

Analyzing and Visualizing Government-Citizen Interactions on Twitter to Support Public Policy-making

ROCÍO B. HUBERT, ELSA ESTEVEZ, and ANA MAGUITMAN, ICIC CONICET-UNS,
Universidad Nacional del Sur Bahía Blanca, Argentina
TOMASZ JANOWSKI, Gdańsk University of Technology, Poland; Danube University Krems, Austria

Twitter is widely adopted by governments to communicate with citizens. It has become a major source of data for analyzing how governments communicate with citizens and how citizens respond to such communication, uncovering important insights about government-citizen interactions that could be used to support public policy-making. This article presents research that aims at developing a software tool called Twitter Analytics for Government Intelligence and Public Participation (TA4GIP) that applies sentiment analysis and visualization techniques to information collected from Twitter and presents the findings to policy-makers and other non-technical users to facilitate understanding and interpretation. The use of the tool is illustrated through the case study of Twitter communication carried by five government secretaries responsible for health, education, social development, labor, and environment sectors in Mexico, and corresponding citizen responses over a nine-month period. The case study demonstrates that TA4GIP helps identify and analyze relevant aspects of government presence and citizen participation on social media, such as abnormal activity, salient topics being discussed, citizen views about enacted public policies, correlations between types of emotions in responses to particular government announcements, topics that generate polarized reactions from citizens, and many others.

CCS Concepts: • **Information systems** → **Social networking sites**; • **Applied computing** → **E-government**; • **Information systems** → **Data mining**; • **Computing methodologies** → **Visual analytics**;

Additional Key Words and Phrases: Digital government, government 2.0, social media, visual analysis, sentiment analysis

ACM Reference format:

Rocio B. Hubert, Elsa Estevez, Ana Maguitman, and Tomasz Janowski. 2020. Analyzing and Visualizing Government-Citizen Interactions on Twitter to Support Public Policy-making. *Digit. Gov.: Res. Pract.* 1, 2, Article 15 (April 2020), 20 pages. <https://doi.org/10.1145/3360001>

1 INTRODUCTION

Social media, particularly Twitter, has become a tool widely used by governments for sharing news about public affairs, providing information about public services, making policy announcements, and in general interacting with citizens and assessing their views on policy and government performance. Through Twitter, citizens

Authors' addresses: R. B. Hubert, E. Estevez, and A. Maguitman, Departamento de Ciencias e Ingeniería de la Computación, Universidad Nacional del Sur, Instituto de Ciencias e Ingeniería de la Computación (UNS-CONICET), San Andrés 800, Campos Altos de Palihue, 8000 Bahía Blanca, Argentina; emails: {rocio.hubert, ece, agm}@cs.uns.edu.ar; T. Janowski, Department of Informatics in Management, Gdańsk University of Technology, ul. Gabriela Narutowicza 11/12, 80-233 Gdańsk, Poland and Department for E-Governance and Administration, Danube University Krems, Dr. Karl Dorrek Straße 30, 3500 Krems, Austria; email: tomasz.janowski@pg.edu.pl.

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

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2639-0175/2020/04-ART15 \$15.00

<https://doi.org/10.1145/3360001>

can express their likes or dislikes to government decisions and announcements, for example, by liking a tweet, retweeting a post, replying to a post, and so on. As such interactions become more engaging and multimedia-rich, they become an important source of inputs to government decision-making. Advanced text mining and data visualization techniques offer many opportunities for making such inputs understandable and interpretable to policy-makers, domain experts, and other non-technical users.

The literature shows that simply processing information and applying visual analytics to the outcomes can provide valuable insights to decision-making in many domains, such as financial monitoring [1], urban planning [2], customer retention [3], service coordination [4], crime analysis [5], and others. Visual analytics can also support policy processes in terms of modeling [6], analysis [7], formulation [8], decision support [9], and so on. In particular, visual representations can reveal patterns and trends resulting from government announcement on social media, citizen reactions to such announcements, and how government responds to citizen reactions in terms of following up, giving up, or correcting their decisions. There are plenty of opportunities to assess decisions from governments and politicians, and reactions from citizens expressed through social media by means of advanced data processing and visualization tools that involve text mining, sentiment analysis, machine learning, and visual analytics. In spite of that, there is a notable lack of off-the-shelf software tools to process information posted and interactions taking place on Twitter, to uncover interaction patterns, and to visualize such patterns to aid decision-making. For example, understanding citizen comments about the latest minister post would typically deploy ad-hoc procedures narrowly dedicated to fulfilling this goal, run by software experts.

The aim of this article is to present research that led to the development of a prototype software tool called Twitter Analytics for Government Intelligence and Public Participation (TA4GIP) to help non-expert users analyze government posts and citizen opinions on Twitter, to introduce this tool, and to assess it through an in-depth case study. TA4GIP implements a number of visualization techniques that help analyze how governments use Twitter to carry out communication with citizens and how citizens respond to such communication. Since there is a wide range of techniques that can be used for this purpose [10], the first step is to test their fitness to this purpose. This was done in our previous work [11]. In particular, we explored: Word Clouds to highlight the main topics discussed by citizens over a period of time, Euler Diagrams to represent different types of information resources shared between government and citizens, Extended Hasse Diagrams to visualize different types and varieties of government-citizen interactions, Line Charts to depict the level of Twitter activity over time, Radar Charts to represent emotional content present in citizen responses, and Bubble Charts to indicate the type and degree of such emotions.

In this article, we advance the preliminary study [11] by introducing TA4GIP and using the case of Mexico and five sectors, managed by the corresponding government secretaries, in the country—health, social development, education, labor, and environment—applying this tool to assess how visual techniques can help analyze government-citizen interactions on Twitter. In particular, we focus on how citizens react to government communication on this platform. The study covers activities that took place over a period of nine months from October 2017 to June 2018 through the official Twitter accounts of the five secretaries. The practical implication of our study is highlighting opportunities created by visual methods to gain deeper insights into data generated by social media, to support decisions in public communication and policy making.

The rest of this article is organized as follows. Section 2 outlines the research methodology, followed by a review of related work in Section 3. Section 4 introduces the design and functionality of TA4GIP, and Section 5 applies this tool to analyze a specific data-set from Mexico. Finally, Section 6 presents a discussion on the findings and Section 7 offers some conclusions.

2 METHODOLOGY

The guiding question underpinning this research is: How can digital technologies help analyze government-citizen interactions on social media and visualize the findings to inform decision- and policy-making?



Table 1. Twitter Accounts of Government Secretaries Responsible for Five Selected Sectors in Mexico

SECTOR	SECRETARY NAME	TWITTER ACCOUNT
Health	Secretaría de Salud	SSalud_mx
Social Development	Secretaría de Desarrollo Social	SEDESOL_mx
Education	Secretaría de Educación Pública	SEP_mx
Labor	Secretaría del Trabajo y Previsión Social	STPS_mx
Environment	Secretaría de Medio Ambiente y Recursos Naturales	SEMARNAT_mx

To address this question, we conducted a series of tasks: (1) Assessing existing literature regarding tools for analyzing and visualizing interactions on social media. The results are presented in Section 3. (2) Assessing existing visualization techniques for government-citizen interactions on social media. The outcomes were discussed in [11]. (3) Designing and developing the software tool to fill the gaps uncovered in previous steps. The requirements and features of the tool are outlined in Section 4. (4) Collecting data from social media to validate the tool, as explained below. (5) Applying the tool to the data collected in the previous step, presented in Section 5.

