

CarSenToGram: Geovisual text analytics for exploring spatio-temporal variation in public discourse on Twitter

Caglar Koylu^{a*}, Ryan Larson^b, Bryce J. Dietrich^c & Kang-Pyo Lee^d

^aGeographical and Sustainability Sciences, University of Iowa, Iowa City, IA, USA;

^bComputer Science, University of Iowa, Iowa City, IA, USA;

^cPolitical Science, University of Iowa, Iowa City, IA, USA;

^dIowa Informatics Initiative, University of Iowa, Iowa City, IA, USA;

*Corresponding author: Caglar Koylu, Assistant Professor,

Department of Geographical and Sustainability Sciences,

University of Iowa, Iowa City, IA, 52242,

caglar-koylu@uiowa.edu

CarSenToGram: Geovisual text analytics for exploring spatio-temporal variation in public discourse on Twitter

Assessing the impact of events on the evolution of online public discourse is challenging due to the lack of data prior to the event and appropriate methodologies for capturing the progression of tenor of public discourse, both in terms of their tone and topic. In this article, we introduce a geovisual analytics framework, CarSenToGram, which integrates topic modeling and sentiment analysis with cartograms to identify the changing dynamics of public discourse on a particular topic across space and time. The main novelty of CarSenToGram is coupling comprehensible spatio-temporal overviews of the overall distribution, topical and sentiment patterns with increasing levels of information supported by zoom and filter, and details-on-demand interactions. To demonstrate the utility of CarSenToGram, in this article, we analyze tweets related to immigration the month before and after the January 27, 2017 travel ban in order to reveal insights into one of the defining moments of President Trump's first year in office. Not only do we find that the travel ban influenced online public discourse and sentiment on immigration, but it also highlighted important partisan divisions within the United States.

Keywords: Geovisual analytics, topic modeling, sentiment analysis, spatio-temporal cartogram, Twitter, public discourse

1. Introduction

Diffusion of policies are often influenced by citizens who express their opinion through public discourse on Twitter. Events significantly impact public discourse which also tend to vary substantially across different political geographies. For example, on January 27, 2017, President Donald Trump suspended the entry of people into the U.S. from seven predominantly Muslim countries. In response, thousands of people flooded airports across the country to protest what the travel ban meant for democracy in the United States and elsewhere, and similar protests were observed on Twitter. Social media data provide a unique opportunity to study geographic variation and evolution of content and sentiment of these publicly shared opinions. However, it is difficult to

assess the impact of events on the evolution of public discourse due to the lack of data prior to the event and appropriate methodologies for capturing the progression of tenor of public discourse, both in terms of their tone and topic. While visual analysis of sentiment and topical themes have become an important area of research in visual analytics, most existing studies have focused on either visualizing sentiment (Zimmerman et al., 2015) or topical themes and their evolution (Cui et al., 2011), diffusion (Wu et al., 2014), and their spatio-temporal patterns (Chae et al., 2012; Koylu, 2018a).

In this article, we introduce a geovisual text analytics framework, CarSenToGram, which integrates topic modeling and sentiment analysis with cartograms to identify the variation in public discourse in terms of the intensity, topical themes and sentiment across space and time. Specifically, we designed CarSenToGram to answer the following questions: (1) What are the major themes of public discourse and sentiment towards a particular topic? (2) How do the intensity, topics and sentiment of public discourse vary across space and time? (3) Which locations (states) have similar public discourse and are more representative of the overall public discourse across all locations (nation)? In order to answer these questions, we first introduce an analytical pipeline for data cleaning and processing that include classification, spatio-temporal aggregation and clustering of topics and sentiment. After the initial processing of data through the analytical pipeline, we introduce CarSenToGram, a geovisual text analytics framework for exploring the intensity, topical and sentiment patterns across space and time. The main novelty of our work is coupling comprehensible spatio-temporal overviews of the overall distribution, topical themes and sentiment patterns with increasing levels of information supported by zoom and filter, and details-on-demand interactions. Thus, while providing spatio-temporal overviews of topics and

sentiment distributions, CarSenToGram also allows the user to view original tweets with respect to their content, topical and sentiment classifications.

In this article, we use CarSenToGram to identify spatio-temporal patterns of public discourse in a keyword-based collection of immigration tweets four weeks before and after President Trump's first travel ban, Executive Order 13769. By analyzing the impact of the travel ban on online public discourse, we not only provide an important foundation for those interested in understanding the dynamics of public discourse on Twitter, but we also help others gain a greater understanding of a very important and recent moment in American political history. Although we use CarSenToGram to explore the online ramifications of the first travel ban, we explain in the conclusion how it can be used more broadly.

2. Related Work

Public Discourse on Twitter

Scholars have used Twitter to study a variety of geo-social phenomena including how tweets can shape both off- and online public discourse (for review, see Zimmer & Proferes, 2014). For example, Shen and Kuo (2014) recently used Twitter to understand how information dissemination can influence online social structures. Similarly, Conover et al. (2011) and Smith et al. (2014) found Twitter exhibited a highly partisan structure with people retweeting and sharing resources to point out their different views. This finding is also consistent with Feller et al. (2011) who found that most topics are discussed by users with competing political preferences.

However, these broader political disagreements are grounded in demographic, social and geographic biases that also permeate the Twitterverse (Pavalanathan & Eisenstein, 2015; Tufekci, 2014). More specifically, Pavalanathan and Eisenstein (2015)

demonstrated that demographic variables such as age and gender interact with geography to create regional linguistic variation. Similarly, Malik et al. (2015) found not only are the users who geotag tweets not representative of the population of the United States, but they also tend to be younger, wealthier, and reside in urbanized areas.

Topic Modeling

Extraction of topical themes of public discourse from Twitter data has increasingly been popular in a variety of domains and applications such as public health, politics, elections, climate change and immigration. Latent Dirichlet Allocation (LDA) have commonly been used to extract topical themes from tweets and geographic patterns of those themes (Ghosh & Guha, 2013; Longley, Adnan, & Lansley, 2015; Pozdnoukhov & Kaiser, 2011). Using a weighting factor based on the term frequency-inverse document frequency (tf-idf) to determine the relative importance of each word (Salton & McGill, 1983), LDA first attributes each word to a set of topics depending on how frequently that word appears within each topic, then classifies each document as a mixture of topics with differing probabilities.

Variation in document sizes has been found to influence the robustness of a topic model, and specifically topic modeling on documents with small number of texts, i.e., short-text, produces unstable document-topic and word-topic probabilities (Hong & Davison, 2010; Yan et al., 2013). Aggregation of tweets into document bins based on keyword and topic similarity (Grant et al., 2011), time (S. Malik et al., 2013), space (Gerber, 2014), user-to-user mentions (Alvarez-Melis & Saveski, 2016; Koylu, 2018a) and location-to-location mentions (Koylu, 2016) have been used to alleviate the short-text problem. However, aggregation of tweets often produces large documents which may also result in unstable classifications. This is because the heterogeneity of content increases proportionally to the size of a document, and documents with very large size

produce document-topic relationships in which a document may belong to large number of topics, and topics become semantically uncertain. Therefore, the documents in a topic model must be small enough so that topic probabilities could vary significantly between documents. Several recent studies have used individual tweets without aggregation to address this problem (Chae et al., 2012; Lansley & Longley, 2016). In this study, we also use tweets as documents to train a series of LDA models.

