOT-RM*. We discarded users who were active in the site for less than 10 seconds or that were in the top 5% of dwell time,⁴ as well as those who deactivated Javascript or could not be geolocated using GeoIP.

Conditions. To evaluate the effect on user behavior of both, location and user interface, we developed two alternative baseline conditions to compare with our proposed design: 1) *Baseline*, where each tweet is rendered independently of the others, and tweets are sorted by their time of publication. Each tweet is displayed inside a box with a bordered color and a legend at the bottom to indicate its originating location; and 2) *Clustered* [33], where tweets are clustered by location. Each location is represented as a box with a bordered color and a legend at the top to indicate the originating location of its tweets. Both baselines are shown in Figure 3. It can be seen that, although they are text-based, they include visual cues that indicate differences in location. Also, both have the same interactive location filters as our condition, *Treemap*.

When a user accessed AT, if it was his/her first visit, a random condition was assigned. Because we tracked users using cookies, in following requests users received the same interface according to their initial assignment. The distribution of users is as follows: *baseline* was assigned to 97 users; *clustered* to 98 users; and *treemap* to 126 users.

Setup. We consider a *between-subjects* design, as participants were exposed to one condition only. When users loaded AT we gathered the following user information: *IP address*, *HTTP Referrer* and *User Agent*. Then, we logged each interaction with elements of the user interface. If the user requested the page through an URL with a location code, we considered it as an initial click on a location filter, because such links

³<http://dev.maxmind.com/geoip/legacy/geolite/>

⁴Those users presumably left their browser windows open.

were generated by the social bot, and the user explicitly chose to follow them instead of the main URL. Finally, the user interface sent a ping every ten seconds to the server to track the time spent on the website even in the absence of interactions, to capture the dwell time of passive users.

Statistical Model. For analysis, we evaluate the effect of location and condition in the dependent variables defined earlier: *distinct locations*, *filter likelihood* and *content events*. We consider the two categorical independent variables *location* (*RM* or *NOT-RM*) and *condition* (*baseline*, *clustered*, or *treemap*). Both are included in the following statistical model:

$$Y \sim C(\text{condition}) \times C(\text{location}).$$

As in the previous study, we consider main effects of each factor as well as the interaction between them. The *C* element in the formula generates *contrasts* as dummy variables.

Over this model we perform *generalized linear model* regressions. The link function used for regressions of *distinct locations* and *content events* is the *Negative Binomial* distribution, which is commonly used to model over-dispersed count data. For *filter likelihood* we use *logistic (logit)* regression, because we model a probability. For each variable, if the statistical interaction was not found to be significant, we performed another regression without the interaction term.

Results

Figure 4 shows the distributions of our variables. Additionally, we include the distribution of the number of days each user has visited the site, because it is used to normalize the variable content events, which is not normalized in the figure.⁵ The following list describes the regression of each variable:

- Variable *distinct_locations* (mean = 0.364, std = 1.013, max = 10): the *N.B.* regression (deviance = 248.44, $\chi^2 = 453$, log-likelihood = -240.45) has intercept $\beta = -1.987$ (95% C.I. [-2.683, -1.290], $p < 0.001$).
- Variable *filter_likelihood* (mean = 0.436, std = 0.497, max = 1): the *logit* regression (log-likelihood = -212.41, $p = 0.002$) has non-significant intercept $\beta = -0.289$ (95% C.I. [-0.779, 0.201], $p = 0.247$).
- Variable *content_events* (mean = 1.688, std = 3.313, max = 33): the *N.B.* regression (deviance = 374.79, $\chi^2 = 564$, log-likelihood = -518.79) has non-significant intercept $\beta = 0.105$ (95% C.I. [-0.328, 0.538], $p = 0.635$).

Table 2 shows the significant coefficients of each regression. We show only those with a p-value lesser than 0.05. We identify each result as R_i , to reference it later in discussion. As effect sizes we analyze Odd-Ratios (OR). In contrast with our first study, no interaction terms were significant.

Discussion of Results

Diversity- and Representation-Awareness. With respect to diversity-awareness, we found that people from *RM* are more likely to be diverse-aware (R_{10} , OR = 2.473), which is coherent with the results from the first user study, given the main

⁵We applied a N.B. regression to the number of days, but no effect of condition/location was found.