To build the dataset for the analysis and for assessing the tool, we selected Mexico as the case study, based on the country's leading position in the region according to the UNDESA's e-Participation Index [12]. We also selected five government sectors: health, education, social development, labor and environment, since these are the main government sectors assessed by the UN e-Government Survey for measuring the Online Service Index (OSI). For each sector, we identified the official name and Twitter account of the responsible secretary. The results are shown in Table 1.

The entire dataset consists of all tweets posted from the official accounts of these secretaries between October 2017 and June 2018, plus citizen comments in response to these tweets. In total, the complete dataset comprises 30,657 tweets, 14,657 (48%) from the secretaries and 16,000 (52%) from citizens.

3 RELATED WORK

To assess the presence of software tools to analyze government-citizen interactions on social media, we conducted a literature review of which results are presented below.

In Reference [13], the authors present a study of Twitter use by municipalities in Saudi Arabia. The study uncovered that the use is mostly serving one-way dissemination of information and providing replies to questions. It concludes that while Saudi Arabia has the highest Twitter use in the world, the opportunity to engage citizens is lost, since municipalities fail to engage them in two-way communication. Likewise, Reference [14] recalls the experience of an e-participation project that applies opinion-mining techniques in a city in Ireland, highlighting the importance of local government, businesses and participating organizations as key stakeholders, and of multichannel communication to ensure constructive, engaged and unbiased deliberation.

In a similar context to our research, Reference [15] discusses a strategy for collecting data from Twitter accounts of the New York and San Francisco city governments and explains a webometric approach to identify e-participation patterns. A related work presents a method that can automatically "identify the public's stance against governmental decisions" [16]. Both papers also show how to automatically summarize citizen opinions and identify collective thinking patterns in e-participation initiatives by using intelligent techniques based on contextualized search, opinion mining and argumentation. Various applications of artificial intelligence and intelligent systems research in e-government and politics 2.0 are also outlined in Reference [17]. In particular, Reference [18] shows how social media analytics tools can be used to collect, monitor, analyze, summarize and visualize social media data. The main difference with the current work is the focus on applying visual analytics techniques.

Other related work focus on the use of social media in the context of e-Participation, using Mexico as a case study. In Reference [19], the authors analyze tweets related to a governor's annual State of the Union speech to



identify and understand the links between governor followers, citizens and media, extracting different clusters and classifying the links into several categories for further analysis. Another work presented in Reference [20] examines the relationship between the voter sentiment expressed through Facebook emoticons and the electoral behavior in the local state election in Mexico by analyzing citizens' reactions to Facebook posts published by the main political parties.

The concept of passive crowdsourcing of e-participation initiatives is discussed in Reference [21]. The concept relies on searching for content and extracting knowledge, ideas and opinions expressed by citizens without any stimulation from government. A related methodology to systematically exploit social media content through the use of text analytics and opinion mining techniques is presented in References [22, 23]. The methodology can be used to aid policy-making. The initiative outlined in Reference [24] explores multiple social media for participative policy-making, enabling a systematic, efficient, centrally managed exploration of multiple social media content by government agencies as well as an integrated management of such media. The framework proposed in Reference [25] helps exploit multiple social media platforms to carry out crowdsourcing for public policy-making, by creating campaigns on multiple social media channels, continuously monitoring citizen comments, and visualizing the analytics and opinion mining results. Opinion mining is used to determine the polarity of citizen responses (positive, negative or neutral), and simulation modeling to estimate the outcomes of citizen proposals and forecast levels of citizens' interest in the policies under discussion. The main difference with the work presented here is the type of visual techniques applied.

Another example of a dedicated e-participation platform is the WeGov toolkit [26], which provides a forum for citizens to interact with policy-makers. Other platforms include UbiPOL, which comprises a workflow for policy-making and business intelligence functionality [27], and the social software infrastructure [28], which integrates citizen-led and government-led deliberations to understand the socio-technical dimensions of e-participation.

In Reference [29], the authors discuss the application of visual analytics to understand information embedded in open government data. Two main differences of this work compared with ours are the data sources—open government data versus Twitter posts, and the type of visual analytics techniques applied.

Finally, sentiment analysis has been used for different purposes, such as predicting social events, increasing citizen participation, and uncovering citizens' political views and priorities. An analysis of blog posts of an e-participation platform shows how sentiments can generate political participation and how emotions increase participation [30]. The frequency of posts made by local governments through their Facebook accounts is combined with sentiment analysis to determine the success of government-citizen communication [31]. Another analysis determines if social events are reflected in the emotions detected on Tweets and whether the detected sentiments can be used to predict riots or manifestations [32]. However, to the best of our knowledge, there are no studies combining visual and sentiment analysis to examine government-citizen interaction on social media as we have done in this article.

4 TA4GIP ARCHITECTURE, FUNCTIONALITY, AND VISUALIZATION

TA4GIP was designed to implement a client-server architecture, with some light-weight functionality processed on the client side and the remaining functionality available on the server that hosts the tool. This section describes the architecture and functionality of the tool (Section 4.1) and supported visualization types (Section 4.2).

4.1 TA4GIP Architecture and Functionality

The main functionality requested at the client side comprises: to follow a Twitter account, to stop following a Twitter account, to select a Twitter account to visualize patterns of interactions, and to select the time frames for collecting the tweets and for analyzing them. Running the tool requires collecting tweets from selected Twitter accounts and over specific periods of time in advance, and uploading such tweets to the internal database used by the tool. Regarding the analysis, once an account and a time frame are selected, all implemented visualizations are shown. This includes the volume of tweets published through the account during the selected period, a histogram



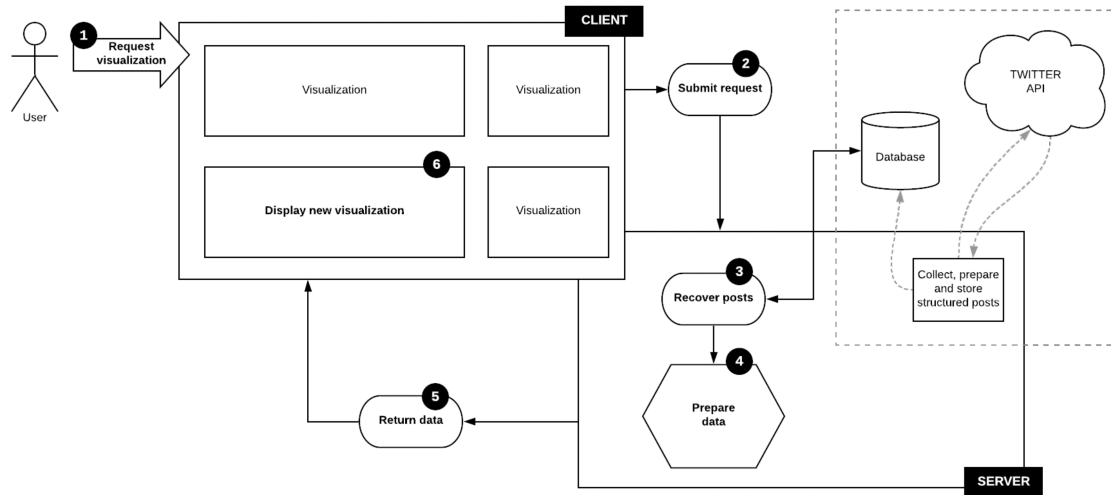


Fig. 1. TA4GIP architecture and work flow.

showing the levels of activity and a line chart depicting types of interactions and emotions. It is possible to zoom into visualizations and see more details for restricted time periods. It is also possible to carry out a deeper analysis of selected interactions, e.g., comments, favorites, and retweets. When selecting a window with a visualization, e.g., a word cloud, a request is sent to the server, which processes the information and updates the selected visualization on the client side.