The studies that focus on public discourse on Twitter are diverse and point out significant variation in terms of the context, topical themes, temporal and spatial patterns. Choi and Park (2013) analyzed co-occurring words to capture discourses relevant to a collective identity during a social protest. Romero et al (2011) analyzed keywords and hashtags in Twitter networks and revealed various topic-specific diffusion patterns. Abel et al. (2011) also showed that characteristics of topics and user profiles significantly vary between different time periods. Kurashima et al. (2013) introduced a topic modeling workflow to estimate a user's interests. Similarly, Steiger et al. (2016) characterized urban activity spaces by analyzing the semantic, temporal and spatial patterns of latent topics derived from LDA. Lansley and Longley (2016) also used LDA to identify the characteristics of places using tweet messages and the time of the day. Finally, Koylu (2018b) introduced an LDA-based framework to extract topics from reciprocal mention tweets and identify how the topics of interpersonal communication vary across space and time. Collectively, these studies not only show LDA can be used to reasonably assess the topical themes of tweets, but it is also quickly becoming the preferred approach.

Sentiment Analysis

Sentiment analysis, or opinion mining, is an area of computational study concerned with identifying people's opinions, emotions, or moods expressed in text. A rapidly growing

interest in sentiment analysis has been observed over the past decade, as it is proven to be practically useful for gaining insight into people's opinions toward individuals, events, topics or issues, especially from a large amount of text data. With that said, sentiment analysis is not without its own limitations, such as the use of sarcasm and instances in which the sentiment is implied, but not explicitly stated. Liu (2012) provides an overview of the key technical issues related to sentiment analysis and various techniques that have been developed to overcome these and other related problems.

The output of sentiment analysis can be positive, negative, or neutral, which can be considered as a classification problem with three classes or be expressed with different intensity levels such as five-star ratings or numbers between -1 and 1, where -1 indicates extremely negative, 1 extremely positive, and 0 neutral. Another two aspects of sentiment analysis are subjectivity and emotion of text (Liu, 2012). Subjectivity score typically takes a value between 0 and 1, where 0 indicates extremely objective (presenting factual information) and 1 extremely subjective (expressing personal feelings or opinions). As with the case of sentiment polarity score described above, subjectivity score can be acquired by binary classifiers based on widely used machine learning approaches. Emotion classification can also be performed by multi-class classifiers. In addition to the dimensions of polarity and subjectivity, other measures of sentiment such as the level of arousal and core emotions such as love, joy, surprise, anger, sadness and fear have successfully been used to classify emotions in textual data (Zimmerman et al., 2015).

There are different levels of sentiment analysis: document level (Bo et al., 2002), sentence level (Wiebe et al., 1999) and entity/aspect level (Hu & Liu, 2004). Depending on the availability of training or benchmark data, approaches to document-

level sentiment classification can be grouped into two broad categories: supervised and unsupervised learning. In this research, we employ document-level sentiment classification which has been widely used in a variety of applications in public health (Tumasjan et al., 2010) and political science (X. Cao et al., 2018). More specifically, we summarize the extreme positive and negative sentiments of tweets before and after an event (January 27, 2017 travel ban) by space (states) and time (week) to identify people's general response to an event (President Trump's executive order). In this way, we summarize individual sentiment classifications by spatio-temporal aggregation and visualization that have been used extensively in cartography and geographic information science (GIScience) but have yet to be employed to understand the changing dynamics of online protests and public discourse during the period before and after the travel ban was announced on January 27, 2017.

Geovisual Analytics

Geovisual analytics integrates cartography and GIScience with analytical and computational methods to derive insight about location-based data (MacEachren, 2013). Due to the growing availability of geo-tagged and user-generated data, researchers have increasingly studied the spatial and temporal evolution of user-generated content on Twitter and other social media sites. However, as this content has grown more and more dynamic, new visual analytical methods are needed to help researchers understand the temporal, spatial, and semantic components of this newly available data (Dork et al., 2008; MacEachren et al., 2011; White & Roth, 2010).

For example, Cao et al (2012) developed a visual analytical tool to detect the temporal trend, social-spatial extent and community response of a topic of interest by analyzing re-tweets. Cao et al (2012) allow the user to pick a keyword or a set of keywords and visualize temporal, spatial and network aspects of re-tweets of selected

keyword(s) using a visual analytics environment. Nelson et al. (2015) developed a similar geovisual analytics technique to link interactive maps with a term polarity view to identify and compare the duality of Republican and Democratic discourses on Twitter by states and topics. Xu et al. (2013) also designed a comparative visual analysis framework to capture competition among topics in social media, and the influences of opinion leaders in formation, convergence and divergence of topics. Koylu (2018a) introduced a geovisual analytics environment consisting of a map display linked with topic-word clouds, a temporal bar chart with topic probabilities, and a time slider to allow interactive exploration of the spatio-temporal patterns of topics during the period of 2016 primary and presidential elections. A commonality between these methods is that they support overview and details-on-demand through interactive filtering of a set of keywords or topics in user-selected temporal and spatial dimensions.

Geovisual analytics has also been used to explore patterns of sentiment. For example, Hoque and Carenini (2016) integrated topic modeling and sentiment analysis with interactive visualization techniques to better understand the evolution of asynchronous conversations in online blogs. After evaluating the utility of their tool, Hoque and Carenini (2016) found that human-computer interactions through an interactive visual interface increased user comprehension of topic model output. Scharl et al. (2013) designed a similar visual analytics dashboard which linked maps with sentiment and text visualizations to help decision makers evaluate the impact of education and public outreach campaigns on environmental literacy. In this article, we combine topic modeling and sentiment analysis with spatio-temporal cartograms in a geovisual analytics framework to allow interactive exploration of changing dynamics in public discourse.

3. Analytical Framework

Figure 1 illustrates the analytical pipeline used to create the geovisual analytics framework, CarSenToGram. Using keywords related to a topic, we first collect tweets that span a period centered on an event. To demonstrate the utility of the proposed framework, we collected immigration tweets from four weeks before and after the first travel ban announced on January 27, 2017. Second, we apply data cleaning and processing steps to filter out irrelevant tweets, and classify each tweet with sentiment and topics. Specifically, we partition tweets into eight-week periods, and derive latent topics by implementing a separate LDA for tweets of each weekly time period. We then cluster the topics across time periods to identify coherent topics across the whole time period. Third, we perform spatio-temporal aggregation (by states and weekly periods) to transform the topic and sentiment outputs into local distribution, topic and sentiment time-series. Each local time-series contain information on the number of tweets and tweets categorized by topics and sentiment for a given temporal (week) and spatial (state) unit. Using the local time-series, we generate rectangular tiles for three cartograms: Distribution View, Topic View and Sentiment View that collectively form the proposed visual analytics framework, CarSenToGram which is named by the combination of the words Cartogram, Sentiment and Topic. With its three views, CarSenToGram provides (1) an overview of spatio-temporal patterns of tweets, topics and sentiment and (2) user interactions to zoom and filter, and perform details-on-demand tasks to view individual tweets with their original content, classified topic and sentiment. Details of the data collection, cleaning, processing as well as the design rationale for CarSenToGram are described in detail in the following subsections.