Result #	Variable	Effect	β	O.R.	95% C.I.	p
R9	distinct_locations (N.B.)	Condition <i>treemap</i>	0.732	2.080	[0.092, 1.373]	0.025
R10	distinct_locations (N.B.)	Location <i>RM</i>	0.950	2.473	[0.348, 1.553]	0.002
R11	filter_likelihood (logit)	Condition <i>treemap</i>	0.7541	2.126	[0.207, 1.301]	0.007
R12	filter_likelihood (logit)	Location <i>RM</i>	-0.494	0.610	[-0.955, -0.033]	0.036
R13	content_events (N.B.)	Condition <i>treemap</i>	0.628	1.873	[0.167, 1.088]	0.008

Table 2: Significant effects from the user study, including coefficients, odds-ratios, and confidence intervals.

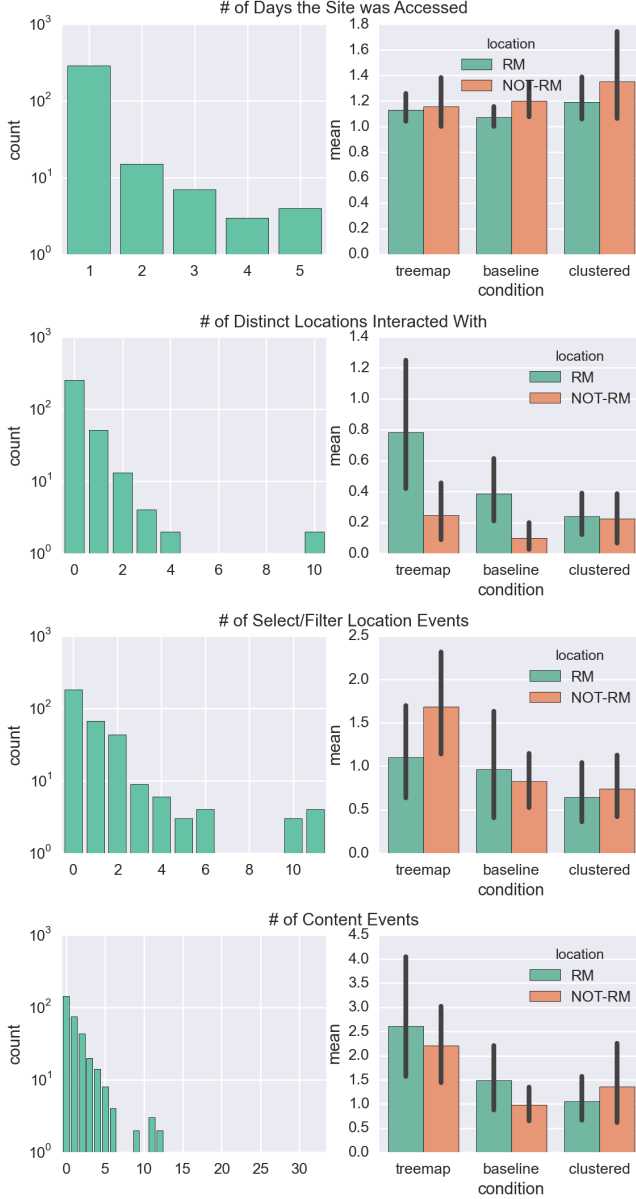


Figure 4: Distribution of each variable analyzed from interaction data. For each variable we include a histogram (left) and a bar plot (right) to compare means between groups.

effect of location with the distinct location variable. Moreover, we found that people exposed to the treemap condition were as likely to be diverse-aware (R9, OR = 2.080), which means

that people from *NOT-RM* exposed to the treemap can be as diversity-aware as people from *RM*.

With respect to representation-awareness, we found that people from *RM* is less likely to select specific locations (R12, OR = 0.60), which means that people from *NOT-RM* is more likely to do so. Moreover, we found that people exposed to the treemap condition were more likely to select specific locations (R11, OR = 2.126).

Interestingness. We did not find an effect of location on the number of content events, which means that users from all locations were equally engaged with content. However, we found that people exposed to the treemap condition were more likely to interact with content (R13, OR = 1.873). Its OR indicates that using the treemap increases the chances of interaction with content by 87.3%.

Our research question asks: *how can we encourage non-centralized exploration?* As interpreted from our first user study, it should never be just about the algorithm, but also about how users respond to what the algorithm returns to them. This response is influenced not only by the algorithm, but also by the user interface used. At the beginning of this section we stated that the main purpose of *AT*'s design was to make users more diversity- and representation-aware. We believe we succeeded, because, regardless of geographical origin (*i. e.*, no statistical interaction), the treemap design showed good properties on the three dimensions under consideration. These properties, in terms of diversity, representation, and user interest, indicate that a good way to encourage such behavior is by allowing users to explore content using non-traditional user interfaces. In fact, the clustered condition did not have any significant effect, even though it was expected [33]. As outlined in our rationale, the number of clusters can make exploration more difficult for users.

