

Understanding Employee Social Media Chatter with Enterprise Social Pulse

N. Sadat Shami¹, Jiang Yang¹, Laura Panc¹, Casey Dugan², Tristan Ratchford², Jamie Rasmussen², Yannick Assogba², Tal Steier³, Todd Soule², Stela Lupushor¹, Werner Geyer², Ido Guy³, Jonathan Ferrar¹

¹IBM, 1 New Orchard Road, Armonk, NY 10504, USA

{ sadat, yangjiang, lpanc, stela.lupushor, jaferrar } @ us.ibm.com

²IBM Research, 1 Rogers St., Cambridge, MA 02142, USA

{ cadugan, tratch, jrasmus, yannick, tssoule, werner.geyer } @ us.ibm.com

³IBM Research, Mt. Carmel, Haifa 31905, Israel { ido, talst } @ il.ibm.com

ABSTRACT

The rise of social media in the enterprise has enabled new ways for employees to speak up and communicate openly with colleagues. This rich textual data can potentially be mined to better understand the opinions and sentiment of employees for the benefit of the organization. In this paper, we introduce Enterprise Social Pulse (ESP) – a tool designed to support analysts whose job involves understanding employee chatter. ESP aggregates and analyzes data from internal and external social media sources while respecting employee privacy. It surfaces the data through a user interface that supports organic results and keyword search, data segmentation and filtering, and several analytics and visualization features. An evaluation of ESP was conducted with 19 Human Resources professionals. Results from a survey and interviews with participants revealed the value and willingness to use ESP, but also surfaced challenges around deploying an employee social media listening solution in an organization.

Author Keywords

Social media; workplace; social analytics; social listening.

ACM Classification Keywords

H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces.

INTRODUCTION

The importance of employee feedback to an organization's well-being is well known and extensively studied [19]. In today's complex business environment constant change is the norm. Traditional top-down decision-making is becoming increasingly difficult because of the complexity

and amount of information needed. Involving employees by enabling constructive, change-oriented communication behaviors, often called 'employee voice' [19], has been shown to improve organizational performance [10, 22]. The notion of employee voice has been studied in many contexts (e.g. [10, 22, 27]). There are various ways in which employees can speak up and provide their input, e.g. direct feedback and discussions with leadership, suggestion boxes, surveys etc.

Different methods of soliciting employee feedback have both advantages and disadvantages. Management willingness to hear employees can create a culture of openness and two-way communication [10]. However, direct feedback to leaders can be hampered by power differences and might only revolve around localized issues. Surveys, on the other hand, are typically anonymous and allow employees to speak more openly without the fear of retribution. When properly administered, surveys can be a powerful tool for understanding employee voice [25]. They allow an organization to obtain responses on targeted questions of interest from a potentially large number of employees. Like any research tool, surveys have their limitations as well. Once a year surveys may not surface emergent issues experienced by an employee population. More frequent (e.g. quarterly) surveys may cause employees to feel survey fatigue and are costly to administer. A proper survey takes time to design, deploy, and analyze. This time requirement may prevent an organization from being aware of or responding to an issue affecting the employee experience right away. Furthermore, if the response rate of a survey is low, it may not be representative of the entire workforce.

The advent of social media in the enterprise provides a new opportunity for organizations to take advantage of the unstructured content being generated by employees. By definition, social media is more informal, and is a medium where employees spontaneously share information and express opinions, often without prompting. Compared to surveys, social media content can potentially be leveraged to provide a more 'real-time' understanding of the employee experience.

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

CSCW'14, February 15 - 19 2014, Baltimore, MD, USA

Copyright 2014 ACM 978-1-4503-2540-0/14/02...\$15.00.

<http://dx.doi.org/10.1145/2531602.2531650>

Social media listening has most commonly been used in domains such as marketing to monitor brand sentiment and engage with customers [9, 20]. It has also proven useful as an early warning system for crises [24] and responding to them [45]. We see great potential in social media listening to provide useful insights about workforce opinions and sentiment. While analyzing social media content is not 'employee voice' in the traditional sense, it does provide an additional source of information for an enterprise to glean business-relevant input from employees. This potential benefit does come with challenges. These include respecting employee privacy, filtering out noise from social media, representativeness, and validity of the data.

Listening to employee social media chatter differs in important ways from other forms of social media listening such as listening to customers. Social media listening by an employer could be viewed by some as employee monitoring with the intent to discipline offending employees. The realization that one's employer may be analyzing public social media chatter can cause an employee to self-censor and temper expressions of negativity. Taken to the extreme, it may reduce levels of participation in public social media.

Fear of speaking up or organizational silence is not a new phenomenon. Studies conducted *before* the advent of social media found that many employees choose the safe option of silence, even when sharing their thoughts would help the organization [33]. Social media is becoming a channel through which many are getting increasingly comfortable to share their thoughts and opinions about various issues. In order to encourage speech, researchers recommend management to explicitly invite and acknowledge employees' ideas through cultivating a culture that makes employees safe enough to contribute fully [11]. By allaying fears of negative repercussions, management can position social media as a real-time channel for employees to express themselves. Unquestionably, this depends on how successful an organization is in engendering the trust of its employees.

In this paper, we introduce Enterprise Social Pulse (ESP) – a social listening tool that allows analysts in the Human Resources (HR) or a related business function of an organization to gain real-time insights from workforce social media content in addition to other channels such as surveys. ESP aggregates employee social media chatter from internal and external sources, combines it with enterprise data, and supports search, data segmentation, filtering, analytics and visualization. The paper is organized as follows. We first review related work on soliciting employee feedback, the use of social media by employees, and tools for understanding employee sentiment from social media. We then discuss the design principles and design requirements that motivated and shaped ESP, and describe its various features. A description of a user study evaluating ESP, and three mini case studies of real-world usage follow. We conclude with a discussion of the issues

encountered in the deployment of ESP, and offer practical suggestions to those interested in deploying such a tool to understand employee chatter.

BACKGROUND LITERATURE

The importance of employee comments and suggestions intended to improve critical work processes are essential to organizational performance [10]. Many scholars have highlighted the virtues of this for organizational health [16, 40] and effective decision making [35]. In a complex world, leaders need input from across the organization to understand the impact of internal practices on the employee experience [43]. Additionally, channels that allow for employee voice are viewed positively, because they signal that employees are valued members of the organization [29]. If employees feel their organization values them, they will be more likely to identify with and trust the organization [12]. Studies have shown that employees evaluate management decisions more favorably when those decisions allow for employee input, even when the input does not have much impact on decision outcomes [46].

