

Characterization of Local Attitudes Toward Immigration Using Social Media

Yerka Freire-Vidal
Social Complexity Research Center
Universidad del Desarrollo
Santiago, Chile
yfreirev@udd.cl

Eduardo Graells-Garrido
Data Science Institute
Universidad del Desarrollo
Santiago, Chile
egraells@udd.cl

ABSTRACT

Migration is a worldwide phenomenon that may generate different reactions in the population. Attitudes vary from those that support multiculturalism and communion between locals and foreigners, to contempt and hatred toward immigrants. Since anti-immigration attitudes are often materialized in acts of violence and discrimination, it is important to identify factors that characterize these attitudes. However, doing so is expensive and impractical, as traditional methods require enormous efforts to collect data. In this paper, we propose to leverage Twitter to characterize local attitudes toward immigration, with a case study on Chile, where immigrant population has drastically increased in recent years. Using semi-supervised topic modeling, we situated 49K users into a spectrum ranging from in-favor to against immigration. We characterized both sides of the spectrum in two aspects: the emotions and lexical categories relevant for each attitude, and the discussion network structure. We found that the discussion is mostly driven by Haitian immigration; that there are temporal trends in tendency and polarity of discussion; and that assortative behavior on the network differs with respect to attitude. These insights may inform policy makers on how people feel with respect to migration, with potential implications on communication of policy and the design of interventions to improve inter-group relations.

ACM Reference Format:

Yerka Freire-Vidal and Eduardo Graells-Garrido. 2019. Characterization of Local Attitudes Toward Immigration Using Social Media. In *Companion Proceedings of the 2019 World Wide Web Conference (WWW '19 Companion)*, May 13–17, 2019, San Francisco, CA, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3308560.3316455>

1 INTRODUCTION

Migration is a phenomenon faced by many countries, which brings a variety of effects; both in the population from which it emigrates and in the receiving population. One of the effects that worries many countries is intolerance and hostile attitudes toward immigrants. These attitudes have been the focus of many research studies, some of which are focused on individual-level psychological and socio-economic factors [8, 47], and others on the contact between immigrant population and locals [7, 30, 32]. The main methods

used in these studies are based on context specific surveys, which makes replication in others societies or countries difficult. The theories that explain the type of attitudes of locals interacting with immigrants can be summarized in two: the Intergroup Contact Theory [1], and the Integrated Threat Theory [40, 49]. The former states that people support multiculturalism and integration. The latter, that people think that immigrants will bring negative effects for their society, including competition for jobs and public services, worsening of the national economy, increase in crime, and the arrival of diseases. Particularly, the attitudes explained by the threat theory can lead to acts of violence, discrimination, and abuse; thus, it is important to understand what factors enhance such attitudes.

However, measuring attitudes is costly and impractical under dynamic scenarios. The most frequent methods are surveys, which are difficult and costly to implement. In this paper, we propose to make use of the information that people publish in Twitter as a proxy of their attitudes toward immigration. It is common to find reactions and attitudes through posts in these platforms, where people express their ideas and opinions voluntarily. We propose to define a spectrum of attitudes based on the two aforementioned theories, and to classify users and tweets into that spectrum. We do so with a semi-supervised topic modeling technique named Topic-Supervised Non-Negative Matrix Factorization [39]. TS-NMF works in a semi-supervised way because some users can be labeled as belonging to each extreme of the spectrum, something that we do with custom-built lexicons for each theory.

We perform a descriptive case study on the Chilean society, because Chile is one of the countries in which migration has reached unprecedented volume in recent years. The statistics show that immigrant population has increased from 0.8% in 1992 to 4.35% in 2017; and where 66.7% of immigrants declare to have arrived mainly in 2016 [31]. For this, Chileans have developed diverse perceptions regarding the number of immigrants in the country and the phenomenon itself. To measure them with our proposed method, we collected more than 206K tweets that discuss immigration in Chile, written by more than 49K users during the year 2017. After inferring user and tweets positions in the spectrum, we performed lexical and network analysis with respect to the spectrum position. In the lexical analysis, we used the “Linguistic Inquiry and Word Count” (LIWC) lexicon [43], typically employed to characterize cognitive and emotional differences in discourse [11, 17, 26]. To analyze the network structure, we estimated the polarization of the retweet and mention networks between users.

As main results, we observed that most of the discussion toward migration in Chile is targeted at Haitian migration, even though other countries have a larger share of the population. We found

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '19 Companion, May 13–17, 2019, San Francisco, CA, USA

© 2019 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-6675-5/19/05.

<https://doi.org/10.1145/3308560.3316455>

lexical differences in how each attitude discussed migration, and those differences were consistent with theories. For instance, social-related words were correlated with empathetic attitudes, job- and money-related words were correlated with threatening attitudes. In the network, the retweet network was polarized, in coherence as predicted by other studies regarding political discussion [10, 20]. Finally, we notice that the amount and tendency of the tweets (the latter reflects the attitude towards immigration) seems to be influenced by relevant news events on national migration issues. These results can inform public policy designers to improve inter-group relations in the country, as well as increasing the understanding of how people feel regarding an important aspect of globalization.

In summary, the contribution of this paper is two-fold. We proposed a methodology to characterize local attitudes toward migration from tweets. Then, we performed a descriptive case study in Chile using this method, obtaining results that are coherent with social theory, with added depth based on the rich information that can be extracted from Twitter.