IMPLICATIONS

In the literature, cultural differences [22] have been acknowledged in the study of communication systems [23] by suggesting specific features and interaction mechanisms pertinent to each culture. As we have found in our studies, Chilean information seekers and content explorers in Twitter are affected by centralization. Their perception of content is different, and their interaction with exploratory interfaces is also different.

This implies that location with respect to centralization (*i. e.*, central or peripheral) introduces an individual difference to be accounted for in system design. If this difference is not considered, user exploration can be biased toward central locations. While centralized exploration is not inherently bad (*i. e.*, there could be genuine interest in central location content given

contingency or current news), the purpose of decentralization is to make users aware of the geographical diversity of the information space. What we do on-line, influences our off-line lives, and thus, such awareness would help users in today's "culture of real virtuality" [9] to reaffirm their identities on the Web by recognizing their local traits and culture on content.

To do so in exploratory settings, we propose the following design changes in micro-blogging platforms:

1. Search results are ranked (presumably) by relevance and popularity. Centralized contexts (such as many Latin American countries) will have biased popularity metrics toward central locations. A diversity-aware algorithm (*e. g.*, *PM*) would help in this case.
2. Currently, Twitter offers two options related to geography in its search page: "Everywhere" (the default) and "Near me". We suggest to change the default value according to the user's context, *i. e.*, use "Near me" when users are from peripheral locations.
3. Since including geographical diversity in timelines is not enough, because users may not be aware of it, techniques to make them *see* this difference can help. We recommend platforms to use visualization techniques such as treemaps to display timelines. The original design by Weskamp [41] has proven to be useful when visualizing news headlines as it "allows many interesting comparisons and readings of how we differ culturally" [28]. Our results provide evidence of the good properties of treemaps, in particular in displaying the differences derived from user location and its relation to centralization.

By considering these implications and suggestions, micro-blogging platforms will allow users to *decentralize* part of their exploratory behavior. Or, they could keep a centralized exploration in a *conscious* way. The important outcome is that users will be able to choose *how* and *what* to explore.

CONCLUSIONS

In his essay "*In praise of shadows*", Junichirō Tanizaki wonders what if the fountain pen, an "*insignificant little piece of writing equipment*", would have been invented in Japan, as in spite of its insignificance it "*had a vast, almost boundless, influence on our culture*". In that line of thinking, we wonder "*what if global web platforms would have emerged in countries with severe systemic biases?*" Perhaps algorithms would have considered those biases and user interfaces would have been adapted to mitigate their effects. To this end, we studied the specific case of Chile, a highly centralized country [17]. With Chilean users, considering their geographical origin from a centralization point of view, we analyzed through carefully designed experiments the differences in users' perception of diversity when exploring geographically diverse timelines.

In a user study with labeling tasks, we found that centralization induces differences in perception of what is interesting, what is informative, and what is diverse. In particular, only users from the centralized location were able to identify the geographical diversity present in the constructed timelines. Inspired by a visual design by Weskamp [41], we used information visualization to make users *aware* of such diversity,

addressing what we called the *diversity- and representation-awareness* problem. We deployed this design on the Web to display timelines generated with a diversity-aware algorithm, and spread information about it using the social bot *@todocl*.

By analyzing logged interaction data, we observed that users behave differently according to their geographical origin. We identified different interaction signals in central and peripheral users. Furthermore, we also observed that the proposed design improves exploration and diversity- and representation-awareness, regardless of geographical origin of users. These results were used to define design implications and suggestions for micro-blogging platforms. We believe that consideration of our implications will have noticeable and positive consequences on information access by users in centralized contexts, which are not uncommon in Latin America [17].

As of September 2015, *@todocl* has more than 5,000 followers. We see this as an indication of the need to provide content aggregators like *AT*, that focus on diversification of specific social-context related aspects [8].

Limitations and Future Work. Critics might rightly say that our experiments suffer from biased user sampling and the lack of counterbalanced designs. Currently, given centralization and the unequal population distribution, finding users from *NOT-RM* is hard, because *RM*'s population is greater and has more access to the Internet in comparison to other locations, making it more difficult to find users willing to participate from non-central locations. Our snowball sampling method provided a way to find a large enough population to gain important insights in the first study. In the second study, all users mentioned by *@todocl* were geographically diverse, and those users retweeted *@todocl*'s tweets, improving the representativity of the sample. In terms of design, we could not validate users until the end of each study. Thus, because several users would have been discarded, the ideal counterbalanced designs would have proven to be unbalanced in practice.