Surveys are the most common form through which organizations solicit employees' opinions. Almost three out of every four large organizations in the United States surveys its employees [23]. Employee surveys allow an organization to diagnose itself, evaluate programs and policies, transmit corporate values, and provide an avenue for feedback on decision making [44]. Employee surveys can help management understand the degree to which an organizational strategy is being implemented and the alignment of the organization's policies to achieving strategic goals [41]. Ever since employee surveys were developed in the 1930s, they continue to be a popular instrument for information gathering, communication, and decision making [31]. The data generated through surveys can be quantified and compared against prior time periods or industry averages.

Research has found that a quick response to issues arising from an employee survey is crucial for building employee commitment [28]. Employees become frustrated when the turnaround time for surveys is slow. Employees can perceive this as inaction. Consequently, traditional surveys, with report back 4-6 months after administration, can be viewed as failing at their objective to open the lines of communication between employer and employee [28].

Real-time analysis of social media content of employees may ameliorate this problem of turnaround time. The use of various forms of internal and external social media such as social network sites, blogs, microblogs, forums, and online communities is proliferating in organizations. According to a survey conducted by global consulting firm McKinsey, 65% of companies reported the use of social technologies in their organizations [5]. Surveys of Microsoft employees conducted annually from 2008 to 2011 found gradual increase of the use of Facebook, LinkedIn and Twitter [1].

Increased adoption of social media by employees presents an opportunity for an organization to better understand employees' thoughts and opinions through social listening.

The analysis of unstructured social media text has been used in a variety of contexts from government services [37], to marketing [21] and finance [15]. The domain of 'crisis informatics' provides a good example of how the noise in social media can be reduced to provide positive outcomes in near real-time. For example, Palen et al. [36] provide a framework by which people can judge the helpfulness of information by establishing source-level and individual-level credibility. In the deluge of social media, accuracy is ideal, but not the measure by which people can utilize information. In their framework, *helpfulness* is the critical criteria rather than producing the 'right' answer. Social media use in organizations where identity is clearly delineated may suffer less from concerns of credibility. Nonetheless shifting away from precision as a gold standard in evaluating social media content is useful, and one that we will draw on in understanding employee chatter.

Researchers have built visualizations that allow near real-time analysis of unstructured social media content. Such visualizations help make sense of the large amount of text generated in social media. For example, Eddi provides an interface for browsing Twitter streams that clusters tweets by topics [4]. TwitInfo is a Twitter based event tracking interface that can collect, aggregate, and visualize tweets about user-specified events as they unfold [30]. In an enterprise context, Streamz uses a faceted search approach to provide advanced capabilities of search, navigation, and visualization over an enterprise activity stream [18].

There has been less focus on tools to understand employee social chatter within the context of an enterprise. Previous work has analyzed textual patterns in email to understand phenomena such as organizational hierarchy [14] and newcomer socialization [48]. De Choudhury and Counts analyzed 204,284 posts from 22,968 Microsoft employees that used an internal microblogging tool called OfficeTalk [7]. Using a psycholinguistic lexicon known as LIWC (<http://www.liwc.net>) to analyze the content, they found that employees generally express positive affect throughout the workday, but exhibit a sharp drop in positive affect when working afterhours.

There are several ways we believe our work provides a novel contribution. We are not aware of any tool that allows an organization to intelligently mine the content of chatter generated by employees on internal and external social media. The design principles embedded in our approach respect employee privacy while allowing the organization to make sense and act on employee chatter. Additionally, we provide practical steps for deploying an employee social media listening tool in an organization.

ENTERPRISE SOCIAL PULSE (ESP)

At its most fundamental level, ESP aims to answer the questions: 'What are my employees talking about? And what is the sentiment of the chatter?' It is a tool meant for analysts in Human Resources or other business function of a company, such as workforce communications, with the responsibility to understand employee chatter and identify input valuable to the organization.

ESP requirements and design principles were generated through persona development in interactive workshops among designers, engineers, and HR professionals. For example, HR needed to compare specific employee segments across business units and countries. This led to a requirement to have filters based on employee demographics. The need to understand employee opinions on particular policies or events required the ability to search. The need to understand emerging chatter led to a requirement of topic extraction and top lists.

Design Principles

Augment, not replace

We do not view social media as a substitute for other ways of collecting input from employees (e.g. surveys). Previous research suggests that social media content can supplement but not replace surveys in areas where precise estimates and correlations are required [34]. Rather, the goal of ESP is to 'fill in the gap' in between periodic surveys without putting a burden on the respondents. In addition to the topics traditionally investigated through surveys and other means such as 'Employee Satisfaction' and 'Business Improvement', ESP is intended to provide a real-time expression of employee thoughts, opinions, and sentiment on any emerging or continuing issue. For example, ESP may identify chatter about a social campaign on 'Women in Tech', and the scale of the chatter can be ever evolving and new relevant topics can keep emerging over time.

Establish authentic employee content

A challenge in understanding employee chatter through social media is identifying employees on social media, especially in external social media. We are only interested in understanding the chatter of current employees. While former employees, former interns and other entities may be discussing topics related to the company, analyzing such discussion is outside the scope of ESP. As we explain in the next section, we used a crowdsourced approach to identify current employees and allow them to opt-into having their external social media content mined. Anonymized social media content is matched with demographic characteristics to allow an analyst to segment the data along different dimensions and obtain fine-grained understanding of employee chatter.

Only mine public data

Social media can be both private and public. ESP only mines content that is made public by the author. In other

words, ESP aggregates content that anyone can read and potentially analyze. ESP does not mine email or instant messaging communication. ESP excludes content from countries where there are restrictions of mining public social media content (e.g. Germany, Austria, Netherlands).

Respect employee privacy

ESP operates under the principle that the company is not interested in what specific individuals are saying. Rather it is interested in what they are saying at an aggregated level. If enough employees are saying the same thing, that likely deserves the attention of the company. ESP goes to great lengths to respect employee privacy. First, all social media content is anonymized at the point of collection (details are provided in the next section). Secondly, ESP has different levels of access control. We restrict view access to individual social media posts (called ‘snippets’ in ESP) because search engines often store social media content, and searching for such content can reveal its author’s identity. Snippet-level access is required when there is a need to make sense of the context in which topics are being discussed to provide richer insights to business leaders. Only a handful of ‘admin’ level users (less than 15) can view social media snippets. These ‘admin’ level users are those analysts who already have access to sensitive employee data because of their job responsibilities. Others with ‘business user’ level access can only view analytics in aggregate form.