On the server side, the main functionality comprises: connecting with Twitter and Facebook interfaces through the respective Application Programming Interfaces (APIs) and collecting user posts; structuring the collected data according to a predefined model comprising users, text, words, hashtags, multimedia files, and numbers of retweets, favorites and comments; and saving the posts in the internal TA4GIP database. For each collected tweet, all associated words and emotions are also retrieved and saved in the database, removing the stop words (articles, prepositions, pronouns, and other non-information-bearing words) and accents beforehand. However, misspellings and grammatical errors are neither removed nor corrected. In the current version of the tool, no cleansing of data is performed. Thus, duplicate tweets are only eliminated if they have the same identifiers to existing tweets, and troll-posted tweets are not identified. In addition, the server responds to client requests by recovering the appropriate tweets and processing selected visualizations. For example, if a user requests multimedia data for a Euler diagram, then it recovers all tweets, summarizes the shared content, and returns the data in the format suitable for visualization. TA4GIP is able to continuously listen and collect new tweets published by the accounts followed.

The architecture and workflow implemented by TA4GIP are depicted in Figure 1. In particular, the figure depicts a typical workflow starting with a user request to a client to add a visualization (1), which triggers the corresponding request to the server (2). In turn, the server responds by recovering the posts from the database (3), preparing the corresponding data (4), and returning this data to the client (5). The client finally displays a new visualization (6). The figure also depicts how using the Twitter API, TA4GIP collects, prepares and stores posts in its internal database, making the content available for new visualization requests from its clients.

4.2 TA4GIP-Supported Visualizations

The tool uses seven techniques to visualize collected data: Line Charts, Word Clouds, Euler Diagrams, Extended Hasse Diagrams, Radar Charts, Bubble Charts, and Sankey Diagrams. The techniques are described in the following.



1. Line Charts provide a graphical representation of the observed data in a time sequence, i.e., as a series of data points connected by line segments. Histograms allow to observe the frequency of a variable within a specific range of values. TA4GIP applies histograms and line charts in the same visualization to depict the level of government activity over time, along with the level of citizen responses, by observing their comments, retweets or favorites, as well as emotions expressed in the comments; examples are shown in Figures 2 and 6. In the rest of this article, such diagrams are referred to as Line Charts.
2. Word Clouds represent graphically the frequencies of words in a passage or corpus of text. TA4GIP applies them to identify and highlight key topics; an example is depicted in Figure 3.
3. Euler Diagrams represent sets and their non-empty intersections. TA4GIP applies Euler Diagrams to visualize links, images and videos shared through selected Twitter accounts; an example is shown in Figure 4.
4. Extended Hasse Diagrams visualize relationships among subsets where the size of the nodes is proportional to the size of the sets. TA4GIP applies them to analyze the combinations of different types of citizen contributions, i.e., comments, retweets or favorites; an example is presented in Figure 5.
5. Radar Charts quantify multiple factors associated with an item or situation. TA4GIP applies them to quantify eight factors representing primary emotions [33] on the examined comments, offering a visual approach to sentiment analysis; an example is shown in Figure 7.
6. Bubble Charts displays points within three dimensions— x , y , and the size. TA4GIP applies them to represent emotions identified in citizen comments where the size of a bubble is proportional to the number of comments associated with an emotion; an example is shown in Figure 8.
7. Finally, Sankey Diagrams are flow charts where the width of arrows is proportional to the volume of flows. TA4GIP applies them to represent relations between the resource shared by a government Twitter account and citizen responses; an example is provided in Figure 9.

Given that different visualizations can be generated simultaneously, we can also perform combined data analysis. Some combinations are already provided by the tool, for instance, correlation between citizen responses and activity levels through Line Charts. Other combinations can be obtained by observing different charts. For example, we can uncover relationships between the resources shared through tweets, i.e., images, videos, or links, and the emotions expressed through citizen responses by looking at the corresponding Euler Diagrams and Bubble Charts. Additional visualization techniques can be used to represent the results of combining different types of analysis. For instance, some techniques help determine if different types of activities are related to each other, such as the Sankey Diagrams representing correlations between citizen responses and different types of resources such as links, images or videos shared by government. An example is depicted in Figure 9.

5 TA4GIP USAGE—MEXICO CASE STUDY

This section describes how TA4GIP can be applied to analyze government-citizen interactions triggered by Twitter announcements carried out by five government secretaries responsible for the health, education, social development, labor, and environment sectors in Mexico, as well as citizen responses to such announcements.

We also illustrate how different visualization techniques supported by TA4GIP can be combined to offer additional analysis and insights into government-citizen interactions through Twitter. The section is structured into six parts: activity analysis (Section 5.1), identification of salient topics (Section 5.2), analysis of resource sharing (Section 5.3), analysis of citizen interactions (Section 5.4), sentiment analysis (Section 5.5), and combined analysis (Section 5.6).

5.1 Analyzing Activity Levels

A simple approach to analyzing government presence on social media is by quantifying activity levels on Twitter and other channels. Visualization of the activity levels of a government institution's Twitter account helps uncover deviations from normal activity levels that occur at certain times, thus detecting topics that raise citizen



attention. Combining context and activity analysis helps highlight the topics that resulted in higher activity levels and, in turn, identify particular events or issues of special interest to government and the society.

For instance, Figure 2 offers a graphical representation of the activity levels by official Twitter accounts of all five government secretaries between October 2017 and June 2018. In addition to assessing the total number of posts per institution, this helps identify important events that occurred during the time frame for each of the sectors.

Some activity peaks for the social development sector are visible in November 2017, specifically on 6 and 14 of November 2017. The examination of the tweets posted by the secretaries and the hashtags used identifies two major events that caused such peaks, namely “National Youth Award” (#PNJ2017), which took place on 4 November 2017, and “National Commission for Social Development” (#CMD2017), which was held on 14 November 2017. The environment sector had a peak of activity on 7 November 2017, associated with a series of posts dedicated to a parliamentary hearing (#comparecencia). We can also observe high activity levels in the education sector during October and November 2017. The peak on 3 October 2017 corresponds to the return to classes after the earthquakes in September 2017 (#RegresoAClases). Likewise, we can observe several activity peaks in the health sector, related to the “advances in health” events (#AvancesEnSalud) conducted by the Health Commission on 10 October 2017, the Doctor’s Day (#DíaDelMédico) on 23 October 2017, and the International Women’s Day on 8 March 2017. A similar situation occurs with the labor sector, with most activities concentrated in October and November 2017. During October 2017, we can observe tweets related to earthquakes and their consequences in Chiapas (#sisimos, #Chiapas), mostly on 5 October 2017, and tweets on 16 October 2017 about measures to prevent and eradicate child labor (#TrabajoInfantil). Tweets concerning the child labor can be also observed during the peak on 6 November 2017, when the Spain-Americas Forum (#FSCTI) fostered alliances between Spanish and Latin American NGOs to fight against child and adolescent labor.

As shown in Figure 2, Line Charts help identify activity levels but also the levels of citizen responses. We can see the evolution of citizen comments, retweets, and favorites taking place over time in response to the official tweets by government secretaries (see Figure 6). The charts also present the levels and polarity of emotions expressed in the comments, along with their progression over time.

5.2 Identifying Salient Topics

Word Clouds are graphical representations that allow visualizing the relevance of a group of words that summarize a text or a text corpus. The more frequent a word occurs in a text or text corpus, the larger it appears in the word cloud; colors are used for design purposes only. Word clouds are an effective way to highlight the most salient topics in a given context, they are useful to visualize the content of a large number of tweets.