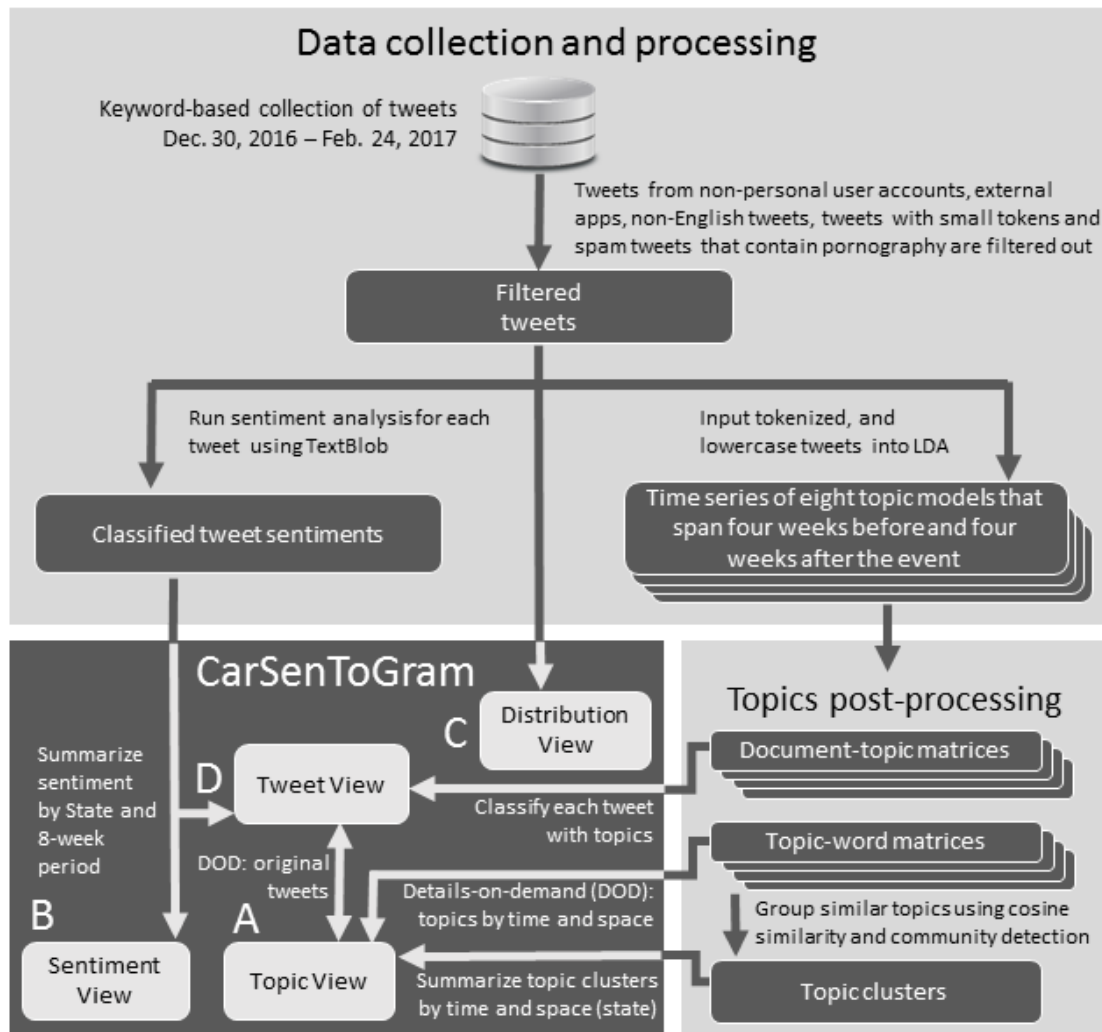


Figure 1 The analytical pipeline for data collection, processing, topic modeling and sentiment analysis for the creation and design of CarSenToGram

3.1 Data Collection and Cleaning

Using the Twitter Streaming API, we collected tweets that contain the keywords related to immigration, and specifically Muslim refugees and immigrants (i.e., “immigration”, “immigrant”, “muslim”, “Islam”, “refugee”) beginning on December 30, 2016 (four weeks before the first travel ban) and ending on February 24, 2017 (four weeks after the first travel ban). Even though we understand that this likely biases our results to those users who are politically interested in the causes and consequences of the travel ban, this is precisely the population we want to study in this article. Because the amount of

tweets generated on the day of the travel ban exceeded the total volume of tweets within the first week of the ban (which was also the maximum volume of tweets across the eight-week period), this resulted in a significant source of bias for temporal analysis of these tweets. Thus, we excluded the tweets that were generated during the day of the travel ban which spanned January the 27th at 4am EST until January the 28th at 4am EST.

There was a total of 72,397,072 “immigration” tweets generated by 6,350,263 users world-wide during the eight-week period. Even though these tweets were written in many languages including Arabic, French, Spanish, Turkish and Persian, we only considered English tweets since our focus was on the reaction to the travel ban in the United States. Ultimately, this left us with 17,501,502 tweets, of which 40 percent were re-tweets. We excluded re-tweets from our topic and sentiment analysis to ensure each tweet was weighted equally. Using the metadata of each tweet, we were able to identify 99% of the tweet locations at the state level, and 27.85% (~3.8 million tweets) of these geo-located immigration tweets were generated in the United States by 759,171 users. Despite Twitter’s language classification we discovered that a significant portion of tweets were non-English or included only emoticons and symbols. We employed an open source language classification library (Lui & Baldwin, 2011) to classify each tweet’s language and further refine our dataset to exclude all non-English tweets. We then filtered out tokens (e.g., linguistic units that represent terms, symbols or words) that are less than 3 characters and tweets with less than three tokens. Finally, we filtered out spam tweets with keywords that are related with nudity or pornography.

After the data cleaning process, we were left with a canonical dataset of 2,088,191 original tweets which were verified to be English, located within the United States and published between December 30, 2016 and February 24, 2017. Ultimately,

this is the largest collection of tweets specifically designed to capture the online public discourse about immigration surrounding the first travel ban. Given the historic and political importance of this event, the data is a contribution in and of itself, but we further use these tweets to introduce our visual analytics framework which helps us understand how major events influence public discourse and sentiment. We use CarSenToGram to explore this important political event, but we explain in the conclusion how our software can be used more broadly.

Figure 2 illustrates the distribution of immigration tweets over time. The large increase corresponding to the day the first travel ban was announced is the evidence of a large public discussion taking place on Twitter. Not only is this a noticeable change in online activity, but the travel ban itself was relatively unexpected. Although there were some grumblings in the White House the day before, the Department of Homeland Security was "caught by surprise" when the first travel ban was announced (Sands, 2018). We expect those on Twitter were equally caught off guard making the ban announcement approximate random assignment. This is important because it gives us some causal leverage, even though we fully acknowledge the inherent limitations with natural experiments such as the strong exchangeability assumption.




Figure 2 Daily frequency of immigration tweets for the eight-week period before and after the first travel ban. The vertical rectangle indicates the day of the travel ban which we omit in our analysis.

3.2 Topic Modeling

Even though the first travel ban produced online discussion, we still know very little about the nature of these conversations. To gain some traction on the latter, we estimated eight LDA topic model series for tweets four weeks before and four weeks after the travel ban. To determine the optimal number of topics and minimize topic overlap, we estimated models with 10, 20, 30, and 40 topics using 2,000 expectation-maximization steps for each of the eight models.

We compared the performance of topic model series with different number of topics in a consecutive manner. For example, we compared the 40-topic model of each week with the 30-topic model of the same week, and the 30-topic model of each week with the 20-topic model of the same week, and so on. In this comparison, we considered

topics as similar if they have cosine similarity equal to or greater than 70%. Cosine similarity is calculated by measuring the cosine of the angle between the two (non-zero) word vectors that form each topic and their term (word) frequencies (Huang, 2008). Cosine similarity between two topics ranges between 0 and 1 with higher values implying greater similarity. We ultimately found that the 20-topic model produced the less topic overlap than the 10-topic model, but had as many distinct topics as the 30-topic model. Therefore, we selected the 20-topic model from each of the weekly collection of tweets, which generated a total of 160 topics (i.e., 20 topics X 8 weeks).