Within ESP, social media content is associated with demographic characteristics of the authors. In order to prevent inferring identity when searching within ESP, snippets are shown only when the search returns more than 20 snippets. When a search result returns less than 20 snippets, a message notifies the user that results cannot be displayed to protect privacy. Additionally there are only two demographical filters that can be applied at a time, which limits the ability to drill into a very narrow employee segment. Since identity information cannot be linked to individual snippets, the company cannot contact an employee based on what he or she said.

Respect Terms of Use of sites mined

ESP adheres to the Terms of Use of the sites that it mines content from.

System Description

Data flow to ensure privacy

ESP currently collects data from one internal source and one external source but can be configured for additional data sources. Internal data is provided by a social platform used in the company. This platform supports microblog status updates, blogs, bookmarks, and community forums. By default, all employees have a profile on this social platform. All content authored by an employee is associated with their identity – nothing is anonymous.

The external data that ESP analyzes comes from Twitter. We are only interested in Twitter content from our employees. Identifying employees of a particular company on Twitter is a non-trivial task. We use the Twitter user search API to do this. Querying the API with the name of the company returns a set of matching users. We also aggregate users on Twitter lists that contain the company name. From this seed list of users, we ask employees to identify those that they know through an internal crowdsourcing service. Employees were interested in finding their colleagues on Twitter and were generally supportive of performing the identification task (more details on the service can be found in [26]). Once an employee is identified, she receives an email asking if she would like to opt into the service. If she declines, her tweets are not collected. So far, out of 4,462 invited users of the crowdsourcing service, 1,827 (41%) employees have opted-in and 177 (4%) explicitly denied being listed. The remaining users are pending as invitations to this service are distributed through email and can often get ignored. Content from employees that have not taken any action on the email is not analyzed in ESP. Employees who opted in come from all business units across 55 countries, which is a fairly representative sample of the organization. Analysis of the demographic characteristics of those that opted out revealed no significant differences from those that opted in.

Once social media content from the internal and external sources are collected, identity is anonymized using an MD5 cryptographic hash function. A given string (e.g. an email address) will always resolve to the same 32 character MD5 hash. Once hashed, there is no technical way to reverse the hash and recover the original string. Any references to an identity within the social media snippet (e.g. @username) are not hashed. However, as mentioned, the actual social media snippet is viewable to a small set of authorized users.

In order to provide useful segmentation of the data, identity information from the corporate directory and the HR data warehouse are also hashed using the same MD5 cryptographic hash function. This enables ESP to match social media content with employee demographic characteristics, while keeping the identity anonymous.

Text Analytics on Social Media

ESP currently provides two types of analysis: Topic extraction and sentiment analysis. These two types of analysis are available during search and for the demographic segmentations supported by our system.

Topic extraction first extracts candidate noun phrases [39] from the snippets using the OpenNLP open source project [2]. These candidates are then assigned with scores using the KL+TB method, which has been previously shown to be effective for term extraction in social streams [6]. This method first assigns a score with a noun phrase based on its Kullback-Leibler (*KL*) measure [3], which is a non-symmetric distance measure between two given

distributions. KL favors noun phrases that maximize the KL divergence between the language model of the current subset of snippets and the language model of the entire ESP dataset. On top of the KL statistical score, a tag-boost (*TB*) is applied, which promotes keywords that are likely to appear as tags, based on a given well-tagged folksonomy. The folksonomy generated from a social tagging application deployed in the company was used to this end.

Sentiment analysis is the task of finding the opinions of authors about specific entities [13]. ESP uses a commercial sentiment analysis engine to parse social media data and assign a sentiment value (e.g. positive, negative, neutral) to each snippet. Sentiment analysis is still an evolving science since human language is filled with irony, sarcasm, humor etc. This is exacerbated in social media. In order to measure the accuracy of the sentiment, a reliability analysis was performed. A random sample of 500 snippets was collected with their associated sentiment as detected by the sentiment analysis engine. A researcher then went through these same 500 snippets and classified their sentiment. An inter-rater reliability was performed comparing the machine and human rater. The level of agreement for positive, negative, and neutral sentiment was 0.76 ($p < 0.001$) using Cohen's Kappa [8]. Realizing that sentiment classification by the machine needs to be improved for the company to make decisions based on ESP data, the team implemented a feature that allows a human analyst to correct or verify the sentiment of a given snippet.

Once social media content is anonymized, matched with anonymized employee records from the corporate directory

and HR data warehouse, it is run through the sentiment analysis engine, and stored in the ESP database. Data is then surfaced through the ESP user interface. Data is stored and available for querying within less than 5 minutes of being generated, providing near real-time access to employee social media content.

ENTERPRISE SOCIAL PULSE FEATURES

Below is a summary of the various features in ESP. All features are available to all users unless otherwise stated.

Organic results and keyword search

When an analyst wants to get a general understanding of employee chatter within a given timeframe, she can view the 9 portlets in the 'Overview' tab of ESP (Figure 1). This provides 'organic' results without having to enter any search term(s). An analyst can also perform a keyword search when she is interested in understanding the chatter around a particular topic. The content of the 9 portlets will change based on the search term entered. Often times multiple keywords are needed to express a topic. ESP allows the aggregation of multiple keywords into a single user defined 'concept'. ESP supports both keyword searches and searches on 'concepts'. Below we briefly describe the contents of the most important portlets.

Overview

This portlet provides summary information about the volume of social media content and associated sentiment. The sentiment analysis engine breaks the social media content into snippets and uses an NLP approach to

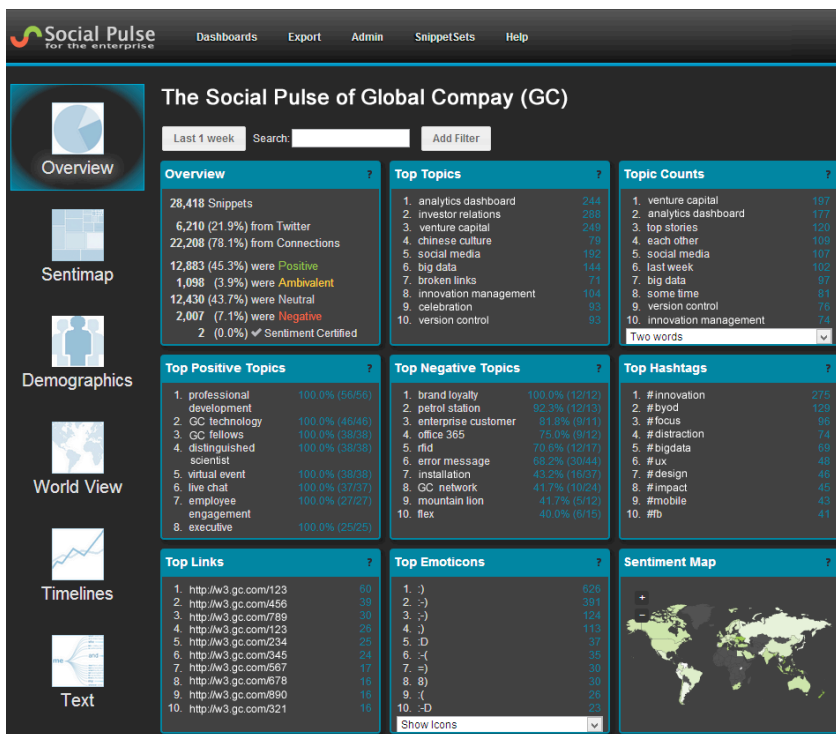


Figure 1. Overview tab in ESP showing 9 portlets.