This paper is structured as follows. Section 2 discusses the related work. Section 3 describes the social theories that guided our analysis. Section 4 describes the data set we analyzed. Section 5 describes the methodology. Section 6 describes the results of applying the methodology to the data set. Section 7 discusses the implications of our work. Finally, Section 8 states our conclusions.

2 RELATED WORK

Migration is a widely studied topic because there are many issues associated to this phenomenon. Some researchers have focused on studying the economic impacts related to migration [23–25, 47], others on social cohesion [30, 32, 33, 48]. Within these studies, those who have focused on integration [36] and racism/xenophobia stand out [7, 27, 44]. Our work seeks to contribute in the latter area, mainly due to the subject of our case study, Chile, a society that in a short time has faced a massive influx of immigrants. Migration in Chile has been a national issue, causing controversy in presidential elections, news, and municipal institutions. However, measuring attitudes is not a simple problem, nor a solved one. Twitter is currently a platform widely used in studies of human behavior, since it provides a valuable source of data. Studies that have used Twitter have allowed to reveal socio-cultural characteristics of users or societies, including the level of integration of immigrants in a city [36], attitudes in response to triggering events, such as terrorist attacks [13], the influence of culture in personal actions [16], political polarization [16], personality traits [46], and personality differences between democrats and republicans [52].

Given this body of research, we propose that Twitter can be used as a proxy to understand human behavior; in our case, the attitudes of Chileans regarding immigration.

3 SOCIAL THEORIES

The attitudes toward immigration are varied and depend on economic, socio-cultural and psychological factors. In this context, psychology and sociology have defined theories that explain the attitudes exhibited by people, who belong to different groups, when interacting with others: the Intergroup Contact Theory [1], and the




Figure 1: Weekly distribution of tweets about immigration in Chile.

Integrated Threat Theory [40, 49]. The attitudes toward immigration are a particular case explained by these theories.

Intergroup Contact Theory. Developed in the book “The Nature of Prejudice” by Gordon W. Allport [1], it postulates that prejudices are reduced when the interaction between different groups meets the following conditions: 1) groups are on equal terms; 2) they have common goals; 3) there is cooperation; and 4) there is support from formal and/or informal institutions. The theory states that intergroup contact reduces the fear and anxiety that exists when people interact with an unknown group [51], and that it promotes empathy and understanding towards the foreign group [50].

This theory has been used to ground several studies: contact between white and black people [7], heterosexuals and homosexuals [27, 28], minority religious groups [42], and locals and immigrants [30]. All these studies conclude that contact improves relationships between groups.

Integrated Threat Theory. In contrast to the Intergroup Contact Theory, the Integrated Threat Theory argues that contact between disparate groups provokes perceptions of threat and contempt [40, 49], for instance, due to competition for work and economic resources [15, 22]. Furthermore, the threat does not have to be real, it can be subjective or fictitious [33].

This theory postulates that, when the interaction conditions are not optimal, the contact between different groups will provoke conflicting and hostile relationships. The concept of “contact” is not limited only to physical contact, it can also be indirect [14], imagined [12], and electronic [2, 53].

Both theories tell us what to search when we look attitudes toward immigration: attitudes motivated by empathy, in favor of immigration; and attitudes motivated by threat, against immigration. As such, we will assume that there are two attitudes, which we label as *empathy* and *threat*.

4 DATA SET DESCRIPTION

In this section we describe our data set of posts from Twitter about migration in Chile.

Twitter is a micro-blogging platform, where users publish tweets (posts) with a maximum of 280 characters. Users may follow others,



Figure 2: Wordcloud of most frequent words in the dataset, after removing stopwords. Color is assigned according to the following categories: words, hashtags, mentions, and URLs.

to see their tweets in their own timelines. Tweets may mention other users, quote other tweets, or retweet another tweet to share it with one's audience. Users can report a screen name, a full name (which can be real or fictitious), a location (real, fictitious, or empty), and a small autobiography, among other attributes. To collect tweets that talk about immigration in Chile we used the Twitter Streaming API using system designed to crawl Chilean tweets [21]. The query parameters were keywords related to immigration (e.g., inmigración, inmigrante, fronteras, racismo, etc.), and origin countries with their respective demonyms (e.g., Haití-haitianos/as, Venezuela-venezolanos/as, Perú-peruanos/as, etc.). Given how generic some of these keywords are, particularly regarding the context of political issues of neighbouring countries, and the presidential elections held in Chile during November and December, we performed extensive manual clean-up of the data set.

In total, our data set is comprised by 206,353 tweets that discuss immigration in Chile during 2017, written by 49,346 users. Figure 1 shows the weekly volume of tweets. As seen on the figure, the amount of tweets has a slight positive trend. Two peaks draw our attention: July 31th, when the news reported a case of an Haitian citizen with Leprosy; and November 19th, when an Haitian citizen rescued a woman who fell from the ninth floor of a building.

Regarding content, Figure 2 shows the most frequent words, after removing stopwords and accents. One can see that words such as Haiti and Haitianos are more relevant than other countries name or demonyms, despite the fact that the largest immigrant population comes from Perú, Colombia and Venezuela [31]. Also, Santiago and Antofagasta are two frequent keywords, the two cities with the largest immigrant population [31]. Other relevant words that appear are: “gobierno” (government), “carabineros” (police-men), “Piñera” (current president, and presidential candidate in 2017) and “proyecto” (project); possibly because during the year an immigration reform was being discussed.