Figure 3 depicts the word clouds generated from the tweets posted by the five secretaries during November 2017. In particular, the most common terms used by the Secretary of Social Development were “social” or “sedesol.” Other terms are #SabíasQue (did you know) and #SedesolPresente (sedesol is here). A common hashtag used in the tweets by the Secretary of Environment and Natural Resources is “#Revillagigedo,” probably related to the signing of the “Revillagigedo National Park Decree,” through which the Government of Mexico committed to the conservation of the Revillagigedo Archipelago. All the words used in the education sector’s word cloud are related to the live hearing (#envivo, #comparecencia) of the Fifth Government Report. The most common words used in the tweets by the Health Secretary are about the project “moving Mexico for your health” (#MoviendoMexicoPorTuSalud). Finally, among labor-related tweets, common words include “work,” “December,” “thirteenth salary,” and “FONACOT credit” (a government financial institution), i.e., #trabajo, #diciembre, #aguinaldo, and #creditoFonacot.

The Word Cloud analysis helps identify major events that affect a given government sector in the period of interest. It should be noticed that while certain terms are common and used at all times, others are specific to an event, as in the example associated with the 2017 earthquake in Mexico.







(a) Social Development




(b) Environment



(c) Education



(d) Health



(e) Labor

Fig. 2. Activity levels of five government secretaries during the period of interest.





5.3 Analyzing Resource Sharing

In Figure 4, colored Euler diagrams depict strategies adopted by the five secretaries for using different resource types in their Twitter communications during the month of January 2018. For instance, the Labor Secretary posted 175 tweets containing links, 384 containing images, and 11 containing videos. It also posted 162 tweets containing both links and images, and one containing links and videos. We can observe that some secretaries, such as the Secretary for Social Development and the Secretary for Labor, use mostly images to share content, while others share images and links. Interestingly, the Education Secretary share a lot more videos than other sectors.

The volume of comments posted by citizens in response to government posts represents a degree of citizen participation. Other measures of participation include instances of citizens retweeting or marking a tweet as favorite.

Digital Government: Research and Practice, Vol. 1, No. 2, Article 15. Publication date: April 2020.

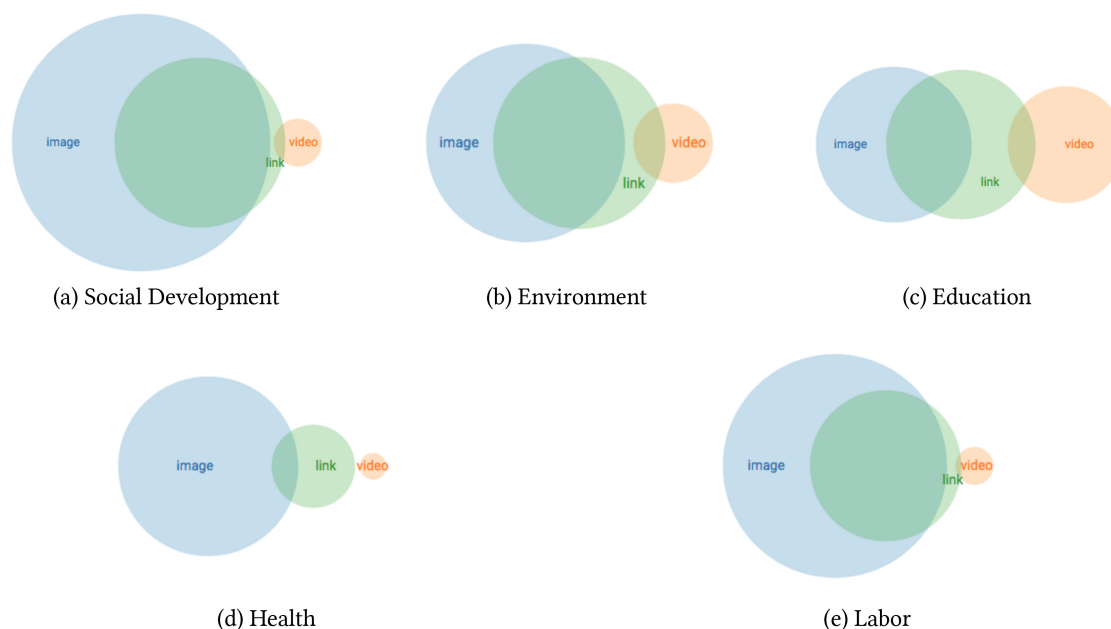


Fig. 4. Different types of resources shared by five government secretaries during the period of interest.

retweets, or favorites. Another level up represents possible combinations of different types of responses and their occurrences. Finally, the upper level of the diagram specifies simultaneous occurrences of comments, retweets, and favorites. Different colors identify different types of interactions, with circles scaled up to represent the numbers of interactions. For clarity, the sizes of circles are not proportional across secretaries.

Figure 5 presents a visualization of the number of comments, favorites, and retweets posted by citizens in response to announcements made by the five secretaries between October 2017 and June 2018. In general, the number of comments is smaller than the number of retweets or favorites. In most cases, the tweets marked as favorites are also retweeted. The Secretary of Environment has the highest level of participation while the Secretary of Education has the lowest level. A comparison between activity levels in Figure 2 and the degree of citizen responses in Figure 5 indicates that the government sectors with the highest levels of activity, such as environment, labor, or social development are also the ones with the highest levels of citizen participation.

The analysis also identifies a correspondence between government activity levels on Twitter and the degree of citizen response. In particular, it can be observed that the Education Secretary has more comments than the Labor Secretary and a similar number as the Health Secretary and the Social Development Secretary, even when all secretaries have lower levels of activity and interaction. While interacting with this diagram through TA4GIP, along with Extended Hasse Diagrams, we can observe the levels of citizen responses over time in Line Charts (Figure 2): time series for retweets (blue), favorites (red), and comments (yellow). We can also observe (Figure 6) that sometimes the number of retweets is significantly higher than the number of favorites, even when these two tend to be similar.

5.5 Analyzing Citizen Sentiments

Sentiment analysis was carried out on citizen comments based on the Plutchik's wheel of emotions [33]. According to Plutchik's theory, there are eight primary emotions: joy, trust, fear, surprise, sadness, disgust, anger, and anticipation. A simple technique to carry out sentiment analysis on a text corpus is to rely on a lexicon that captures word-emotion associations and to count the number of occurrences in the corpus associated with each