Figure 2. Sunburst visualization of the number of employees posting social media content by country in a week.

determine the sentiment score of the snippet. If there are one or more positive phrases detected – the sentiment of the snippet is determined as positive. Similarly, negative sentiment is assigned to a snippet when one or more negative phrases are found. Ambivalent sentiment indicates that an equal amount of positive and negative phrases have been found. Neutral means the snippet does not contain any positive nor negative phrases.

Top Topics

This portlet helps users identify prominent and trending topics in the set of snippets being analyzed, based on the KL+TB measure explained previously. The KL+TB score can sometimes be higher for topics with lower counts (listed next to the topics).

Topic counts

This portlet lists top topics based on strict frequency counts. It allows showing one-word topics, two-word topics, or three-word topics. It offers a different way to identify important topics in parallel with Top Topics.

Top positive topics and Top negative topics

These two portlets list top topics with extreme sentiment scores. The system first calculates top-k topics using the KL+TB measure (currently k is set to 1,000) that pass a count threshold (currently set to 10) and considers those that have the highest portion (shown next to the topics) of positive or negative snippets as Top Positive or Top Negative, respectively. This feature is particularly helpful in identifying topics with strong sentiment and thus of great interest, but potentially with less volume than topics in the general Top Topics list.

Sentiment map

This portlet visualizes the geographical distribution of sentiment such that users can quickly compare across countries. In the world map, each country's color reflects its ratio of positive (more green) to negative (more red) snippets.

Data segmentation and filtering

ESP supports segmentation of data based on characteristics such as business unit, geography, job category, work location, years of service, manager/non-manager, performance rating etc. In fact, there are 40 filters to segment the data. However only two can be applied at a time given the privacy considerations discussed earlier. Any search on ESP can be filtered based on these dimensions. Search and segmentation based on filters are available in all the tabs in ESP. For example, the 'Demographics' tab in ESP provides a breakdown of the set of snippets being analyzed by characteristics such as work location, manager / non-manager, performance rating, business unit, country etc. Figure 2 shows a sample visualization of the 'employees by country' portlet in the 'Demographics' tab

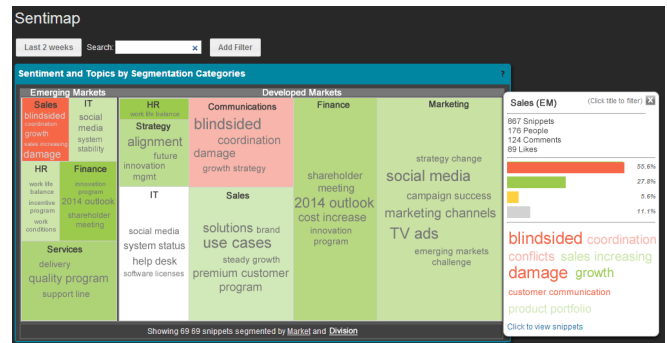


Figure 3. 'Sentimap' Tree Map visualization.

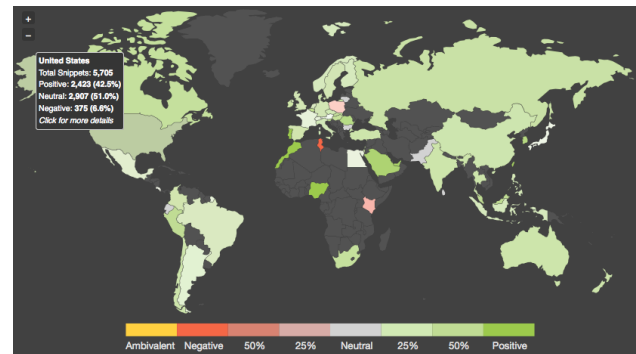


Figure 4. 'World View' for a given week. Sentiment color is overlaid on top of countries, and hovering over a country reveals more details.

for a given set of snippets. Additionally, users can specify or search within a custom date range.

Custom dashboard

An analyst may want to follow a particular business unit, geography or other demographic segment. ESP has the ability to set up all the 9 portlets in the 'Overview' tab pre-filtered based on an analyst's needs and saved such that an analyst can monitor the social stream of interest without having to specify filter(s) every time.

Analytics and visualization

ESP supports a variety of interactive visualization features. We briefly describe them below.

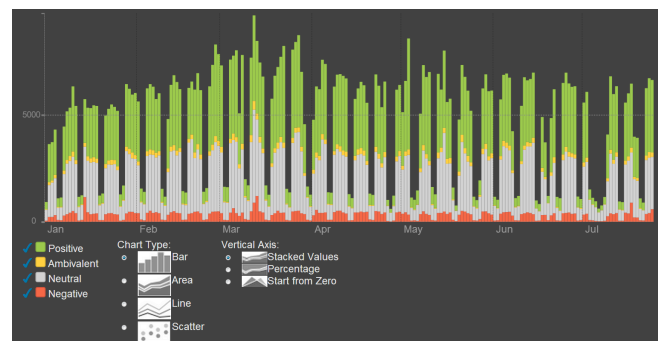


Figure 5. 'Timelines' tab showing sentiment over time.

keyword being searched (Figure 6). Only users with 'admin' level access are allowed to view this tab. This visualization lets an analyst search for a word and see common phrases that appear before or after it, weighted by occurrence in the set of snippets being analyzed.

Data export

An analyst may want to perform advanced statistics or visualize the data in a manner different from ESP. In order to support this, ESP allows up to 10,000 rows of the social media snippets augmented by selected demographic data to be exported as a comma or tab delimited file. This feature is only available to users with 'admin' level access.

Table of snippets

Snippets can be sorted by date, source, sentiment, number of likes and number of comments. They can be individually selected to fix their sentiment or added to a collection of snippets (snippetset) for export and further analysis. This is only available to users with 'admin' level access (Figure 7).