The Web is increasingly being more accessed from mobile devices than from desktops. In our case, of 1335 detected mobile users, 68% could be geolocated using the IP address. Considering those users in a new mobile-friendly version of our site will allow us to study differences in a mobile context.

Finally, we will approach further studies from two perspectives. On the one hand, the lack of qualitative feedback in our second study imply that we cannot fully demonstrate a solution to diversity-awareness, although we believe that we have found evidence toward that conclusion. On the other hand, our results apply to one country only. Thus, we will seek to generalize our results by replicating our studies in other Latin American countries.

Source Code. We have published the code needed to replicate this experiment in <https://github.com/carnby/aurora>.

Acknowledgments. We thank Pajarito, Denis Parra, Luz Rello, and Sergio Salgado for their helpful comments. This work was partially funded by Grant TIN2012-38741 (Understanding Social Media: An Integrated Data Mining Approach) of the Ministry of Economy and Competitiveness of Spain.

References

1. Abel, F., Hauff, C., Houben, G.-J., Stronkman, R., and Tao, K. Semantics + filtering + search = Twitcident. Exploring information in social web streams. *Proceedings of the 23rd ACM Conference on Hypertext and Social Media*. 2012, 285–294.
2. Aiello, L. M., Deplano, M., Schifanella, R., and Ruffo, G. People are strange when you're a stranger: impact and influence of bots on social networks. *Links*, 697(483,151), 2012: 1–566.
3. Archambault, D., Greene, D., Cunningham, P., and Hurley, N. ThemeCrowds: multiresolution summaries of Twitter usage. *Proceedings of the 3rd international workshop on Search and mining user-generated contents*. 2011, 77–84.
4. Atienza, M., and Aroca, P. Concentración y crecimiento en Chile: una relación negativa ignorada. *EURE (Santiago)*, 38(114), 2012: 257–277 (In Spanish. Title translation: Concentration and growth in Chile: an ignored negative relation).
5. Bostock, M., Ogievetsky, V., and Heer, J. D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12), 2011: 2301–2309.
6. Bruls, M., Huizing, K., and Van Wijk, J. J. Squarified treemaps. Springer, 2000.
7. Butler, J. Precarious life: the powers of mourning and violence. Verso, 2006.
8. Buxton, B. Sketching user experiences: getting the design right and the right design. Morgan Kaufmann, 2010.
9. Castells, M. The rise of the network society: the information age: economy, society, and culture. John Wiley & Sons, 2011.
10. Center, P. R. Internet seen as positive influence on education but negative on morality in emerging and developing nations. <http://www.pewglobal.org/2015/03/19/internet-seen-as-positive-influence-on-education-but-negative-influence-on-morality-in-emerging-and-developing-nations/>. 2015.
11. Cheng, Z., Caverlee, J., and Lee, K. You are where you tweet: a content-based approach to geo-locating Twitter users. *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*. 2010, 759–768.
12. Crabtree, A., Chamberlain, A., Grinter, R. E., Jones, M., Rodden, T., and Rogers, Y. Introduction to the special issue of “The turn to the wild”. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 20(3), 2013: 13.
13. De Choudhury, M., Counts, S., and Czerwinski, M. Identifying relevant social media content: leveraging information diversity and user cognition. *Proceedings of the 22nd ACM Conference on Hypertext and Hypermedia*. 2011, 161–170.
14. Eaton, K. Designing subnational institutions regional and municipal reforms in postauthoritarian Chile. *Comparative Political Studies*, 37(2), 2004: 218–244.
15. Faridani, S., Bitton, E., Ryokai, K., and Goldberg, K. Opinion space: a scalable tool for browsing online comments. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2010, 1175–1184.
16. Freeman, L. C. A set of measures of centrality based on betweenness. *Sociometry*, 1977: 35–41.
17. Galiani, S., and Kim, S. Political centralization and urban primacy: evidence from national and provincial capitals in the Americas. In: *Understanding long-run economic growth: geography, institutions, and the knowledge economy*. University of Chicago Press, 2008, 121–153.
18. Garcia-Gavilanes, R., Mejova, Y., and Quercia, D. Twitter ain't without frontiers: economic, social, and cultural boundaries in international communication. *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 2014, 1511–1522.
19. Garcia-Gavilanes, R., Quercia, D., and Jaimes, A. Cultural dimensions in Twitter: time, individualism and power. *International AAAI Conference on Weblogs and Social Media*. 2013.
20. Graells-Garrido, E., and Lalmas, M. Balancing diversity to counter-measure geographical centralization in microblogging platforms (*short paper*). *25th ACM Conference on Hypertext and Social Media*. 2014.
21. Hecht, B., Hong, L., Suh, B., and Chi, E. H. Tweets from Justin Bieber's heart: the dynamics of the location field in user profiles. *Proceedings of the 2011 annual conference on Human Factors in Computing Systems*. 2011, 237–246.
22. Hofstede, G., Hofstede, G. J., and Minkov, M. Cultures and organizations, software of the mind: intercultural cooperation and its importance for survival (3rd edition). McGraw-Hill Professional, 2010.
23. Kayan, S., Fussell, S. R., and Setlock, L. D. Cultural differences in the use of instant messaging in Asia and North America. *Proceedings of the 2006 20th anniversary conference on Computer Supported Cooperative Work*. 2006, 525–528.
24. Kollman, K. Perils of centralization: lessons from church, state, and corporation. Cambridge University Press, 2013.
25. Krugman, P. The role of geography in development. *International regional science review*, 22(2), 1999: 142–161.
26. Kulshrestha, J., Kooti, F., Nikraves, A., and Gummadi, K. P. Geographic dissection of the Twitter network. *Proceedings of the 6th International AAAI Conference on Weblogs and Social Media*. 2012.