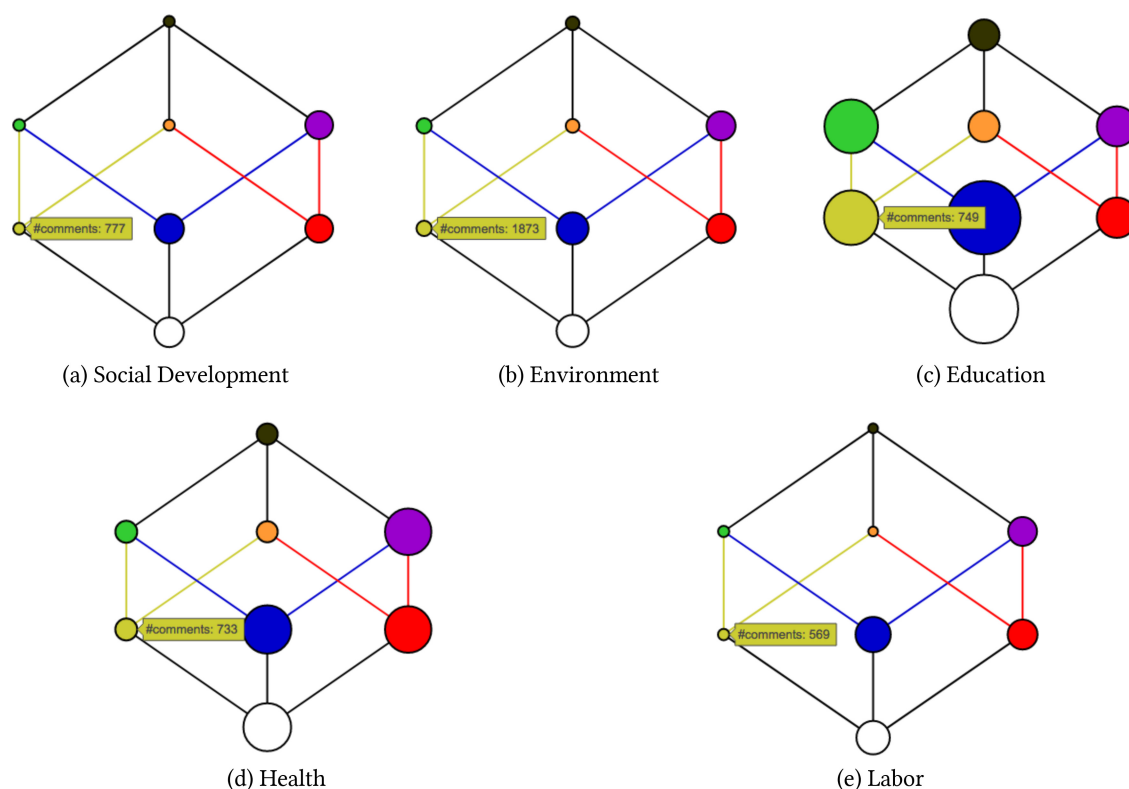


Fig. 5. Citizen responses through comments, retweets, and favorites to the tweets made by the five government secretaries over the period of interest.

emotion. This technique allows estimating the prevailing emotions in the analyzed text corpus. Our analysis was carried out using the Spanish version of the NRC Affect Intensity Lexicon [34]. For illustration, Table 2 contains examples of the Spanish words associated with the eight primary emotions with the corresponding English translations.

This analysis was applied to assess the general mood of citizens in response to the tweets posted by particular government secretaries. After applying this technique to the comments posted by citizens, we generated graphical representations to visualize the prevailing emotions found in the analyzed comments. Figure 7 presents the resulting graphical representations for the five secretaries. The figure allows us to conclude that trust is the most common emotion associated with citizen responses for all five sectors, while in some sectors, such as environment and labor, anticipation is almost as common as trust. We can also note that fear is high in response to the health sector tweets.

Another way to analyze emotions is by means of Bubble Charts. The level of emotions in a given day is represented by a bubble: the greater the bubble, the higher the value of the emotion. The bubble is placed vertically depending on the number of posts published by a secretary under analysis. With this visualization, we can identify the prevailing emotions expressed by citizens in different time periods and the evolution of emotions over time. In Figure 8(a), we can identify high levels of trust as the predominant emotion for the health sector. The visualization allows deselecting some emotions, so we can deselect trust to observe other emotions in the chart (Figure 8(b)). In this figure, we can identify several bubbles of sadness and surprise but also of joy and anticipation. We have seen earlier that fear is more prevalent in the health sector than in other sectors. Additional

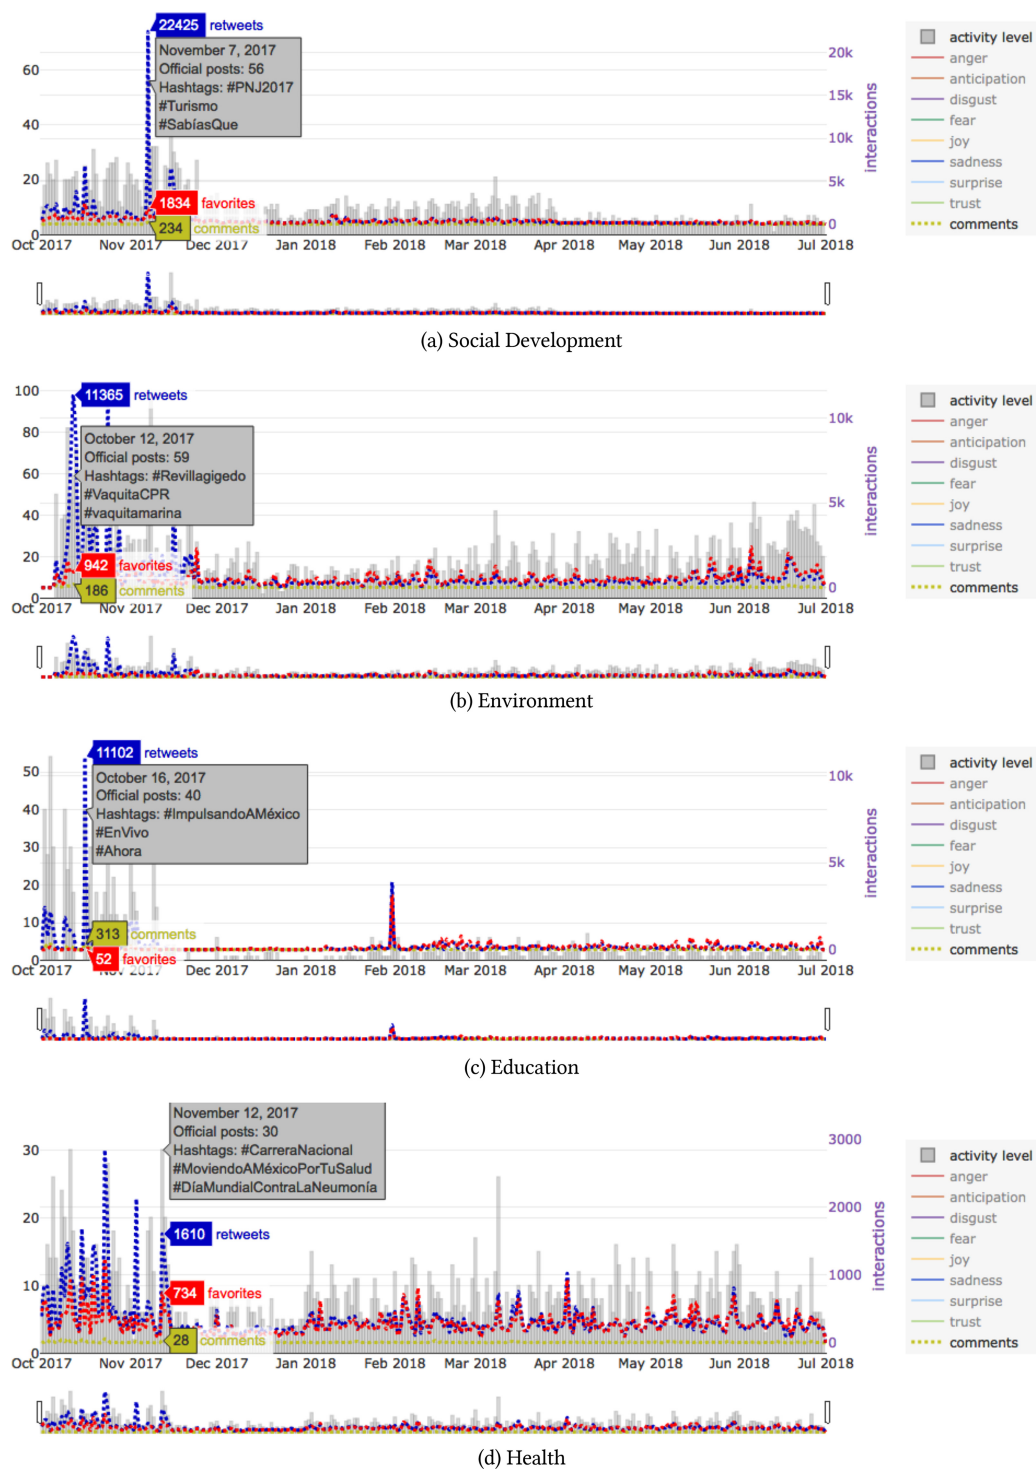



Fig. 6. Continued.