Implementation

The system is implemented as a J2EE web application across two server nodes. One server manages the real-time data collection and serves as database server, the other server manages the web application and indexing of the data. We used Lucene (<http://lucene.apache.org/>) as a basic indexing technology but applied the organization's own adaptation of faceted search to support the various segmentation options described above. Sentiment analysis is done with a commercial sentiment analysis package. The user interface and visualizations are mostly based on the d3-dojo (<https://github.com/sdqali/d3-dojo>) and jQuery (<http://jquery.com/>) libraries.

DATA VOLUME

Figure 8 shows the volume of social chatter for the entire months of February through April 2013. We see a consistent pattern of more positive than negative snippets across time. Figure 9 shows how different job roles have different levels of posting behavior during the same time period. Software developers, and employees in Sales and Marketing & Communications are the most active. The data shows that there is representation of all business units and job roles in the company, albeit in varying degrees. The demographic distribution of posts largely matches the employee distribution, with software developers and sellers being a large proportion of the employee population of the organization.

USER EVALUATION

We conducted a user study to evaluate how well ESP meets the needs of HR professionals interested in understanding employee social media chatter. We utilized both quantitative and qualitative methods for this purpose.

Participants

We recruited 19 HR professionals of a large Information Technology company through email invitation. The majority of participants had been at the company for over five years. Participants originated from 4 countries (USA, Canada, China, Australia) and 10 of them were female. They held job responsibilities such as workforce analytics, employee engagement, employee & labor relations, talent partner, workforce enablement, HR location leader, and HR director. They had non-technical backgrounds.

Training

Participants were first required to complete an online course on data privacy. The course lasted about half an hour. The goal of the course was to educate participants on handling sensitive employee data. Following the course, participants were required to attend an hour-long training session on how to use ESP. In addition to receiving a tutorial on how to use the different features of ESP, this training session allowed participants to ask any questions they had. The training was recorded and the video was made available to access afterwards as needed. Additionally, we were available for one-on-one follow-up sessions with participants in case they needed further instructions or clarifications on using ESP. For the purposes of the study all participants were temporarily given admin access.

Task

Participants were given two weeks to familiarize themselves with ESP and complete a research task. The task consisted of thinking of any business question that participants wanted to answer, and then using ESP to explore that question. Results of the analysis were to be documented in a report. The report, similar to the task, was open-ended and participants could choose any format, framework, or length for their reports. The goal of this task-oriented method was not to study or compare task completion effectiveness but rather encourage authentic use of ESP in order to collect qualitative data about the value and usage of ESP for HR professionals. Some examples of business questions our participants explored include understanding employee sentiment on professional development and training, performance reviews, work life balance, LGBT and marriage equality, the onboarding experience, seller experience etc.

Survey and Interview

Upon submission of their reports, participants were asked about their experience using ESP through a survey. The survey consisted of 40 questions and took approximately 15 minutes to complete. The majority were Likert-scale questions to evaluate individual or overall features, with some open-ended ones asking for insights and comments. 18 of 19 participants completed the survey. At the conclusion of the survey, participants had the option of signing up for an interview. 12 of 19 participants signed up to be interviewed.

We conducted semi-structured interviews to elicit richer understanding of users' experience and expectations. We followed a semi-structured interview protocol, referred to interviewees' survey responses and reports, and probed when necessary. Two people conducted most interviews, with one person guiding the interview while the other took notes. Interviewers were not involved in the engineering and design of ESP. This was mentioned to interviewees to encourage them to be frank and honest in their responses. All interviews were recorded with the participants' permission and lasted on average about 45 minutes. Repeatedly listening to recordings and working over transcripts allowed us to identify emergent themes from interviews. Overall, the study lasted two weeks, with an additional day for those that signed up to be interviewed.

RESULTS

Survey: Evaluating ESP features

An overwhelming majority (17 of 18) of our participants see themselves using ESP going forward. About a third would use ESP monthly, while another third would use it weekly and a quarter of participants would use after an event of interest. 4 of 18 respondents are currently conducting searches/analysis on employee social media chatter (without using ESP). Most of those conducting social media searches believe that ESP can considerably

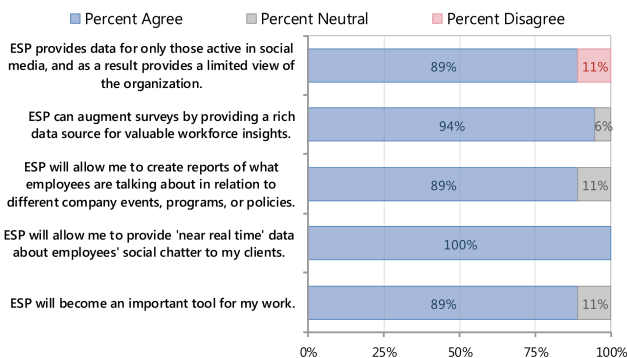


Figure 10. Survey results. N = 18.

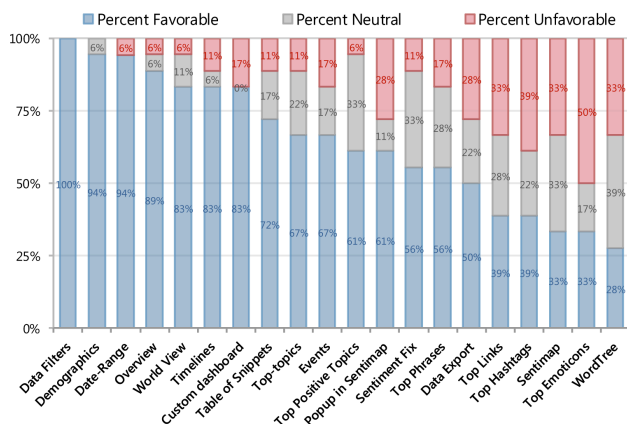


Figure 11. Usefulness of various ESP features. N = 18.

improve their understanding of employee chatter. We were pleased to hear this from that sub-population as they are familiar with the challenges of social media listening. The responses to our Likert-scale questions (Figure 10) reveal that the majority of respondents view ESP as a valuable tool to gather workforce insights. However, concerns remain about the representativeness of the data, which is consistent with prior research [34].

Figure 11 shows responses to a question asking about the usefulness of different ESP features ("How useful do you consider this feature to your analytic/research work?"). The scale ranged from 1=Not useful at all, to 5=Very useful. Values of 4 and 5 are pooled into the 'favorable' category while values of 1 and 2 are 'unfavorable', with a value of 3 considered 'neutral'. Data filters were found unanimously to be the most useful (100%). Also scoring high are the 'Demographics' tab, the ability to search/filter within a date range, and the 'Overview' tab. The 'Word Tree' was considered the least favorable. This visualization uses Java. Many participants reported experiencing issues with Java on their browsers.