To make better sense of these topics, we created a network of topic similarity both within and between all weekly models from all periods using cosine similarity. We further employed a topic clustering technique utilized by Koylu (2018a) in order to group coherent topics across time periods. Within the topic similarity network, a node represents a topic in a weekly model and an edge represents a binary undirected link that indicates the degree to which two topics are similar. To determine the similarity threshold, we experimented on 50, 60, 70, and 80% cosine similarity. The threshold of 70% produced an optimal number of 6 communities with a modularity score of 0.47. Therefore, we selected 70% as the threshold to restrict the number of edges, meaning only edges with a cosine similarity of at least 70% were included in the final network. This resulted in a topic similarity network of 160 topics and 1,417 edges.

Using this network and a modularity-based community detection algorithm (Clauset, Newman, & Moore, 2004) we derived six topic clusters which are listed in Table 1. Even though these topics capture the general tenor of the online discussion surrounding the first travel ban, 22 topics were classified as outliers because they did not have a cosine similarity over 70% with any other topic. We grouped these topics into a cluster of outliers which we named as “other”. We inferred the labels of the topic

clusters using the combination of the words that appear the most and reading sample tweets from each topic category. Ultimately, we used the clusters outlined in Table 1 to create the Topic View of CarSenToGram.

Table 1 Topic clusters of immigration tweets. The top twenty words with the highest term frequencies in the eight-week period were selected to illustrate each topic cluster.

Topic	Words
Muslim immigrants	Muslim, Trump, countries, immigration, Islam, refugee, Obama, immigrant, white, country, muslims, refugees, terrorist, women, America, Trump's, brotherhood, woman, banned, attack
Islam and Religion	Islam, Muslim, religion, muslims, radical, peace, women, sharia, Islamic, law, America, ISIS, Christian, Christians, country, Christianity, terrorism, kill, Trump, Allah
Trump's order	Immigration, Trump, Trump's, muslim, order, refugee, immigrant, illegal, executive, Donald, president, legal, judge, court, policy, protest, federal, immigrants, law, wall
Refugee crises	Refugee, Muslim, refugees, crisis, Syrian, children, immigration, Trump, resettlement, Canada, camp, program, U.S, child, border, scheme, Islam, America, dubs
Immigration policy	Immigration, Obama, Muslim, policy, Cubans, Cuban, foot, Americans, wet, discrimination, dry, reject, refugee, special, brotherhood, making, administration, ends, Trump, patriotic
Protests	Muslim, immigration, rally, immigrant, protest, rights, Times Square, Trump's, thousands, support, NYC, Trump, Americans, Washington, park, gather, protestors, work, march, video

3.3 Sentiment Analysis

In order to extract the sentiment from immigration tweets, we used TextBlob ("TextBlob: Simplified Text Processing," 2018) which is a commonly-used wrapper for the natural language toolkit or NLTK ("Natural Language Toolkit - NLTK 3.3 documentation," 2018). TextBlob provides a convenient interface for sentiment analysis, in which users can create a TextBlob object with a sentence of interest, then the TextBlob allows internal analysis of the sentence, and the created object contains different kinds of useful information such as sentiment, part-of-speech tagging, tokenization, etc. The output of the sentiment analysis consists of two types of values as sentiment, one for polarity score between -1 (extremely negative) and 1 (extremely

positive) and the other for subjectivity score between 0 (extremely objective) and 1 (extremely subjective). We classified each tweet with a polarity score while we disregarded the subjectivity score. We first used the polarity values of each tweet both in Tweet View, which illustrates the original tweets with their topic and sentiment classification. Secondly, we summarize the distribution of negative and positive sentiments per space (i.e., state) and time (i.e., week) partitions in Sentiment View.

4. CarSenToGram Design and Pattern Exploration

We designed CarSenToGram as a web-based application which can be accessed using the following link: <https://geo-social.com/carsentogram>. We implemented CarSenToGram using JavaScript, HTML, and CSS and with open-source JavaScript libraries including React.js and D3.js. The rationale for our design is based on the general research questions which we translate from the domain specific language such as states, weeks, immigration and travel ban to a generic vocabulary of visualization tasks. Recall, our research questions are as follows:

1. What are the major themes of public discourse and sentiment towards a particular topic (immigration)?
2. How do the intensity, topics and sentiment of public discourse vary over space (states) and time (four weeks before and after the travel ban)?
3. Which locations (states) have similar public discourse and are more representative of the overall public discourse across all locations (nation)? How different local time-series from the global time-series?

In order to answer these questions, we designed CarSenToGram with three overview panels: Distribution View, Topic View and Sentiment View. Each of the three views consists of a spatio-temporal cartogram that allows comparison of patterns across locations (states) and time periods (weeks). In addition to the rectangle tiles that allow

comparison of the patterns per state and week, each view includes a global time-series tile (i.e., all states combined) that allows the user to compare the overall distribution of tweets, their topics and sentiment with the corresponding distribution patterns of each state. Thus, the global time-series tile allows us to answer the third research question specifically. In addition to these overviews, the user can select a location (state) and open Tweet View to visualize the original tweet content with the classified topic and sentiment.

The top panel of the interface (<https://geo-social.com/carsentogram>) includes controls to switch (1) between the three views of Topic, Sentiment and Distribution, (2) absolute and relative scaling for the time-series distributions, and (3) tile sizing factor to determine the size of each rectangle for the cartograms. Alongside with the tile sizing, absolute scaling allows comparison across states and weeks. On the other hand, relative scaling allows comparison of the percentage of different topics or extreme sentiments across states and the eight-week period. While it is useful to compare the relative proportion of topics or extreme sentiment, relative scaling does not take into account the frequency of tweets, and therefore, does not allow comparisons of the relative involvement of states over the eight-week period. Finally, tile sizing parameter is used to alter the size of each rectangle using either the total count of tweets or the population of each state. By default, we use absolute scaling, and number of tweets for tile sizing. In order to prevent the creation of very small rectangle tiles for states that have low tweet or population counts (e.g., Wyoming, or North Dakota), we use min-max scaling and control the minimum size a rectangle (state). The values in between the maximum and minimum are scaled to fit the distribution.

Here, we first describe the common visual variables and parameters used consistently by the three views of CarSenToGram: Distribution View (Figure 3), Topic

View (Figure 4) and Sentiment View (Figure 6). Each state is illustrated by a rectangle tile whose size is set by the total number of tweets for Distribution and Topic Views. Each rectangle is divided into eight temporal bar charts, each of which illustrates the frequency of immigration tweets per week. Dashed line in the middle is used to reference the time of the event, the travel ban. Because absolute scaling is selected as default in all of the views, bar heights (time-series) for each state are scaled based on the maximum value in a local time-series (i.e., the highest aggregate count of tweets among the eight-week period for a given state). While this does not allow direct comparisons of bar heights (time-series) across states, the area of each bar is comparable since we control the size of each rectangle (state) using the count of tweets.