(e) Labor

Fig. 6. Activity levels and citizen responses for the five government secretaries over the period of interest.

Table 2. Spanish Words Associated with Eight Primary Emotions and Their Translations to English [34]

	SPANISH (ENGLISH TRANSLATION)	SPANISH (ENGLISH TRANSLATION)
JOY	felicidad (happiness)	genial (great)
	feliz (happy)	maravilloso (wonderful)
TRUST	lealtad (loyalty)	ayudar (help)
	responsable (responsible)	agradable (nice)
FEAR	asustado (scared)	peligroso (dangerous)
	cobardía (cowardly)	terror (terror)
SURPRISE	asombro (astonishment)	incidente (incident)
	impredecible (unpredictable)	sorprendente (surprising)
SADNESS	condolencia (condolences)	dolor (pain)
	depresión (depression)	duelo (mourning)
DISGUST	deshonestidad (dishonesty)	suciedad (dirt)
	corrupción (corruption)	aberración (aberration)
ANGER	enfado (anger)	queja (complaint)
	molesto (annoyance)	crimen (crime)
ANTICIPATION	intento (attempt)	ansia (craving)
	conjetura (guess)	anhelo (longing)

insights can be gained by analyzing the evolution of emotions over time. In Figure 8(c), we can identify that fear is particularly high during October. We can even compare different emotions during a period of interest. For example, Figure 8(d) shows the presence of fear and anger during the month of October 2017, with peaks for fear near 23 October 2017. Anger is also present when fear is, which may imply the correlation of these two emotions.

5.6 Combining Graphs and Conducting Deeper Analysis

Some diagrams can be used to analyze different aspects in isolation, as explained in the previous sections. For example, the Line Charts in Section 5.1 provide an immediate view of the relation between the level of activity and the level of citizen responses. In such visualizations, we can also analyze the growth of different emotions presented in the comments, to sense citizen moods before certain events are triggered by government tweets.

One might also be interested in exploring a relationship between the information sharing approach adopted by government and citizen responses. For this purpose, we use a Sankey Diagram to show the flow of resources

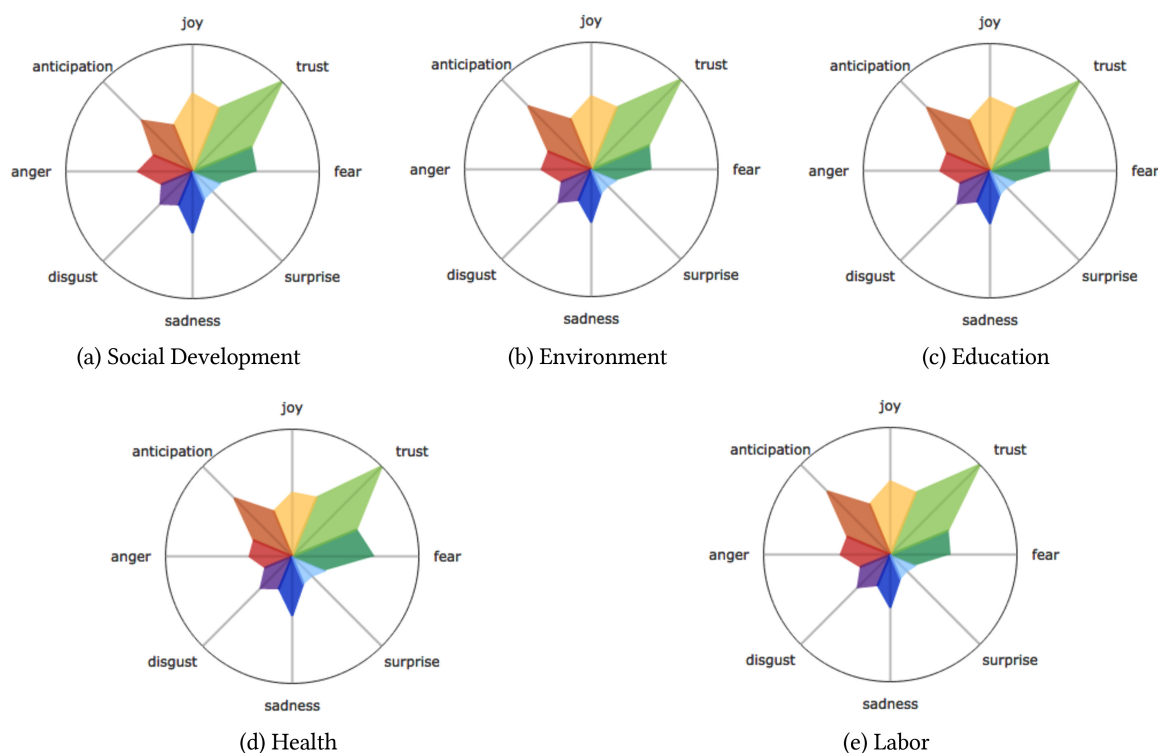


Fig. 7. Emotions expressed by citizens in their responses to the tweets made by the five government secretaries over the period of interest.

shared by the secretaries' Twitter accounts and the corresponding responses from citizens. Figure 9 shows the proportion of the tweets published by a secretary (black vertical line at the left) that share an image (vertical line in light blue), a link (vertical line in green), a video (vertical line in orange), or a text (vertical line in purple). Then, for each of them, we can observe the responses from citizens: retweets (in blue), favorites (in red) or comments (in yellow).


For example, in Figure 9, we can observe that the Social Development, Health, and Labor Secretaries share mostly images while the Education Secretary shares mostly links. We can also conclude that sharing videos is not a very common practice for all secretaries. As previously, the most common responses from citizens are retweets and tweets marked as favorites. The exception is the Education Secretary, where citizens mostly comment the posted tweets rather than marking them as favorites. In general, we have a large flow of tweets that share images, a smaller flow sharing links, and the remaining are sharing videos or include text only. If we observe the incoming flow in citizen responses, then the proportions are preserved. For example, for the Social Development Secretary, we can observe that a large number of retweets are in response to the tweets with images, this is probably because the number of tweets with images is higher than the rest of the tweets.

Through Figure 9, we can also observe that some secretary tweets receive no reaction/interaction from citizens, particularly in the labor sector. Receiving a large number of comments is desirable for all secretaries, since it allows them to know citizen views about public policies they promote. The secretary with the best results in this regard is the Education Secretary. This may be sector-specific but, following this example, other secretaries could also analyze and adapt the communication strategy by the Education Secretary to improve their interaction with citizens.