Interviews: Deeper investigation

Our interview sessions provided further fine-grained detail and richer insights about the experience of using ESP and the challenges encountered.

General positive reaction

The interviews confirmed the positive findings of the survey.

"Until you have a tool like Social Pulse where you can actually read the comments yourself and interpret what that reaction is - now you have some strong intellectual capital to use to either explain to our executives that everything is okay, or explain to them what some of the root cause issues are being identified."

Participants felt ESP provided a unique service not available through other tools.

"Without this tool, how would I be able to access all the forums, communities, blogs to see the information? It's not possible."

The majority of participants appreciated the flexibility of the tool and its drill-down capabilities.

"There is a really good blend of high level synopsis stuff as well as down in the weeds detail oriented stuff. And you can go to whatever level of detail you choose to."

Value of data segmentation

Several participants mentioned the value of demographic segmentation to understand employee chatter.

"It is useful to get those kind of demographics - what part of the business is something coming from? Is it a global issue, or is it a predominantly North American issue? ..."

The issues that you are going to experience in India and China could be significantly different in that the maturity of the business is so young, you know, the issues that you are facing in a major market may be totally different. The things you want to do or focus on will be different depending on your demographics.”

“Demographics will tell you a lot about the data. Does a manager have a different view than an individual contributor? Does a high performer have a different view than a low performer? Or, does everyone have the same opinion, and what does that tell you? Being able to classify into meaningful groups helps give you more context to draw conclusions about what the data is actually telling you.”

Value of real-time

The majority of participants felt the main value in ESP was in its capability to provide real-time insights.

“I left Social Pulse running in the background, and in the course of an hour, I could see the comments increasing. It was very real-time, which was impressive to me.”

Different uses of the tool

In addition to the business questions participants showcased in their reports, they were often inspired to talk about more possibilities of using ESP, for example:

“In a micro way, following an incident of some type I think it will be very insightful. Where is social media lighting up regarding that incident or topic?”

“I can see particularly in the acquisition/integration side of my job incorporating some of the Social Pulse analysis and sentiment tracking. Morale and change-management in acquisitions is a very big topic. It’s commonly reported that people issues and culture change are the reasons for failure in acquisitions.”

Lack of data

Some participants found that there wasn’t enough social chatter on their topics of interest to provide business value.

“The more detailed I got in my search, I would pretty much go from 20 snippets to 0.”

Participants noticed that there were usually more positive comments than negative comments, and that employees may be self-censoring.

“The one thing that I worry about is that there aren’t many negative comments. Obviously we’d like to work in a positive environment. But I believe, my own theory is that there is more negativity on some of these subjects than what is appearing [in ESP]. That causes me concern that we would be drawing conclusions when people are self-filtering.”

Lack of confidence in sentiment analysis

Participants understood the limitations of the sentiment analysis algorithms. However, this did not keep them from raising concerns.

“If I’m taking this very seriously, then it is imperative that I take these snippets as seriously as well. And for that the sentiments have to be correct, you know... It’s important that it’s correct.”

BRIEF CASE STUDIES OF POST EVALUATION USAGE

Based on the results of the evaluation, the organization has decided to rollout ESP to a wider population of HR professionals and other employees whose job responsibilities involve understanding the employee experience and/or analyzing social media chatter. Below we describe three mini case studies that showcase the usage of ESP to deliver business value.

Understanding overall health of the company

A monthly report based on exported data from ESP summarizing social media trends at the company such as changes in volume, sentiment, and top topics segmented by geography and business unit is sent to senior leaders, including the CEO. After the first of such reports was sent, the senior leadership has continued to request monthly updates and analysis of changes over time. The adoption of ESP reports as part of a wider set of reports sent to senior leaders to help understand the impact of change initiatives within the company is suggestive of the insights ESP can provide. Senior leaders roll out different communications, change initiatives, programs and policies in the company, and they are interested in understanding if these generate any employee chatter. The ability to search social media content using specific keywords related to these communications, programs and policies, and analyze chatter over time, provides useful input to decision makers in understanding if employees are adopting the change, and sharing information about it through social media.

Understanding reputation as an employer

A specific business unit of the organization was interested in the perception of their employees about the work experience in growth markets. A team examined the internal Human Resources Management System (HRMS) to analyze data on employees voluntarily leaving the business unit in growth market countries. They supplemented this data with anonymous snippet commentary from ESP. Searching for keywords related to career, recognition, strategy, and job/work allowed the team to surface content that provided additional context and richness to the quantitative data extracted from the HRMS. This provided insights into specific attrition levers that the business unit could use to implement pre-emptive strategies to address voluntary attrition of high performers. Additionally, we found that certain keywords (e.g. re-org) are not used much, likely because they touch on sensitive issues. A further challenge with a keyword-based approach is that employees

might be discussing the topic of the keyword without using the specific keyword itself. To address this limitation, we encouraged employees to use hashtags. This allows them to view related content by clicking on a hashtag in the social platform used by the company, and also enables ESP to aggregate related snippets.

Fine-tuning and aligning internal and external messaging

The business unit of the company that specializes in outsourcing wanted to know if the language used in relation to the term was consistent across the company. This can improve the brand messaging of the business unit around outsourcing. A team uses ESP to measure this on a monthly basis. Additionally, they examine if the volume of conversation around outsourcing is growing, and identify demographics of employees engaging in these conversations. Using ESP to analyze data over several months, the team found that the terms ‘outsourcing’, ‘company’ and ‘BPO’ were most used externally, while ‘outsourcing’ and ‘sourcing’ were mentioned most internally. Through this analysis, the team recommended the use of the term ‘sourcing’ in all internal and external messaging in order to have consistent communication. Moreover, they were able to identify peaks in conversation and the reasons behind the peaks. For example, ESP data revealed that a service agreement contract led to a significant peak in volume toward the end of July 2013. Most of the posts on outsourcing were from sellers and employees in the software business unit, accounting for 42% of the conversation in total.

DISCUSSION

With the advent of social technologies, communication is increasingly going from one-way and two-way to multi-directional. Employees are already active on social media. As more people that have grown up with social technologies enter the workforce, social media in general and internal social media in particular will become a strong platform for employee expression. Organizations are beginning to realize the value of instant, real-time data on what their employees are thinking. Consequently, tools are needed to mine social content in order to deliver insights to the organization about the employee experience. Our design of ESP was a means to that end of delivering immediacy of insight. ESP promotes understanding employee social chatter while preserving the privacy of employees.