Figure 3 illustrates the Distribution View, a spatio-temporal cartogram that depicts the rate of immigration tweets within the eight-week period. We use two color hues: orange and blue bars to represent the frequency of tweets before and after the first travel ban respectively. In all but a handful of states, orange bars of varying sizes are followed by noticeably larger blue bars. This suggests the dramatic increase in immigration tweets immediately after the travel ban was announced, this notable increase was not isolated to states in which individuals were immediately affected by the ban. For example, increases in states like Massachusetts and New York where large airport protests occurred seem to mirror similar increases in states like Colorado and Missouri where airport protests did not take place. Also, the comparison of the global time-series distribution (U.S. tile) with the local time-series of states (state tiles) reveals that the states of Washington, Texas and Florida exhibit a similar pattern to the overall distribution of tweets that include all states. Undoubtedly, Figure 3 does not give us any information about what is being said, but even a cursory glance would provide strong

evidence that the first travel ban produced a national, rather than regional, discussion of immigration on Twitter.



Figure 3 Distribution View illustrates the relative involvement of each state in producing immigration tweets over the eight-week period before and after the travel ban.

Figure 4 illustrates CarSenToGram's Topic View, which is also a spatio-temporal cartogram designed to provide an overview of the topic content. The same as the Distribution (Figure 3), each state is illustrated by a rectangle tile whose size is determined by the total number of tweets produced by that state. The tile is then divided into eight temporal partitions with stacked bars which illustrate the weekly topic-series that correspond to four weeks before and four weeks after the travel ban. Stacked-bar colors are determined by the topic clusters illustrated in keyword clouds derived during the topic modeling phase (Table 1).

Figure 4 allows us to identify the general topics of public discourse on immigration, and the prevalence of these topics across states and the eight-week period. Here, we find the topic associated with Trump's Executive Order dominating the discussion after the first travel ban was announced. This topic – highlighted in orange in the Topic View – includes several factual words like "executive" and "order" which likely capture the public's general reaction to the travel ban itself. Other words like "legal," "illegal," and "court" likely speak the constitutional questions that were later addressed by the courts. Indeed, while several topics include words like "protest," these legal terms are unique to this topic which make them a defining feature. With that said, the relative distribution of each topic seems to be approximately the same before and after the travel ban. This suggests that the topics discussed in relation to immigration did not change in response to the first travel ban. Rather, the travel ban seemed to magnify the issues that were already being debated on Twitter.

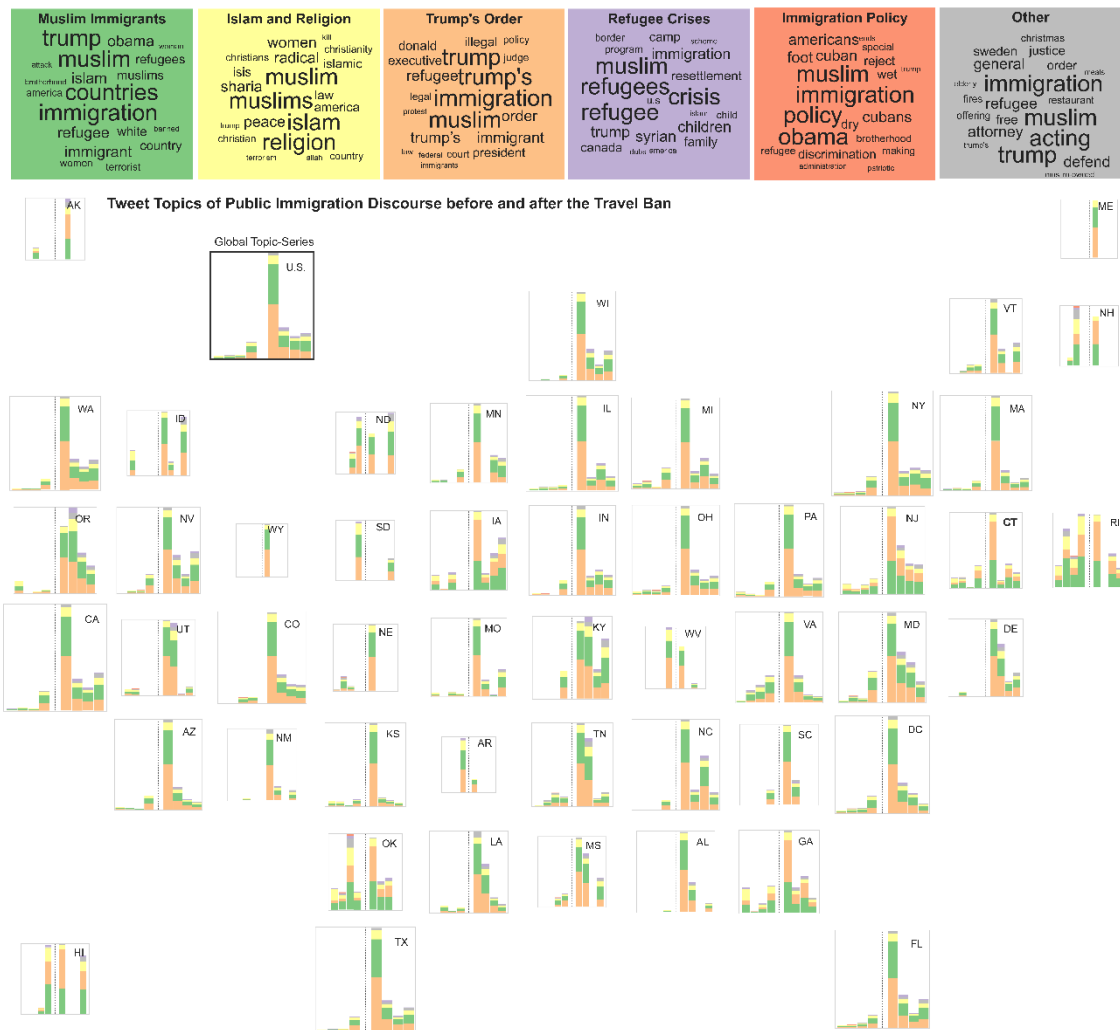


Figure 4 Topic View is dynamically linked with a series of keyword clouds that illustrate the top keywords in topic clusters over the eight-week period before and after the travel ban.

By dynamically linking the cartogram with the keyword cloud and sentiment for individual tweets, Topic View provides increasing levels of information by zoom and filter, and details-on-demand (Shneiderman, 1996). Here, users can click on each state and time partition (e.g., any of the stacked bars in a state) to open the Tweet View to display the original tweets with their respective sentiment and topic classification (Figure 5). Users can also click on a state and time cluster in the topic panel of keyword clouds (Figure 4) to highlight all state and time partitions that belong to that cluster in

the cartogram. Instead of visualizing each topic probability per state and time partitions separately, our topic clustering and visualization provide a fundamental overview of how similar public discourse was between states both before and after the ban was announced. Topic modeling results are sometimes misleading or hard to interpret as the context of the discourse may get lost due to the bag-of-words approach or the short-text problem. However, we address this limitation by allowing the user to gain an understanding of the context by simultaneously viewing the overall distribution of topics across space and time, and the original tweet content alongside with the classified sentiment and topic.