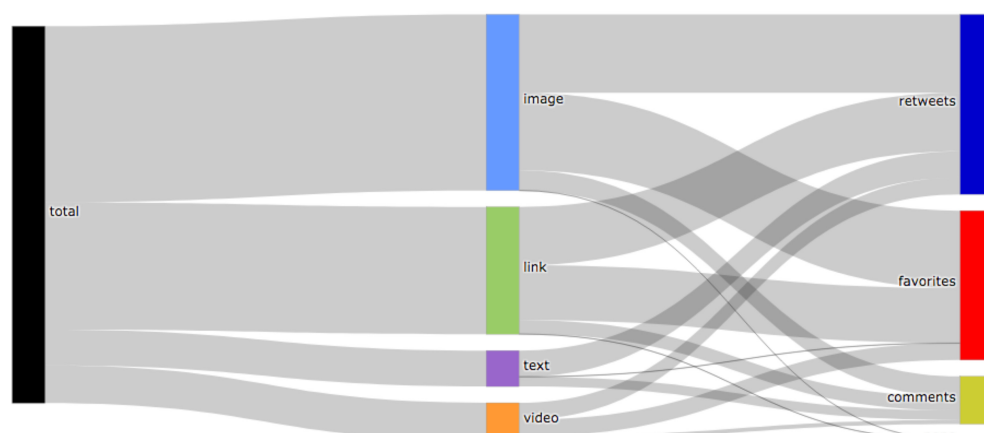





Fig. 8. Emotions expressed by citizens in their responses to the tweets made by the Health Secretary and their evolution over the period of interest.



(a) Social Development




(b) Environment




(c) Education

Fig. 9. Continued.



(d) Health



(e) Labor

Fig. 9. Relationships between links, video, images, and text shared by the five secretaries and citizen comments, retweets, and favorites.

Other forms of combined analysis could include sentiment analysis to identify correlations between types of emotions such as trust or fear in responses to particular policy announcements, or could just focus on detecting the topics that generated the most polarized reactions from citizens, e.g., the responses expressing joy or anger.

6 DISCUSSION

The presented tool, TA4GIP, implements different techniques to analyze government-citizen interactions on social media, specifically on Twitter. Line Charts provide a simple way to instantly observe the level of activity across time and thus help detect peaks, i.e., abnormally high or abnormally low activity levels. The notes embedded in such graphs help identify the topics addressed by showing the most used hashtags. Line Charts also help compare interactions and emotions—we can identify time ranges, activity peaks and citizen emotions used

to respond to posts. Once identified, we can deepen the analysis using other information sources and visualizations. Additionally, we can identify peaks in citizen comments, but these are not analyzed in the scope of this research. Word Clouds are a natural way to identify the topics addressed in government communication over a time period. They can also help identify the most salient issues concerning citizens and what kind of information government should provide to address such concerns. Euler and Sankey diagrams are both good options for identifying communication strategies. The former helps identify multimedia resource strategies adopted by government. The latter depicts the distribution in multimedia resource use and the influence of this distribution on citizen responses. By relating information strategies and citizen responses, government can adopt new ways and new channels for delivering information to improve communication with citizens. Extended Hasse diagrams help identify how citizens communicate with government Twitter accounts, how much they like the published content (favorites), how they analyze posts (retweets), and what they think about information posted (comments). Of interest is also emotional analysis of the retweets, since a shared tweet can be positive or negative, depending on the associated comments. The emotions expressed by citizens can be observed through Radar Charts. Finally, Bubble Charts depict how the sentiments' polarity and intensity change over time.

One limitation of the visualizations presented in this article is their static nature. Although visualizations can be generated for different time periods, it is difficult to visually detect changes over time, especially in relations between communication strategies and levels of citizen interactions. We are currently assessing how to add dynamism and interactivity to our visualizations, to show how the context evolves over time, to improve user experience, and to deepen analysis. We are also exploring changes to on-the-screen arrangements of visualizations to be able to see some of them simultaneously. Another limitation of the tool is related to data recovery, pre-processing, and analysis. Recovered tweets belong to a nine-month period, which prevents the evaluation of major events occurring in different months that could be compared, such as citizen behavior toward the earthquake on September 2017 or toward the hurricane on July–October 2018 (time range not included in this analysis). The preprocessing includes identification of words and multimedia resources used, citizen responses, and detection of duplicate tweets. New versions of the tool should also detect trolling and spamming to limit the influence of noise in the analysis. In addition, we are exploring scenarios for gathering additional information to help understand the context for posting the tweets, including the percentage of users and interactions in relation to the entire population, to avoid under- or over-estimation of the displayed values and the consequent misinterpretation of the results. For the analysis stage, we applied a simple sentiment analysis, which has some flaws. One is related to determining the number of words associated with particular emotions. If the number of words associated with one emotion differs significantly from others, then we could receive unbalanced results, e.g., having a large number of tweets related to “trust” in every analysis. New functionality is also needed for misspellings, words used to express opposite emotions, locally used variations of words, and ambiguities. The currently used sentiment analysis technique can be improved to deal with such issues, or more elaborate techniques can be used instead.

7 CONCLUSIONS

This article discusses the application of visualization techniques to examine government-citizen interactions on Twitter. The techniques are implemented within a software tool called TA4GIP–Twitter Analytics for Government Intelligence and Public Participation—to support this type of analysis. To illustrate the use of the tool, we collected all tweets posted by five secretaries of the Government of Mexico and citizen responses to such tweets over a nine-month period, and applied the implemented visualization techniques to the collected data. We conclude that the tool is useful for analyzing various aspects of government-citizen interactions on social media. These include government activity levels, most frequently used words, abnormal activities, salient topics discussed and general citizen moods. The implemented techniques can also be combined, making possible the analysis of additional features, e.g., the emotions expressed in citizen responses to the tweets posted by particular government secretaries, or relationships between the types of multimedia resources shared and citizen responses.



Such combined analysis augments the capabilities of the tool to derive insights to how citizens view particular government decisions or public policies, identify correlations between citizen responses in terms of the polarity and intensity of the emotions expressed, and particular government announcements, detect topics that generate polarized reactions from citizens, and others.

The prototype software tool described in this article represents a major step in the development of a fully fledged tool that collects information from Twitter, processes this information, and presents the findings in a user-friendly way to facilitate understanding and interpretation. The focus of this work was on describing the visualization techniques used to analyze information and present the finding to non-technical users. A fully fledged tool should support cleansing of the collected data to detect trolling and spamming, and gathering additional information to help understand the context of the information visualized, including, e.g., the percentage of users and interactions visualized in relation to the entire population. Without this contextual information, the visualization can lead to underestimation or overestimation of the displayed values, and the consequent misinterpretation of the results.

The proposed techniques exemplify different approaches to graphically representing the analyzed data. As the case study used in the article concerns a single country and was applied over a limited time period, a broader investigation is needed to validate the generalizability and accuracy of the techniques. Also, the techniques are only a small subset of a range of methods that could be used to visually analyze government-citizen interactions. A more complex analysis could be enabled by integrating these and other techniques and supporting different types of interactions between users and diverse aspects of the presented information. Further research is also needed to explore combined approaches, to apply statistical functions, to triangulate the results produced by various forms of visualizations, to apply the proposed techniques for benchmarking communication strategies between countries with different languages and ways of expressing emotions, and to extend this analysis to other social media channels.

As part of our future work, we plan to explore other techniques to develop visualizations for dynamically changing data, to introduce various forms of combined analysis, to support the piloting of policy initiatives by monitoring citizen reactions to government announcements, and to cover longer periods. We also plan to recover tweets related to events of interest through the use of hashtags, to explore other natural language processing algorithms such as topic modeling to discover the main discussed topics, and to improve the lexicon used for sentiment analysis by adaptation to specific study domains. To improve sentiment analysis, we plan to refine the concept of emotions for a better understanding of their semantics in each government sector; for example, defining the meaning of “trust” in the environment and labor sectors. Finally, the presented techniques can be applied to other social media platforms like Facebook or Instagram, considering their own mechanisms of interaction between users and available retrieval methods through native Application Programming Interfaces.