Our interviews revealed that HR professionals did not really have any tool to systematically gather and analyze employee social media content prior to ESP. The employees that performed searches would frequent employer review sites, or do manual searches on Google or Twitter. However there was no way to verify that the people posting were actually employees of the company. Perhaps this explains why our study participants were so enthusiastic about continuing to use ESP in the future. The

in-situ evaluation of ESP using employees to perform a real business task lends ecological validity to the findings. The continued use of ESP, such as those mentioned in the mini case studies, shows the promise of ongoing business insights.

The wider use of ESP in the organization paves important avenues for future research. Tools such as ESP are only as good as the data employees are willing to provide through their participation in social channels. Without strong support from management and a culture that encourages employees to voice their thoughts, tools such as ESP cannot be effective. We plan to conduct a longitudinal study to draw out the nuances in which ESP may play a role in understanding the employee experience. At the very least, the existence of ESP provides a signal to employees that the organization is interested in listening to their thoughts to drive positive change. In order to strengthen that signal, we are planning to make some of the ESP visualizations publicly available inside the organization.

A challenge with social media is the representativeness of the data. Surveys also face this challenge when low response rates are encountered. Our participants clearly felt that ESP provides over representation of employee segments active in social media, and under representation of others. With ESP, the results are shown along with the underlying employee demographics of the data. This helps analysts interpret the data better, and aids decision-making. Contextualizing the data in such a manner allows a business leader to determine how much credibility she will lend to the social data during decision-making.

In deploying ESP, the greatest challenge we faced was with ensuring we respect employee privacy. We have discussed the technical solution in previous sections. Just as important is handling the organizational processes. We went through a couple of steps to ensure global deployment of ESP. First, we had a series of meetings with the Chief Privacy Officer and her legal team to make sure ESP conforms to the company’s Social Media Policy and respects all labor laws. We worked with work councils in the European Union to determine which data fields could be collected. We adhered to all country specific work council restrictions. We developed a Terms of Use (ToU) of ESP in consultation with the company’s legal team. The ToU clearly specifies the purpose of ESP and how data will be used. The salient aspects of the ToU include the prohibition of sharing ESP confidential data with others that do not have a business need, and that violation of agreed usage terms can result in disciplinary action.

An important component of the deployment plan was to communicate to employees through an article published on the Intranet and shared through social channels. The article includes a description of ESP, who has access to it, and what it is being used for. It is important to be transparent with employees and obtain consent when required. We have plans for deploying a public widget on the Intranet

incorporating data from some of the portlets in the 'Overview' tab, allowing all employees to see the type of analytics and visualization an analyst sees. However, this was not part of the initial deployment.

One of the findings of our user study was that many employees appeared to be reluctant to express overtly negative sentiment. Out of the 2.6M snippets in ESP to date, 51.8% are neutral, 37.8% are positive, 7.2% are negative, and 3.1% are ambivalent. This is consistent with prior studies showing that even the most proactive or satisfied employees are likely to "read the wind" as to whether it is safe and/or worthwhile to speak up in their organization [32]. At the same time, research has found that individuals have a strong need for control over decisions that affect them [17, 38]. An important way in which employees gain a sense of control over their environment is by expressing their opinions [29]. Through such expression they can make their opinions heard and influence decisions that affect them. Our motivation behind creating ESP was to allow organizations to make the opinions and sentiment of their employees more accessible. Although ultimately it comes down to a matter of trust, we believe the privacy features enabled in ESP can encourage more employees to express themselves without fear of negative consequences. Furthermore, an organization would rather have employees express negative sentiment behind the firewall rather than repress it or express it externally.

Another issue that surfaced in our user study was the lack of relevant content related to a specific business question. There are two issues here: one involving noise, and another regarding the sparseness or complete lack of data.

With respect to noise, we refer to Palen et al.'s notion of 'helpfulness' and moving away from precision as a gold standard in retrieved social media results [36]. In this context, helpfulness is *relative to what is needed*. In relation to employee chatter, often times what is needed is some clue that can allow an analyst to dig deeper and determine if something is of concern or not.

With respect to the sparseness or lack of data, we anticipate that as more employees embrace social media, there will be more content to mine. We plan on adding more data sources as ESP has the ability to manage multiple data sources. Passive listening may not surface data from all employee segments. To compensate for under representation of certain employee segments in the data, a more 'active' and directed approach could be taken. For example, a quick poll targeted at an employee segment that is under-represented in ESP data can help provide more complete data in answering a business question. In our future work, we plan to integrate quick polls into ESP to further enrich the capabilities of surfacing employee sentiment.

CONCLUSION

In this paper we have described a new tool for mining internal and external social media to understand employee

chatter while respecting employee privacy. ESP was found to be very useful in providing near real-time data about employee social chatter, which can allow analysts to create better reports describing the employee experience. We believe that ESP will be beneficial for companies with an active social media footprint and the business need to understand employee chatter. Technological challenges such as the limitations of sentiment analysis, and cultural barriers such as self-censorship and the lack of relevant social media content remain. As technology advances, one can see the technological challenges dissipate. The cultural barriers may need more time to overcome.

REFERENCES

1. Archambault, A. and Grudin, J. A longitudinal study of facebook, linkedin, and twitter use. In *Proc. CHI 2012*, ACM Press (2012).
2. Baldridge, J. *The openmlp project*, 2005.
3. Berger, A. and Lafferty, J. Information retrieval as statistical translation. In *Proc. SIGIR 1999*, ACM Press (1999).
4. Bernstein, M. S., Suh, B., Hong, L., Chen, J., Kairam, S. and Chi, E. H. Eddi: interactive topic-based browsing of social status streams. In *Proc. UIST 2010*, ACM Press (2010).
5. Bughin, J. and Chui, M. The rise of the networked enterprise: Web 2.0 finds its payday. (2010), Retrieved May 27, 2013, from http://www.mckinsey.com/insights/high_tech_telecoms_internet/the_rise_of_the_networked_enterprise_web_2_0_finds_its_payday.
6. Carmel, D., Uziel, E., Guy, I., Mass, Y. and Roitman, H. Folksonomy-based term extraction for word cloud generation. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 3, 4 (2012).
7. Choudhury, M. D. and Counts, S. Understanding affect in the workplace via social media. In *Proc. CSCW 2013*, ACM Press (2013).
8. Cohen, J. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20 (1960), 37-46.
9. Crawford, K. Following you: Disciplines of listening in social media. *Continuum: Journal of Media & Cultural Studies*, 23, 4 (2009), 525-535.
10. Detert, J. R. and Burris, E. R. Leadership behavior and employee voice: Is the door really open? *Academy of Management Journal*, 50, 4 (2007), 869-884.
11. Detert, J. R. and Edmondson, A. C. Why Employees Are Afraid to Speak. *Harv. Bus. Rev.*, 85, 5 (2007), 23-25.
12. Eisenberg, R., Fasano, P. and Davis-LaMastro, V. Perceived organizational support and employee