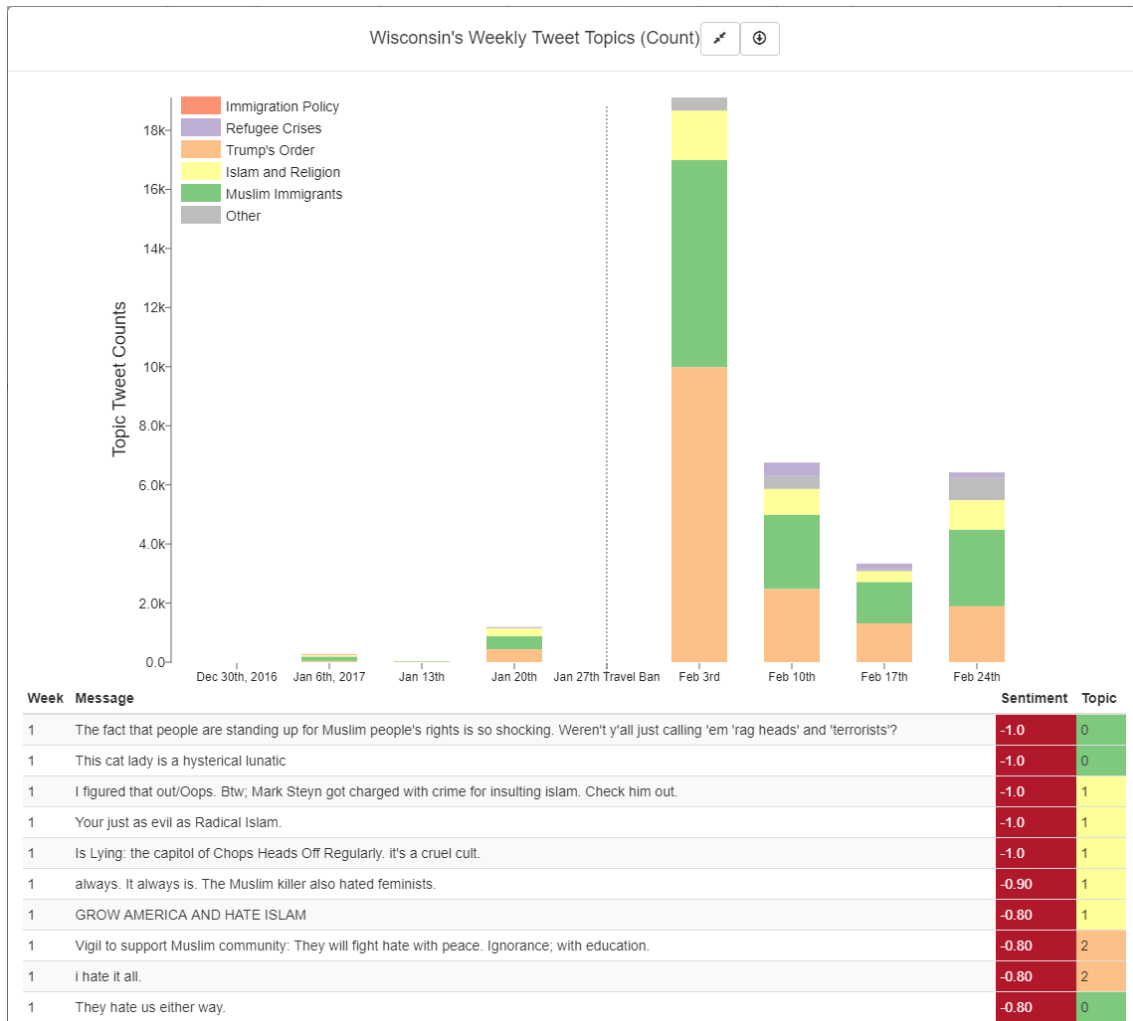


Figure 5 Tweet View in CarSenToGram for details-on-demand. Users can click on each state and weekly period to highlight the top keywords within the eight-week period and display the original tweets with their respective sentiments.

Figure 6 illustrates CarSenToGram's Sentiment View which summarizes the distribution of positive and negative sentiments per state across the eight-week period. Sentiment View allows the user to assess changes in sentiment across states and the eight-week time period. In order to create the sentiment time-series, we first grouped the scores returned by TextBlob into negative (-1 to -0.33), neutral (-0.331 to 0.33), and positive (0.331 to 1) tweets. Because we are primarily interested in the valence expressed towards the first travel ban, we removed the neutral group and only

highlighted the positive and negative sentiment expressed. The area of each rectangle is determined by the total number of positive and negative tweets (excluding the neutral tweets). Each rectangle is divided horizontally into two halves. While the upper half is color-coded by blue to illustrate the count of positive tweets, the lower half is color-coded by orange to illustrate the count of negative tweets (given that absolute scaling is used). The height of each bar in a local sentiment-series is then determined by the maximum count of positive or negative tweets within each local series.

Figure 6 allows the user to explore the spatio-temporal distribution of extreme sentiment across states and the eight-week period. From this figure we observe a significant disagreement over the travel ban regardless of any state. Indeed, we observe both positive and negative sentiment expressed towards immigration before and after the first travel ban even in states which consistently vote Democratic (e.g., California and Massachusetts). The same can be said for predominantly Republican states. For example, even in Alabama where President Trump won 62.7% of the popular vote, an approximately equal number of immigration tweets expressed positive and negative sentiment. Ultimately, this suggests the discussion of immigration is just as polarizing as what previous scholars have found in other realms of the Twitterverse (Conover et al., 2011). Moreover, the first travel ban did not seem to change the tone of debate. Before and after the ban was announced, the distribution of positive and negative tweets remained constant, suggesting the travel ban may have only intensified a discussion that was already taking place on Twitter. Finally, the comparison of the global sentiment-series to local sentiment-series of each state revealed that the states of New York and Washington were most representative of all states for public sentiment on immigration.

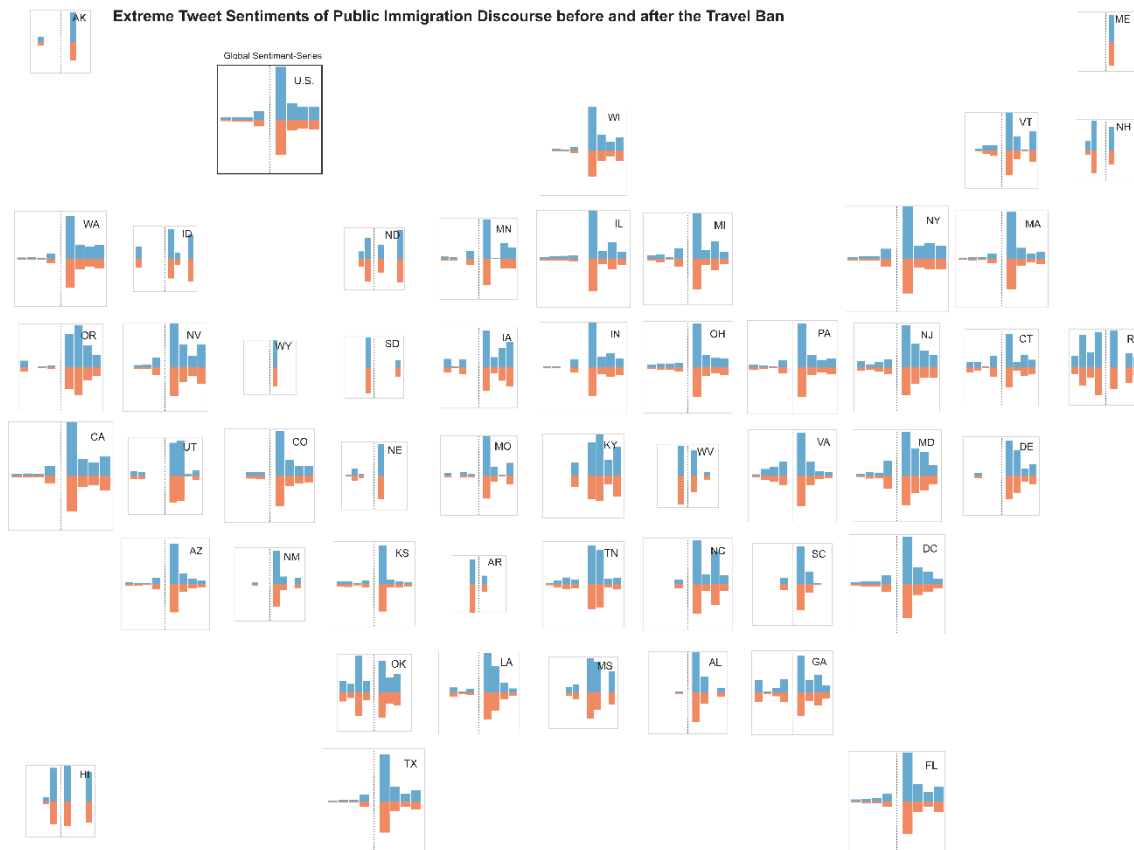


Figure 6 Sentiment View illustrates the number of positive (blue) and negative (orange) immigration tweets. Each state is represented by a rectangular tile whose size is determined by the total number of extreme tweets (positive and negative).