We explored the data set to seek for words, phrases, and hashtags that could be mapped to the *empathy* and *threat* attitudes. In *empathy* we chose terms that indicated that immigrants are welcome and will be received in equal conditions (e.g., “we are all immigrants”). In *threat* we chose terms and words that showed that immigrants are not welcome and qualified them negatively (e.g., “illegal immigrants”). Table 1 shows some examples of the the

Table 1: Examples of training terms for each attitude.

Attitude	Training Terms
<i>Empathy</i>	#todossomosmigrantes, #stopxenophobia, #chilesinbarreras, #chileterecibe, #bienvenidosmigrantes, @oimchile, bienvenidos a chile, #derribandomuros, @sjmchile, ...
<i>Threat</i>	#vendepatria, #nomasinmigrantes, #nomasilegales, #inmigrantesilegales, inmigrantes delincuentes, inmigracion descontrolada, indeseables, ...

terms we associated to both attitudes. These labeled terms are not necessarily frequent, however, the methodology that we describe in the next sections allows to propagate these labels through a topic model.

5 METHODOLOGY

In this section we describe how to characterize users and tweets according to their attitude toward immigration. We define how to apply machine learning techniques to user profiles to derive user-attitude and term-attitude associations. Then, we define how to characterize attitudes from sentiment, lexical and network perspectives.

5.1 Attitudes and Topic Modeling

Topic models are a family of techniques used to discover the underlying semantic structure of a corpus by identifying and quantifying the importance of representative themes in all documents [6]. Topic models assume that each text document is generated by a set of topics which have a determined distribution. At the same time, each topic is defined by a set of words, which also have a particular distribution for each topic.

A popular topic modeling technique is Non-negative Matrix Factorization [38]. NMF works by constructing a k -rank factorization of a positive document-term matrix V into $W \times H$. Matrices W and H are estimated by minimizing the following objective function:

$$D_{NMF}(W, H) = \|V - W \times H\|_F^2, \quad W, H \geq 0, \quad (1)$$

where $\|\cdot\|_F$ is the Frobenius norm. In topic modeling, k , W and H have a special interpretation: k is the number of topics, W_{ij} quantifies the relevance of topic j in document i , and H_{ij} quantifies the relevance of term j in topic i .

Typical topic modeling applications select different numbers of k based on metrics such as perplexity. However, the meaning of topics is not always interpretable, as the factorization may follow latent patterns not necessarily aligned with human expectations. Based on the social theories described in Section 3, we propose to guide the learning procedure to seek for two topics: one that represents *empathy*, and another that represents *threat*. In such cases, supervised methods could be employed, however, these methods require a fully labeled data set, not available in our case. Since it is possible to map specific terms (words, phrases, hashtags, URLs, *etc.*) into these two topics, we propose to use a semi-supervised version of NMF known as Topic-Supervised NMF [39]. TS-NMF defines the

minimization problem as follows:

$$D_{TS}(W, H) = \|V - (W \circ L)H\|_F^2, \quad W, H \geq 0, \quad (2)$$

where \circ is the Haddamard product operator, and L is a supervision matrix, defined as $L_{ij} = 1$ if topic j contributes to the document i , and $L_{ij} = 0$ if the topic j does not contribute to the document i . Thus, TS-NMF allows to provide examples of documents labeled with known topics, and to restrict the latent representation of the corpus to align with the labeled examples.

In our context, we work with user profiles, *i.e.*, the concatenation of tweets by a single user is one document. As terms we consider hashtags, mentions, URLs, and n -grams with n up to four. This allows us to define how specific phrases are mapped to each topic. The user corpus is represented as a document-term matrix D weighted with TF-IDF [4], and then row-normalized with L2 norm. To label users in the supervision matrix, we construct a list of seed terms for each theory. Then, for each row in D we estimate a preliminary attitude score for each topic, by adding the values of the cells of the corresponding seed terms. All users with a score above a certain threshold are labeled with the corresponding topic. In our experiments, we defined a threshold of 0.25, implying that only users who strongly used the seed terms of each topic were labeled.

As result, we obtain $D = U \times T$, where the rank of U and T is two. In our context, each topic is an attitude, the matrix U contains the user-attitude associations, and the matrix T contains the term-attitude associations (transposed). We interpret these associations as probabilities.

5.2 Attitude Tendency and Polarity

To characterize attitudes, we calculate two metrics common in the sentiment analysis literature to measure the leaning and amount of sentiment: tendency and polarity [35]. Tendency is defined as:

$$\text{tendency}(u) = P(\text{empathy} | u) - P(\text{threat} | u), \quad (3)$$

where, $P(\text{attitude} | u)$ is the association between user u and the corresponding attitude. Note that the definition is analog for terms. For tweets, tendency is defined as:

$$\text{tendency}(\text{tweet}) = \sum_{\text{term} \in \text{tweet}} \text{tendency}(\text{term}). \quad (4)$$

Note that tendency values close to zero do not imply a neutral attitude, as there could be non-zero contributions in both topics. To clarify this fact, we consider attitude polarity as the amount of associations to both attitudes, defined for users as:

$$\text{polarity}(u) = P(\text{empathy} | u) + P(\text{threat} | u). \quad (5)$$

The definition for terms is analog. For tweets, polarity is defined as:

$$\text{polarity}(\text{tweet}) = \sum_{\text{term} \in \text{tweet}} \text{polarity}(\text{term}). \quad (6)$$

In this way, tendency will allow us to group users/tweets (according to their attitude), while polarity will allow us to measure the intensity of the discussion (how polarized is the attitude).