27. McCullagh, P. Regression models for ordinal data. *Journal of the royal statistical society. Series B (Methodological)*, 1980: 109–142.
28. Meirelles, I. Design for information: an introduction to the histories, theories, and best practices behind effective information visualizations. Rockport Publishers, 2013.
29. Merton, R. K., and Sztompka, P. On social structure and science. University of Chicago Press, 1996.
30. Ministerio de Desarrollo Social, G. d. C. CASEN Survey. [http : / / observatorio . ministeriodesarrollosocial . gob . cl / casen _ obj . php](http://observatorio.ministeriodesarrollosocial.gob.cl/casen_obj.php). 2011.
31. Munson, S. A., Zhou, D. X., and Resnick, P. Sidelines: an algorithm for increasing diversity in news and opinion aggregators. *International Conference on Weblogs and Social Media*. 2009.
32. Newman, M. E. A measure of betweenness centrality based on random walks. *Social networks*, 27(1), 2005: 39–54.
33. Park, S., Kang, S., Chung, S., and Song, J. NewsCube: delivering multiple aspects of news to mitigate media bias. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2009, 443–452.
34. Quattrone, G. A., and Jones, E. E. The perception of variability within in-groups and out-groups: Implications for the law of small numbers. *Journal of Personality and Social Psychology*, 38(1), 1980: 141.
35. Quercia, D., Capra, L., and Crowcroft, J. The social world of Twitter: topics, geography, and emotions. *The 6th International AAAI Conference on Weblogs and Social Media*. 2012.
36. Quercia, D., Ellis, J., Capra, L., and Crowcroft, J. Tracking gross community happiness from tweets. *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*. 2012, 965–968.
37. Rout, D., Bontcheva, K., Preotiu-Pietro, D., and Cohn, T. Where’s@ wally?: a classification approach to geolocating users based on their social ties. *Proceedings of the 24th ACM Conference on Hypertext and Social Media*. 2013, 11–20.
38. Rubin, M., and Badea, C. They’re all the same!... but for several different reasons: a review of the multi-causal nature of perceived group variability. *Current Directions in Psychological Science*, 21(6), 2012: 367–372.
39. Takhteyev, Y., Gruz, A., and Wellman, B. Geography of Twitter networks. *Social Networks*, 34(1), 2012: 73–81.
40. Valenzuela, S., Arriagada, A., and Scherman, A. Facebook, Twitter, and youth engagement: a quasi-experimental study of social media use and protest behavior using propensity score matching. *International Journal of Communication*, 8, 2014: 25.
41. Weskamp, M. newsmap. [http : // newsmap . jp](http://newsmap.jp). 2004.
42. Ziegler, C.-N., McNee, S. M., Konstan, J. A., and Lausen, G. Improving recommendation lists through topic diversification. *Proceedings of the 14th International Conference on World Wide Web*. 2005, 22–32.