5. Discussion and Conclusion

In this article, we introduced a geovisual text analytics framework named CarSenToGram, which allows exploration of spatio-temporal distribution, topical and sentiment patterns of public discourse on a particular topic before and after a major event. While our design is generic and allows capturing the variation in public discourse and sentiment between locations and time periods, we demonstrated the utility of CarSenToGram by analyzing the changing dynamics of the online public discourse on immigration before and after the January 27, 2017 travel ban. We not only provide an important foundation for those interested in understanding spatio-temporal changes on

Twitter, but we also help others gain a greater understanding of a very important (and recent) moment in American political history.

CarSenToGram provides an overview first, and then zoom and filter, and details-on-demand interactions to link the holistic patterns of topical themes and sentiment across space and time with the details of original tweet content. Using the principles of increasing level of information, CarSenToGram contributes to visual analytics with a visual text analysis framework for exploring the changing dynamics of public discourse across space and time.

Although understanding the consequences of the first travel ban is undoubtedly important, we use this application to demonstrate the broader utility of CarSenToGram. Researchers can use CarSenToGram to examine a corpus of tweets about any phenomenon such as the reputation of a person, product or company. For example, one can analyze the progression of public discourse before and after any major event such as the 2017 Delta Airlines computer outage or the Volkswagen emissions scandal in 2015. Another example application of this tool would be using news articles published about a certain topic such as immigration. Not only can one identify the topical themes and sentiment using a corpus of news articles, but the three main components of our software -- Distribution View, Topic View and Sentiment View -- give researchers the ability to examine spatio-temporal dynamics before and after important events, potentially providing important causal insights.

With that said, we are still actively improving our software. First, CarSenToGram allows comparison of states both in terms of topics and sentiments separately. Topic View allows a detail-on-demand functionality which displays individual tweets with their sentiment. However, in order to provide an overview of the relationship between topical and sentiment similarity, we plan to dynamically link the

cartograms of sentiment with topics. This would allow us to explore sentiments of topics and the changing dynamics over time and space. Second, we have analyzed the change of topics and sentiment using a time granularity of a week. However, one can look into change by finer temporal granularity such as by daily or even hourly. Such analysis could be employed with online topic modeling to extract topics and sentiments in real-time to gauge the reactions to certain events using CarSenToGram. Third, we plan to conduct an evaluation to assess the utility and usability of CarSenToGram in capturing insight on semantic, sentiment and spatio-temporal patterns. Finally, we plan to more fully understand how the proximity to actual protest events influences online discussions, like those found on Twitter immediately after the travel ban.

Undoubtedly, there is considerably more work to be done, but CarSenToGram and the broader analytical approach we have outlined in this paper will help future scholars (including ourselves) better understand public discussions of major events, like the first travel ban. In doing so, we move both the methodological and theoretical literature in an important new direction. Reactions to major political events occur both on- and offline. They also occur across space and time. However, the way we currently visualize data does not fully capture these interwoven dynamics. By giving researchers the ability to work with data in real-time, CarSenToGram could help move geovisual analytics towards real-time insight derivation. Even though we use our software to understand the public reaction to the first travel ban, we look forward to seeing how future scholars use our broader analytic approach to answer their own research questions.

References

Abel, F., Gao, Q., Houben, G.-J., & Tao, K. (2011). Analyzing user modeling on twitter for personalized news recommendations. In *User Modeling, Adaption and Personalization* (pp. 1-12): Springer.

- Alvarez-Melis, D., & Saveski, M. (2016). Topic Modeling in Twitter: Aggregating Tweets by Conversations. *ICWSM, 2016*, 519-522.
- Bo, P., Lee, L., & Vaithyanathan, S. (2002). *Thumbs up?: sentiment classification using machine learning techniques*. Paper presented at the Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP).
- Cao, N., Lin, Y.-R., Sun, X., Lazer, D., Liu, S., & Qu, H. (2012). Whisper: Tracing the spatiotemporal process of information diffusion in real time. *Visualization and Computer Graphics, IEEE Transactions on*, 18(12), 2649-2658.
- Cao, X., MacNaughton, P., Deng, Z., Yin, J., Zhang, X., & Allen, J. G. (2018). Using Twitter to Better Understand the Spatiotemporal Patterns of Public Sentiment: A Case Study in Massachusetts, USA. *International journal of environmental research and public health*, 15(2), 250.
- Chae, J., Thom, D., Bosch, H., Jang, Y., Maciejewski, R., Ebert, D. S., & Ertl, T. (2012). *Spatiotemporal social media analytics for abnormal event detection and examination using seasonal-trend decomposition*. Paper presented at the Visual Analytics Science and Technology (VAST), 2012 IEEE Conference on.
- Choi, S., & Park, H. W. (2013). An exploratory approach to a Twitter-based community centered on a political goal in South Korea: Who organized it, what they shared, and how they acted. *New Media & Society*.
- Clauset, A., Newman, M. E. J., & Moore, C. (2004). Finding community structure in very large networks. *Physical Review E*, 70(6). doi:10.1103/PhysRevE.70.066111
- Conover, M., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F., & Flammini, A. (2011). *Political Polarization on Twitter*. Paper presented at the ICWSM.
- Cui, W., Liu, S., Tan, L., Shi, C., Song, Y., Gao, Z., . . . Tong, X. (2011). Textflow: Towards better understanding of evolving topics in text. *IEEE transactions on visualization and computer graphics*, 17(12), 2412-2421.
- Dork, M., Carpendale, S., Collins, C., & Williamson, C. (2008). Visgets: Coordinated visualizations for web-based information exploration and discovery. *Visualization and Computer Graphics, IEEE Transactions on*, 14(6), 1205-1212.
- Feller, A., Kuhnert, M., Sprenger, T. O., & Welp, I. M. (2011). *Divided They Tweet: The Network Structure of Political Microbloggers and Discussion Topics*. Paper presented at the ICWSM.
- Gerber, M. S. (2014). Predicting crime using Twitter and kernel density estimation. *Decision Support Systems*, 61, 115-125.
- Ghosh, D., & Guha, R. (2013). What are we 'tweeting' about obesity? Mapping tweets with topic modeling and Geographic Information System. *Cartography and Geographic Information Science*, 40(2), 90-102. doi:10.1080/15230406.2013.776210
- Grant, C. E., George, C. P., Jenneisch, C., & Wilson, J. N. (2011). *Online Topic Modeling for Real-Time Twitter Search*. Paper presented at the TREC.
- Hong, L., & Davison, B. D. (2010). *Empirical study of topic modeling in twitter*. Paper presented at the Proceedings of the First Workshop on Social Media Analytics.
- Hoque, E., & Carenini, G. (2016). Interactive Topic Modeling for Exploring Asynchronous Online Conversations: Design and Evaluation of ConVisIT. *Acm Transactions on Interactive Intelligent Systems*, 6(1). doi:10.1145/2854158
- Hu, M., & Liu, B. (2004). *Mining and summarizing customer reviews*. Paper presented at the Proceedings of ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD).