5.3 Lexical Characterization

The previous metrics give an overview of user and tweet attitudes. The next step is to characterize grouped tweets belonging to each

attitude according to their tendency. To do so, we use a psycholinguistic lexicon named “Linguistic Inquiry and Word Count” [43]. LIWC is a lexicon used to study emotional, cognitive and structural components contained in a text. In its Spanish version, it contains 7,515 words classified in one or more of 72 categories. Categories are classified into four dimensions: 1) standard linguistic processes (*e.g.*, articles, prepositions, pronouns, *etc.*); 2) psychological processes (*e.g.*, positive and negative emotions); 3) relativity (*e.g.*, time, verb tense, motion, space); and 4) personal matters (*e.g.*, sex, death, home, occupation, *etc.*). LIWC categories are organized hierarchically, for instance, all words related to the category *anger* are also organized in the categories of *negative emotions* or *affect* words.

We seek to estimate the association of tweets by tendency groups to LIWC categories. After classifying tweets into groups, we estimate how associated the words in LIWC are to each group. Note that specific events may entice a more active discussion by either group, increasing the amount of tweets, thus, we need a way to control the association with these activity patterns. In previous work, this has been done to estimate gross community metrics with z-scores [34, 45]. In our case, the definition is as follows:

$$Z_{l,t'} = \frac{P_{l,t'} - \mu_l}{\sigma_l}, \quad (7)$$

where, $Z_{l,t'}$ is the association of LIWC category l with the tendency t' , $P_{l,t'}$ is the mean of fraction of words in l in each tweet with tendency t' , μ_l is the mean of fraction of words in l in all tweets, and σ_l is the standard deviation of the fraction of words in l in all tweets. Hence, this relative metric allows us to compare behavior between groups, by controlling for external variability.

5.4 Network Assortativity

The previous definitions capture the behavior in expression, however, the social aspect of Twitter allows to also capture network behavior. We focus on two different networks: the mention network, related to discussion, and the retweet network, related to information diffusion. In both networks, node are users, and links are weighted relations between users. Each node has as attributes its associations to each attitude. In the mention network, a directed link between users u_1 and u_2 exists if u_1 mentions u_2 in one or more tweets. The link weight is the number of times this happens. In the retweet network, a directed link between users u_1 and u_2 exists if u_1 republishes content by u_2 . The link weight is the number of times that one user retweets another. These kind of networks are commonly analyzed to understand polarization [10]. To be able to analyze connectivity, we will focus on the Largest Strongly Connected Component of each network.

To analyze the networks structure, we estimate the assortativity coefficient with respect to each attitude. The assortativity coefficient is the Pearson correlation coefficient of numerical attributes between pairs of linked nodes (this numerical attributes are the attitudes given by the model). It measures the similarity of connections in the graph with respect to the given numeric attribute [41]. Hence, the assortativity coefficient measures whether people relations are homophilic with respect to attitude. This behavior is commonly found in networks [5], and it has been documented in Twitter political discussion [10], including in Chile [20].



Figure 3: Most associated words to each attitude according to the TS-NMF model. Note that only single words are displayed, to avoid repetition in n-grams.

In the next section we apply this methodology to the data set described in Section 4, covering an entire year of discussion about immigration in Chile.

6 RESULTS

Here we present the results of applying the methodology from Section 5 to the data set from Section 4.

Term Associations. Figure 3 shows the association of words with each attitude, empathy on top, threat on bottom. One can see that words associated to empathy include “integración” (integration), “salud” (health), and “educación” (education), reflecting their empathetic attitude. Words associated to threat include “delinquentes” (delinquents), “control” (control), and “ilegales” (illegals), reflecting a feeling of threat. Also, empathy group uses the word “Migrantes” (migrants) and threat group uses “Inmigrantes” (immigrants), which can be interpreted as that the empathy group is concerned about the general phenomenon (migration includes emigration and immigration), while the threat group only for the particular phenomenon (immigration).

Tendency and Polarity. Figure 4 shows the distribution of tendency and polarity for users. One can see that the distributions are fairly symmetric, with peaks in the center of the distribution. Figure 5 shows the tendency and polarity of tweets during the year under study, estimated using LOWESS. One can see that the tendency trend exhibits two interesting periods, before and after the news about the Leprosy case of an Haitian in July 31th. In the first




Figure 4: Top: tendency distribution for users. Bottom: polarity distribution for users.

period, tendency is slightly negative (threat), with an arguably low variability. In the second period, variability increases, and a small negative trend appears, even though at a point in time it reaches its maximum value (*i.e.*, maximum empathy) at the beginning of October. This could be explained by a news event reported in October 6th, about a Colombian citizen that gave birth on the street because a taxi driver expelled her from his car.

It is interesting that both news are related with the Integrated Threat theory and Intergroup Contact theory, respectively. On the one hand, the first event shows the immigrant as a threat, being a possible source of contagion of a disease (Leprosy). On the other hand, the second event shows the immigrant being a victim of violence and discrimination, which arguably makes people more empathetic. Regarding polarity, the trend exhibits a gradual increase in time, with two interesting peaks. The first one reflects the Leprosy case, and the second one reflects the presidential elections, where migration was a common topic in discussion.