REFERENCES

- [1] M. D. Flood, V. L. Lemieux, M. Varga, and B. L. L. William Wong. 2016. The application of visual analytics to financial stability monitoring. *J. Financ. Stab.* 27 (2016), 180–197.
- [2] Y. Lerman and Y. Lebendiger. 2017. Visualizing space syntax analyses for decision makers Lessons from Israel. In *Proceedings of the 11th International Space Syntax Symposium (SSS'17)*. 81.1–81.14.
- [3] L. Pépin, P. Kuntz, J. Blanchard, F. Guillet, and P. Suignard. 2017. Visual analytics for exploring topic long-term evolution and detecting weak signals in company targeted tweets. *Comput. Industr. Eng.* 112 (2017), 450–458.
- [4] D. Parygin, N. Sadovnikova, M. Kalinkina, T. Potapova, and A. Finogeev. 2017. Visualization of data about events in the urban environment for the decision support of the city services actions coordination. In *Proceedings of the 5th International Conference on System Modeling and Advancement in Research Trends (SMART'16)*. 283–290.
- [5] C. H. Ku, A. Iriberry, and G. Jena. 2017. Visual analytics for crime analysis and decision support. In *Decision Management: Concepts, Methodologies, Tools, and Applications*, vol. 3–4. IGI Global, 1412–1435.
- [6] J. Kohlhammer, K. Nazemi, T. Ruppert, and D. Burkhardt. 2012. Toward visualization in policy modeling. *IEEE Comput. Graph. Appl* 32, 5 (2012), 84–89.



- [7] T. Ruppert, J. Bernard, and J. Kohlhammer. 2013. Bridging knowledge gaps in policy analysis with information visualization. In *Proceedings of the Series of the Gesellschaft für Informatik (GI'13)*. 92–103.
- [8] T. Ruppert, J. Bernard, H. Lucke-Tieke, and J. Kohlhammer. 2015. Towards a tighter coupling of visualization and public policy making. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST'14)*. 271–272.
- [9] L. G. Forni et al. 2016. Exploring scientific information for policy making under deep uncertainty. *Environ. Model. Softw.* 86, (2016), 232–247.
- [10] G. D. Sun, Y. C. Wu, R. H. Liang, and S. X. Liu. 2013. A survey of visual analytics techniques and applications: State-of-the-art research and future challenges. *J. Comput. Sci. Technol.* 28, 5 (2013), 852–867.
- [11] R. B. Hubert, E. Estevez, A. Maguitman, and T. Janowski. 2018. Examining government-citizen interactions on Twitter using visual and sentiment analysis. In *Proceedings of the ACM International Conference on Digital Government Research: Governance in the Data Age*.
- [12] United Nations Department of Economic and Social Affairs. 2016. *United Nations E-Government 2016: E-Government in Support of Sustainable Development*.
- [13] A. Yousuf Al-Aama. 2015. The use of Twitter to promote e-participation: Connecting government and people. *Int. J. Web Based Commun.* 11, 1 (2015), 73–96.
- [14] L. Porwol, P. O'Donoghue, J. Breslin, C. Coughlan, and B. Mulligan. 2012. Social inclusion and digital divide: EParticipation dilemmas in municipalities. In *Proceedings of the ACM International Conference on Theory and Practice in Electronic Governance*. 389–392.
- [15] K. Jung, S. J. Park, W.-N. Wu, and H. W. Park. 2015. A webometric approach to policy analysis and management using exponential random graph models. *Qual. Quant.* 49, 2 (2015), 581–598.
- [16] G. Stylios et al. 2010. Public opinion mining for governmental decisions. *Electron. J. e-Gov.* 8, 2 (2010), 203–214.
- [17] H. Chen, P. J. H. Hu, H. F. Hu, E. L. Chu, and F. M. Hsu. 2009. AI, e-government, and politics 2.0. *IEEE Intell. Syst.* 24, 5 (2009), 64–86.
- [18] W. Fan and M. D. Gordon. 2014. The power of social media analytics. *Commun. ACM* 57, 6 (2014), 74–81.
- [19] R. Sandoval-Almazán and D. Valle-Cruz. 2016. Understanding network links in Twitter: A Mexican case study. In *Proceedings of the ACM International Conference on Digital Government Research*. 122–128.
- [20] R. Sandoval-Almazán and D. Valle-Cruz. 2018. Facebook impact and sentiment analysis on political campaigns. In *Proceedings of the ACM International Conference on Digital Government Research*.
- [21] Y. Charalabidis, E. N. Loukis, A. Androutsopoulou, V. Karkaletsis, and A. Triantafillou. 2014. Passive crowdsourcing in government using social media. *Transform. Gov. People Process. Policy* 8, 2 (2014), 283–308.
- [22] Y. Charalabidis, G. Gionis, E. Ferro, and E. Loukis. 2010. Towards a systematic exploitation of web 2.0 and simulation modeling tools in public policy process. In *Proceedings of the International Conference on Electronic Participation*. 1–12.
- [23] Y. Charalabidis and E. Loukis. 2011. Transforming government agencies' approach to e-participation through efficient exploitation of social media. In *Proceedings of the European Conference on Information Systems*. 1–12.
- [24] Y. Charalabidis and E. Loukis. 2012. Participative public policy making through multiple social media platforms utilization. *Int. J. Electron. Gov. Res* 8, 3 (2012), 78–97.
- [25] L. Spiliotopoulou, Y. Charalabidis, E. N. Loukis, and V. Diamantopoulou. 2014. A framework for advanced social media exploitation in government for crowdsourcing. *Transform. Gov. People Process. Policy* 8, 4 (2014), 545–568.
- [26] S. Joshi, T. Wandhoefer, and M. Thamm. 2011. Rethinking governance via social networking: The case of direct vs. indirect stakeholder injection. In *Proceedings of the 5th International Conference on Theory and Practice of Electronic Governance*. 277–280.
- [27] A. Kó and T. D. Medeni. 2012. UBIPOL approach to improve policy making—Strengths and weaknesses. In *Proceedings of the European, Mediterranean, and Middle Eastern Conference on Information Systems*.
- [28] L. Porwol, A. Ojo, and J. G. Breslin. 2014. Social software infrastructure for e-participation. *Gov. Info. Q.* 2014.
- [29] F. Artigas and S. Ae Chun. 2013. Visual analytics for open government data. In *Proceedings of the 14th Annual International Conference on Digital Government Research*. 298–299.
- [30] P. Parycek, R. Schollhammer, and J. Schossbock. 2015. Emotions and E-participation of young people: A preliminary study. In *Proceedings of the Annual Hawaii International Conference on System Sciences*. 2415–2424.
- [31] S. Hofmann, D. Beverungen, M. Räckers, and J. Becker. 2013. What makes local governments' online communications successful? Insights from a multi-method analysis of Facebook. *Gov. Info. Q.* 30, 4 (2013), 387–396.
- [32] T. Lansdall-Welfare, V. Lamos, and N. Cristianini. 2012. Effects of the recession on public mood in the UK. In *Proceedings of the 21st Annual Conference on World Wide Web Companion (WWW'12)*. 1221–1226.
- [33] R. Plutchik. 2001. The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. *Amer. Sci.* 89, 4 (2001), 344–350.
- [34] S. M. Mohammad and P. D. Turney. 2010. Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon. In *Proceedings of the NAACL HLT 2010 Workshop to Computational Approaches to Analysis and Generation of Emotion in Text (CAAGET'10)*. 26–34.

Received February 2019; revised June 2019; accepted August 2019