- Huang, A. (2008). *Similarity measures for text document clustering*. Paper presented at the Proceedings of the sixth new zealand computer science research student conference (NZCSRSC2008), Christchurch, New Zealand.
- Koylu, C. (2016). *Extracting and Visualizing Geo-Social Semantics from the User Mention Network on Twitter*. Paper presented at the The Ninth International Conference on Geographic Information Science (GIScience), Montreal, Canada.
- Koylu, C. (2018a). Modeling and visualizing semantic and spatio-temporal evolution of topics in interpersonal communication on Twitter. *International Journal of Geographical Information Science*, 1-28. doi:10.1080/13658816.2018.1458987
- Koylu, C. (2018b). Uncovering Geo-Social Semantics from the Twitter Mention Network: An Integrated Approach Using Spatial Network Smoothing and Topic Modeling. In S.-L. Shaw & D. Sui (Eds.), *Human Dynamics Research in Smart and Connected Communities* (pp. 163-179). Cham: Springer International Publishing.
- Kurashima, T., Iwata, T., Hoshide, T., Takaya, N., & Fujimura, K. (2013). *Geo topic model: joint modeling of user's activity area and interests for location recommendation*. Paper presented at the Proceedings of the sixth ACM international conference on Web search and data mining.
- Lansley, G., & Longley, P. A. (2016). The geography of Twitter topics in London. *Computers, Environment and Urban Systems*, 58, 85-96.
- Liu, B. (2012). Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies*, 5(1), 1-167.
- Longley, P. A., Adnan, M., & Lansley, G. (2015). The geotemporal demographics of Twitter usage. *Environment and Planning A*, 47(2), 465-484.
- Lui, M., & Baldwin, T. (2011). *Cross-domain feature selection for language identification*. Paper presented at the Proceedings of 5th international joint conference on natural language processing.
- MacEachren, A. M. (2013). Cartography as an Academic Field: A Lost Opportunity or a New Beginning? *Cartographic Journal, The*, 50(2), 166-170.
- MacEachren, A. M., Robinson, A. C., Jaiswal, A., Pezanowski, S., Savelyev, A., Blanford, J., & Mitra, P. (2011). *Geo-twitter analytics: Applications in crisis management*. Paper presented at the Proceedings of the 25th international cartographic conference. Paris, France.
- Malik, M. M., Lamba, H., Nakos, C., & Pfeffer, J. (2015). Population bias in geotagged tweets. *People*, 1(3,759.710), 3,759.710-757,233.531.
- Malik, S., Smith, A., Hawes, T., Papadatos, P., Li, J., Dunne, C., & Shneiderman, B. (2013). *TopicFlow: visualizing topic alignment of Twitter data over time*. Paper presented at the Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, Niagara, Ontario, Canada.
- Natural Language Toolkit - NLTK
3.3 documentation. (2018). Retrieved from <http://www.nltk.org>
- Nelson, J. K., Quinn, S., Swedberg, B., Chu, W., & MacEachren, A. M. (2015). Geovisual analytics approach to exploring public political discourse on Twitter. *ISPRS International Journal of Geo-Information*, 4(1), 337-366.
- Pavalanathan, U., & Eisenstein, J. (2015). Confounds and consequences in geotagged twitter data. *arXiv preprint arXiv:1506.02275*.
- Pozdnoukhov, A., & Kaiser, C. (2011). *Space-time dynamics of topics in streaming text*. Paper presented at the Proceedings of the 3rd ACM SIGSPATIAL International Workshop on Location-Based Social Networks.

- Romero, D. M., Meeder, B., & Kleinberg, J. (2011). *Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter*. Paper presented at the Proceedings of the 20th international conference on World wide web.
- Salton, G., & McGill, M. J. (1983). Introduction to modern information retrieval.
- Sands, G. (2018). DHS 'caught by surprise' when Trump issued travel ban 1.0, report finds Retrieved from <https://abcnews.go.com/Politics/dhs-caught-surprise-trump-issued-travel-ban-10/story?id=52474592>
- Scharl, A., Hubmann-Haidvogel, A., Weichselbraun, A., Lang, H.-P., & Sabou, M. (2013). *Media Watch on Climate Change--Visual Analytics for Aggregating and Managing Environmental Knowledge from Online Sources*. Paper presented at the System Sciences (HICSS), 2013 46th Hawaii International Conference on.
- Shen, C.-W., & Kuo, C.-J. (2014, 2014). *Analysis of Social Influence and Information Dissemination in Social Media: The Case of Twitter*.
- Shneiderman, B. (1996). *The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations*. Paper presented at the Proceedings of the 1996 IEEE Symposium on Visual Languages.
- Smith, M. A., Rainie, L., Shneiderman, B., & Himelboim, I. (2014). Mapping Twitter topic networks: From polarized crowds to community clusters. *Pew Research Center, 20*.
- Steiger, E., Resch, B., & Zipf, A. (2016). Exploration of spatiotemporal and semantic clusters of Twitter data using unsupervised neural networks. *International Journal of Geographical Information Science, 30*(9), 1694-1716. doi:10.1080/13658816.2015.1099658
- TextBlob: Simplified Text Processing. (2018). Retrieved from <http://textblob.readthedocs.io/en/dev>
- Tufekci, Z. (2014). Big Questions for Social Media Big Data: Representativeness, Validity and Other Methodological Pitfalls. *ICWSM, 14*, 505-514.
- Tumasjan, A., Sprenger, T. O., Sandner, P. G., & Welp, I. M. (2010). Predicting Elections with Twitter: What 140 Characters Reveal about Political Sentiment. *ICWSM, 10*, 178-185.
- White, J. J., & Roth, R. E. (2010). *TwitterHitter: Geovisual analytics for harvesting insight from volunteered geographic information*. Paper presented at the Proceedings of GIScience.
- Wiebe, J., Bruce, R. F., & O'Hara, T. P. (1999). *Development and use of a gold-standard data set for subjectivity classifications*. Paper presented at the Proceedings of the Association for Computational Linguistics (ACL).
- Wu, Y., Liu, S., Yan, K., Liu, M., & Wu, F. (2014). Opinionflow: Visual analysis of opinion diffusion on social media. *IEEE transactions on visualization and computer graphics, 20*(12), 1763-1772.
- Xu, P., Wu, Y., Wei, E., Peng, T.-Q., Liu, S., Zhu, J. J., & Qu, H. (2013). Visual analysis of topic competition on social media. *IEEE Transactions on Visualization and Computer Graphics, 19*(12), 2012-2021.
- Yan, X., Guo, J., Lan, Y., & Cheng, X. (2013). *A biterm topic model for short texts*. Paper presented at the Proceedings of the 22nd international conference on World Wide Web.
- Zimmer, M., & Proferes, N. J. (2014). A topology of Twitter research: disciplines, methods, and ethics. *Aslib Journal of Information Management, 66*(3), 250-261.

Zimmerman, C., Stein, M.-K., Hardt, D., & Vatraru, R. (2015). Emergence of things felt: Harnessing the semantic space of Facebook feeling tags.