LIWC Analysis. Figure 6 shows the differences of cognitive and emotional categories from LIWC in tweets grouped by tendency: *empathy* contains all tweets with tendency ≥ 0 ; *threat*, otherwise. For each category and group, we estimated the z-score for all tweets each month. As a general observation, one can see that both groups tend to have opposite behaviors. For instance, tweets in the empathy group are positively associated to the *sociability*, *family*, and *positive emotions* category more than tweets in the threat group. Conversely, tweets in the threat group are positively associated with *money*, *job*, and *inhibition* categories. This could be explained by the threat theory, as immigrants can be perceived as an economic threat and labor competition. Also, inhibition category can be interpreted by the desire to prohibit the arrival of more immigrants or to prevent them from accessing social benefits.




Figure 5: Trend distributions (top: tendency, bottom: polarity) for all tweets in the data set. Each tweet is a point, the x -position encodes its publication date, the y -position encodes its tendency or polarity. The line is the LOWESS interpolation of tendency and polarity.

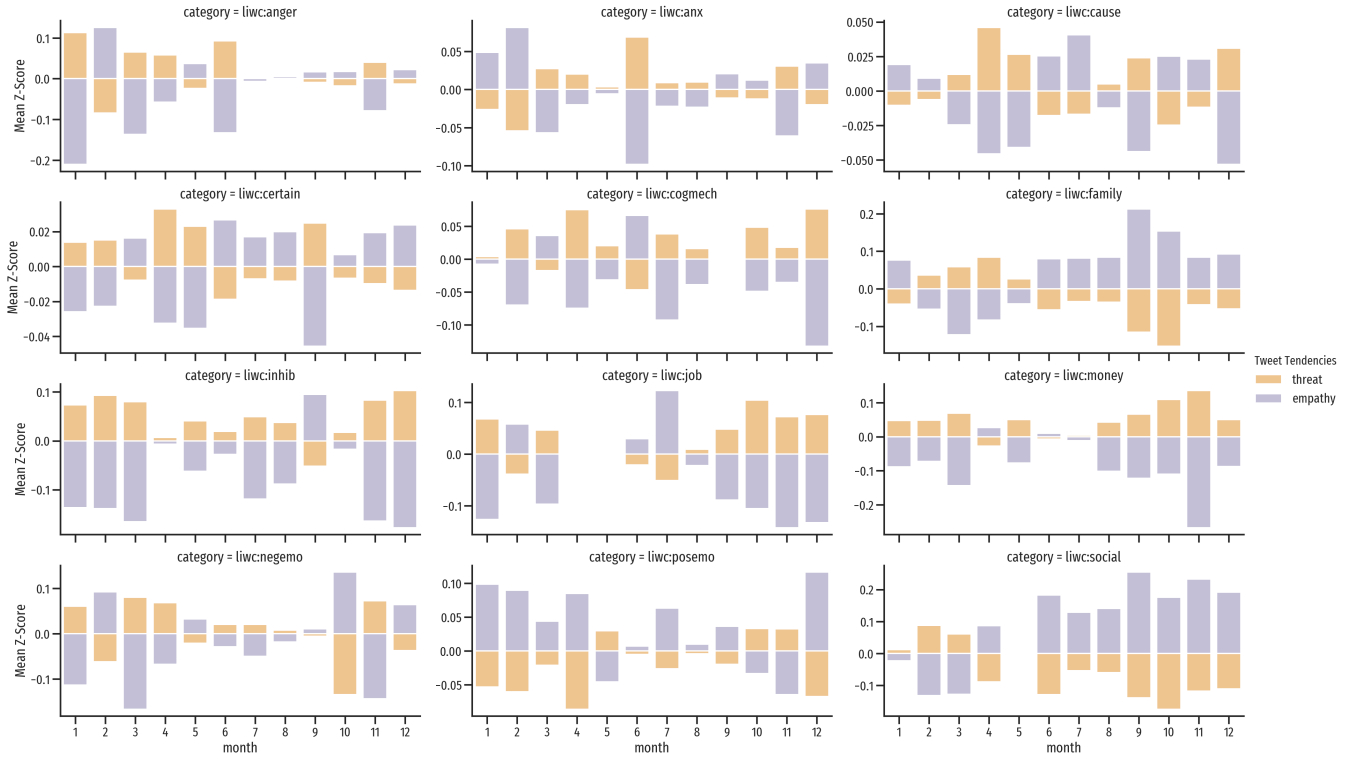


Figure 6: Association between attitudes (*empathy* and *threat*) and LIWC categories, per month. Each bar represents the association between groups, estimated with z -scores of fraction of words from each LIWC category and all other words. Purple bars indicate empathy associations, orange bars indicate threat associations.

Mention and Retweet Networks. The largest SCC of the retweet network has 1,239 nodes and 6,441 edges, while the largest SCC of the mention network has 1,868 nodes and 10,201 links. Figure 7 visualizes both networks using Hierarchical Edge Bundling [29]. This method allows us to make explicit the adjacency relations between users, as similar edges are bundled to decrease visual clutter. In the figure, each link is colored according to tendency of

the source node (purple: empathy group, orange: threat group). Note that the visual encoding makes explicit the community structure in the retweet network and the heterogeneity of the mention network.

The assortativity coefficient for the retweet network are 0.26 (empathy) and 0.14 (threat), implying that homophilic behavior exists, but it is not as strong as in other topics (for instance, the discussion about abortion in Chile is greater [20]), and it is not

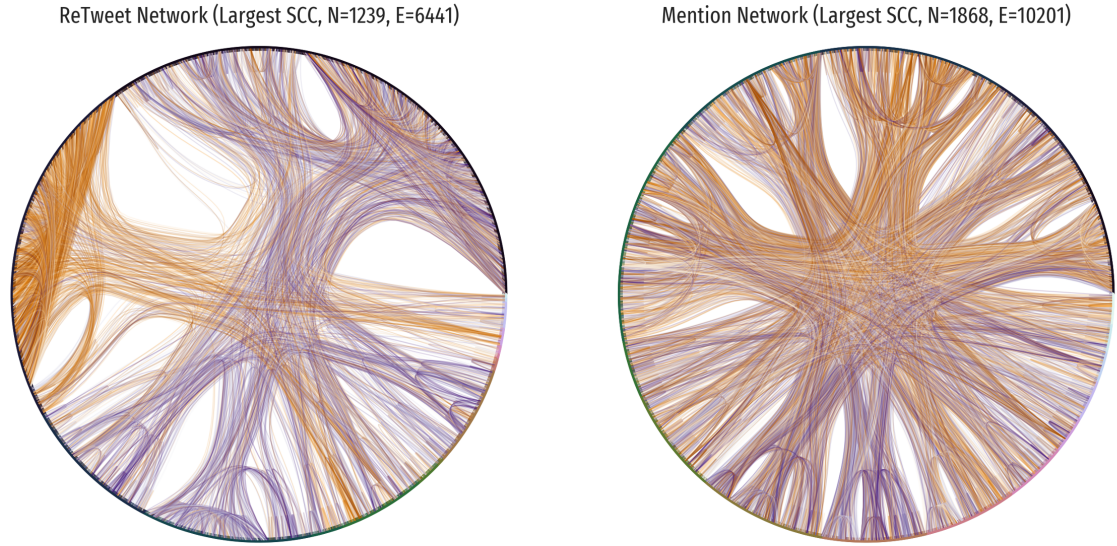


Figure 7: ReTweet Network (left) and Mention Network (right). Each node is a circle in the outside, sorted according to the connectivity patterns to other nodes. Edges are lines that join nodes, where color is the attitude of the source node (purple: empathy, orange: threat). This encoding allows to group edges that are similar in terms of connectivity between groups.

equal in both groups. As hinted by the visualization, in the mention network the results are small: 0.06 (empathy) and 0.08 (threat). Thus, the retweet network is more segregated than the mention network. This could be explained because retweets are expected to be seen by all followers, and are a key factor in information diffusion, while mentions and replies are not. For instance, one user may send tweets to another holding an opposite position, but if there is no reply, then the interaction is not meaningful.

7 DISCUSSION AND FUTURE WORK

Migration is a controversial issue in Chile, and, although there are some studies about Chileans attitudes toward immigration [9, 37], they do not cover recent migration patterns. To complement knowledge about this topic, we defined a way to classify and measure attitudes, enabling to study the dynamics of perception with respect to immigration and performed a descriptive study of how immigration is perceived in Chile, according to Twitter discussion.

Our results may inform policy and intervention design, as it quantifies how people feel and communicate with respect to immigration. This is relevant, as there exists several contact strategies to improve relationships between social groups [44]. For instance, the discussion we analyzed is mostly targeted at Haitian migration. A majority of them is from Afro-Haitian descent, an ethnicity that was almost non-present in Chile.

There are two key aspects that need further exploration, and that limit the scope of our results: the representativity of Twitter, and the validation of the TS-NMF model. In terms of representativity, Twitter is a biased sample of the population [3]. As such, our results only cover this sample, even though it is not known to which degree nor to which sub-populations it represents. Having these biases into

account will surely improve the interpretation of results. However, one aspect that needs to be considered is that Twitter is within the most popular applications in Chile [18], and that it reflects some cultural aspects, such as the country's centralization [19]. In terms of validation, the lack of ground truth or approximate measures of the problem stands in the way of effectively measuring the model accuracy, leaving us only with a qualitative evaluation.

Besides working on the limitations of our approach, there are two lines of future work that we devise. On the one hand, it would be relevant to understand the relationship between attitudes and actual presence of immigrants in a place. This would provide a way to measure real and imagined threat attitudes [33]. On the other hand, there is a potential influence of news events in attitudes. Given the rise of *fake news* and *post-truth* media, this would provide a way to measure the effect of such phenomena on how people feel with respect to a specific issue, migration in this case.

8 CONCLUSIONS

In this paper, we have characterized attitudes toward immigration by locals in Chile. We used a semi-supervised topic modeling technique (TS-NMF [39]) to identify attitudes grounded in two social theories, the Intergroup Contact Theory [1], and the Integrated Threat Theory [40, 49]. Then, we measured differences in attitudes using psycho-linguistic lexicons and interaction networks. As result, we found consistent behaviour with respect to social theory. There is still work to do in the evaluation and representativeness of our model, including the definition of a suitable ground-truth perception to validate our proposal. We believe our results help to inform the design of public policy and interventions to improve relations between groups in a country.

REFERENCES

- [1] Gordon Willard Allport, Kenneth Clark, and Thomas Pettigrew. 1954. The nature of prejudice. (1954).
- [2] Yair Amichai-Hamburger and Katelyn YA McKenna. 2006. The contact hypothesis reconsidered: Interacting via the Internet. *Journal of Computer-Mediated Communication* 11, 3 (2006), 825–843.
- [3] Ricardo Baeza-Yates. 2018. Bias on the web. *Commun. ACM* 61, 6 (2018), 54–61.
- [4] Ricardo Baeza-Yates and Berthier Ribeiro-Neto. 2011. *Modern Information Retrieval: the concepts and technology behind search*, 2nd. Edition. Addison-Wesley, Pearson.
- [5] Pablo Barberá. 2015. Birds of the same feather tweet together: Bayesian ideal point estimation using Twitter data. *Political Analysis* 23, 1 (2015), 76–91.
- [6] David M Blei. 2012. Probabilistic topic models. *Commun. ACM* 55, 4 (2012), 77–84.
- [7] Kendrick T Brown, Tony N Brown, James S Jackson, Robert M Sellers, and Warde J Manuel. 2003. Teammates on and off the Field? Contact with black teammates and the racial attitudes of white student athletes 1. *Journal of applied social psychology* 33, 7 (2003), 1379–1403.
- [8] Peter Burns and James G Gimpel. 2000. Economic insecurity, prejudicial stereotypes, and public opinion on immigration policy. *Political science quarterly* 115, 2 (2000), 201–225.
- [9] Héctor Carvacho. 2010. Ideological configurations and prediction of attitudes toward immigrants in Chile and Germany. *International Journal of Conflict and Violence (IJCV)* 4, 2 (2010), 220–233.
- [10] Michael Conover, Jacob Ratkiewicz, Matthew R Francisco, Bruno Gonçalves, Filippo Menczer, and Alessandro Flammini. 2011. Political polarization on Twitter. *Proc. of ICWSM* 133 (2011), 89–96.
- [11] Glen Coppersmith, Mark Dredze, and Craig Harman. 2014. Quantifying mental health signals in Twitter. In *Proceedings of the Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*. 51–60.
- [12] Richard J Crisp and Rhiannon N Turner. 2009. Can imagined interactions produce positive perceptions?: Reducing prejudice through simulated social contact. *American psychologist* 64, 4 (2009), 231.
- [13] Kareem Darwish, Walid Magdy, Afshin Rahimi, Timothy Baldwin, and Norah Abokhodair. 2017. Predicting Online Islamophobic Behavior after# ParisAttacks. *The Journal of Web Science* 3, 1 (2017).
- [14] John F Dovidio, Anja Eller, and Miles Hewstone. 2011. Improving intergroup relations through direct, extended and other forms of indirect contact. *Group processes & intergroup relations* 14, 2 (2011), 147–160.
- [15] Victoria M Esses, John F Dovidio, Lynne M Jackson, and Tamara L Armstrong. 2001. The immigration dilemma: The role of perceived group competition, ethnic prejudice, and national identity. *Journal of Social issues* 57, 3 (2001), 389–412.
- [16] Ruth Garcia-Gavilanes, Daniele Quercia, and Alejandro Jaimes. 2013. Cultural dimensions in Twitter: Time, individualism and power. *Proc. of ICWSM* 13 (2013).
- [17] Roberto González-Ibáñez, Smaranda Muresan, and Nina Wacholder. 2011. Identifying sarcasm in Twitter: a closer look. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: Short Papers-Volume 2*. Association for Computational Linguistics, 581–586.
- [18] Eduardo Graells-Garrido, Diego Caro, Omar Miranda, Rossano Schifanello, and Oscar F Peredo. 2018. The WWW (and an H) of Mobile Application Usage in the City: The What, Where, When, and How. In *Companion of the The Web Conference 2018 on The Web Conference 2018*. International World Wide Web Conferences Steering Committee, 1221–1229.
- [19] Eduardo Graells-Garrido and Mounia Lalmas. 2014. Balancing diversity to counter-measure geographical centralization in microblogging platforms. In *Proceedings of the 25th ACM conference on Hypertext and social media*. ACM, 231–236.
- [20] Eduardo Graells-Garrido, Mounia Lalmas, and Ricardo Baeza-Yates. 2015. Finding intermediary topics between people of opposing views: a case study. In *Social Personalisation & Search*, Christoph Trattner, Denis Parra, Peter Brusilovsky, and Leandro Balby Marinho (Eds.). CEUR, Santiago, Chile.
- [21] Eduardo Graells-Garrido, Mounia Lalmas, and Ricardo Baeza-Yates. 2016. Encouraging diversity-and representation-awareness in geographically centralized content. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*. ACM, 7–18.
- [22] Shang E Ha. 2010. The consequences of multiracial contexts on public attitudes toward immigration. *Political Research Quarterly* 63, 1 (2010), 29–42.
- [23] Jens Hainmueller and Michael J Hiscox. 2007. Educated preferences: Explaining attitudes toward immigration in Europe. *International organization* 61, 2 (2007), 399–442.
- [24] Jens Hainmueller and Michael J Hiscox. 2010. Attitudes toward highly skilled and low-skilled immigration: Evidence from a survey experiment. *American political science review* 104, 1 (2010), 61–84.
- [25] Gordon H Hanson, Kenneth Scheve, and Matthew J Slaughter. 2007. Public finance and individual preferences over globalization strategies. *Economics & Politics* 19, 1 (2007), 1–33.
- [26] GACCT Harman and Mark H Dredze. 2014. Measuring post traumatic stress disorder in Twitter. In *ICWSM* (2014).
- [27] Gregory M Herek. 1986. The instrumentality of attitudes: Toward a neofunctional theory. *Journal of Social Issues* 42, 2 (1986), 99–114.
- [28] Gregory M Herek and John P Capitanio. 1996. “Some of my best friends” Inter-group Contact, concealable stigma, and heterosexuals’ attitudes toward gay men and lesbians. *Personality and social psychology bulletin* 22, 4 (1996), 412–424.
- [29] Danny Holten. 2006. Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE Transactions on visualization and computer graphics* 12, 5 (2006), 741–748.
- [30] Daniel J Hopkins. 2010. Politicized places: Explaining where and when immigrants provoke local opposition. *American political science review* 104, 1 (2010), 40–60.
- [31] National Statistics Institute (INE). 2018. Síntesis resultados Censo 2017. [urlhttps://www.censo2017.cl/descargas/home/sintesis-de-resultados-censo2017.pdf](https://www.censo2017.cl/descargas/home/sintesis-de-resultados-censo2017.pdf).
- [32] Seth K Jolly and Gerald M DiGiusto. 2014. Xenophobia and immigrant contact: French public attitudes toward immigration. *The Social Science Journal* 51, 3 (2014), 464–473.
- [33] Jeffrey S Kopstein and Jason Wittenberg. 2009. Does familiarity breed contempt? Inter-ethnic contact and support for illiberal parties. *The Journal of Politics* 71, 2 (2009), 414–428.
- [34] Adam DI Kramer. 2010. An unobtrusive behavioral model of gross national happiness. In *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 287–290.
- [35] Onur Kucuktunc, B Barla Cambazoglu, Ingmar Weber, and Hakan Ferhatosmanoglu. 2012. A large-scale sentiment analysis for Yahoo! answers. In *Proceedings of the fifth ACM international conference on Web search and data mining*. ACM, 633–642.
- [36] Fabio Lamanna, Maxime Lenormand, María Henar Salas-Olmedo, Gustavo Romanillos, Bruno Gonçalves, and José J Ramasco. 2018. Immigrant community integration in world cities. *PLoS one* 13, 3 (2018), e0191612.
- [37] Duncan Lawrence. 2015. Crossing the Cordillera: immigrant attributes and Chilean attitudes. *Latin American Research Review* 50, 4 (2015), 154–177.
- [38] Daniel D Lee and H Sebastian Seung. 1999. Learning the parts of objects by non-negative matrix factorization. *Nature* 401, 6755 (1999), 788.
- [39] Kelsey MacMillan and James D Wilson. 2017. Topic supervised non-negative matrix factorization. *arXiv preprint arXiv:1706.05084* (2017).
- [40] Todd D Nelson. 2009. *Handbook of prejudice, stereotyping, and discrimination*. Psychology Press.
- [41] Mark EJ Newman. 2003. Mixing patterns in networks. *Physical Review E* 67, 2 (2003), 026126.
- [42] Josef Novotny and Filip Polonsky. 2011. The Level of Knowledge about Islam and Perception of Islam among Czech and Slovak University Students: does Ignorance Determine Subjective Attitudes? *Sociologia* 43, 6 (2011), 674–696.
- [43] James W Pennebaker, Martha E Francis, and Roger J Booth. 2001. Linguistic inquiry and word count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates* 71, 2001 (2001), 2001.
- [44] Thomas F Pettigrew and Linda R Tropp. 2006. A meta-analytic test of intergroup contact theory. *Journal of personality and social psychology* 90, 5 (2006), 751.
- [45] Daniele Quercia, Jonathan Ellis, Licia Capra, and Jon Crowcroft. 2012. Tracking gross community happiness from tweets. In *Proceedings of the ACM 2012 conference on computer supported cooperative work*. ACM, 965–968.
- [46] Daniele Quercia, Michal Kosinski, David Stillwell, and Jon Crowcroft. 2011. Our Twitter profiles, our selves: Predicting personality with Twitter. In *Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third International Conference on Social Computing (SocialCom), 2011 IEEE Third International Conference on*. IEEE, 180–185.
- [47] Kenneth F Scheve and Matthew J Slaughter. 2001. Labor market competition and individual preferences over immigration policy. *Review of Economics and Statistics* 83, 1 (2001), 133–145.
- [48] Paul M Sniderman, Pierangelo Peri, Rui JP de Figueiredo Jr, and Thomas L Piazza. 2000. The Outsider: Politics and Prejudice in Italy.
- [49] Cookie White Stephan and Walter S Stephan. 2013. An integrated threat theory of prejudice. In *Reducing prejudice and discrimination*. Psychology Press, 33–56.
- [50] Walter G Stephan and Krystina Finlay. 1999. The role of empathy in improving intergroup relations. *Journal of Social issues* 55, 4 (1999), 729–743.
- [51] Walter G Stephan and Cookie White Stephan. 1985. Intergroup anxiety. *Journal of social issues* 41, 3 (1985), 157–175.
- [52] Karolina Sylwester and Matthew Purver. 2015. Twitter language use reflects psychological differences between democrats and republicans. *PLoS one* 10, 9 (2015), e0137422.
- [53] Fiona A White and Hisham M Abu-Rayya. 2012. A dual identity-electronic contact (DIEC) experiment promoting short-and long-term intergroup harmony. *Journal of Experimental Social Psychology* 48, 3 (2012), 597–608.