- diligence, commitment, and innovation. *Journal of Applied Psychology*, 75 (1990), 51-59.
13. Feldman, R. Techniques and applications for sentiment analysis. *Commun. ACM*, 56, 4 (2013), 82-89.
 14. Gilbert, E. Phrases that signal workplace hierarchy. In *Proc. CSCW 2012*, ACM Press (2012), 1037-1046.
 15. Gilbert, E. and Karahalios, K. Widespread worry and the stock market. In *Proc. ICWSM 2010*, (2010).
 16. Glauser, M. J. Upward information flow in organizations: Review and conceptual analysis. *Human Relations*, 37 (1984), 613-643.
 17. Greenberger, D. B. and Strasser, S. The development and application of a model of personal control in organizations. *Academy of Management Review*, 11 (1986), 167-177.
 18. Guy, I., Steier, T., Barnea, M., Ronen, I. and Daniel, T. Swimming against the streamz: search and analytics over the enterprise activity stream. In *Proc. CIKM 2012*, ACM Press (2012).
 19. Hirschman, A. O. *Exit, voice, and loyalty: Responses to decline in firms, organizations and states*. Harvard University Press, Cambridge, MA, 1970.
 20. Hoffman, D. L. and Fodor, M. Can you measure the ROI of your social media marketing. *MIT Sloan Management Review*, 52, 1 (2010), 41-49.
 21. Kaschesky, M., Sobkowicz, P., Lobato, J. M. H., Bouchard, G., Archambeau, C., Scharioth, N., Manchin, R., Gschwend, A. and Riedl, R. Bringing Representativeness into Social Media Monitoring and Analysis. In *Proc. HICSS 2013*, IEEE Computer Society (2013), 2003-2012.
 22. Kiker, D. S., Kiker, M. B. and Perry, M. Constructive challenge: employee voice, helpfulness and task performance on organizational rewards. *Journal of Management and Marketing Research*, 13 (2013).
 23. Kraut, A. Moving the needle: Getting action after a survey. In A. Kraut (Ed.), *Getting action from organizational surveys: New concepts, technologies and applications*. San Francisco: Jossey-Bass, 2006.
 24. Lamos, V. and Cristianini, N. Tracking the flu pandemic by monitoring the social web. In *International Workshop on Cognitive Information Processing (CIP) 2010*, (2010), 411-416.
 25. Lawler, E. *High-Involvement Management. Participative Strategies for Improving Organizational Performance*. Jossey-Bass Inc. Publishers, 1986.
 26. Lee, T. Y., Dugan, C., Geyer, W., Ratchford, T., Rasmussen, J., Shami, N. S. and Lupushor, S. Experiments on Motivational Feedback for Crowdsourced Workers. In *Proc. ICWSM 2013*, (2013).
 27. LePine, J. A. and Van Dyne, L. Voice and cooperative behavior as contrasting forms of contextual performance: Evidence of differential relationships with Big Five personality characteristics and cognitive ability. *Journal of Applied Psychology*, 86, 2 (2001).
 28. Lilienthal, P. If you give your employees a voice, do you listen? *The Journal for Quality and Participation*, 25, 3 (2002), 38-40.
 29. Lind, E. A. and Tyler, T. R. *The social psychology of procedural justice*. Plenum, New York, 1988.
 30. Marcus, A., Bernstein, M. S., Badar, O., Karger, D. R., Madden, S. and Miller, R. C. Twitinfo: aggregating and visualizing microblogs for event exploration. In *Proc. CHI 2011*, ACM Press (2011).
 31. McConnell, J. *How to Design, Implement, and Interpret an Employee Survey*. Amacom, New York, 2003.
 32. Miliken, F. J., Morrison, E. W. and Hewlin, P. F. An exploratory study of employee silence: Issues that employees don't communicate upward and why. *Journal of Management Studies*, 40 (2003), 1453-1476.
 33. Morrison, E. W. and Milliken, F. J. Organizational silence: A barrier to change and development in a pluralistic world. *Academy of Management Review*, 25, 4 (2000), 706-725.
 34. Murphy, J., Kim, A., Hagood, H., Richards, A., Augustine, C., Kroutil, L. and Sage, A. Twitter Feeds and Google Search Query Surveillance: Can They Supplement Survey Data Collection? In *Association for Survey Computing 6th International Conference*, (2011).
 35. Nemeth, C. J. Managing innovation: When less is more. *California Management Review*, 40, 1 (1983), 59-74.
 36. Palen, L., Vieweg, S. and Anderson, K. M. Supporting "Everyday Analysts" in Safety-and Time-Critical Situations. *The Information Society*, 27, 1 (2011), 52-62.
 37. Paris, C. and Wan, S. Listening to the community: social media monitoring tasks for improving government services. In *Ext. Abstracts CHI 2011*, ACM Press (2011), 2095-2100.
 38. Parker, L. E. When to fix it and when to leave: The relationship among perceived control, self-efficacy, dissent and exit. *Journal of Applied Psychology*, 78 (1993), 949-959.
 39. Rijkhoff, J. *The noun phrase*. Oxford University Press, 2004.
 40. Saunders, D. M., Sheppard, B. H., Knight, V. and Roth, J. Employee voice to supervisors. *Employee Responsibilities and Rights Journal*, 5 (1992), 241-259.
 41. Schneider, B., Ashworth, S., Higgs, C. and Carr, L. Design, validity and use of strategically focused employee attitude surveys. *Personnel Psychology*, 49, 3 (1996), 695-705.
 42. Shneiderman, B. Tree visualization with tree-maps: 2D space-filling approach. *ACM Transactions on Graphics (TOG)*, 11, 1 (1992), 92-99.

43. Senge, P. M. *The fifth discipline: The art and practice of the learning organization*. Doubleday, New York, 1990.
44. Smith, F. *Organizational Surveys: The Diagnosis and Betterment of Organizations Through Their Members*. Lawrence Erlbaum 2003.
45. Starbird, K. and Palen, L. Working and sustaining the virtual "Disaster Desk". In *Proc. CSCW 2013*, ACM Press (2013).
46. Tyler, T. R., Rasinski, K. and Spodick, N. The influence of voice on satisfaction with leaders: Exploring the meaning of process control. *Journal of Applied Social Psychology*, 15 (1985), 700-724.
47. Wattenberg, M. and Viegas, F. B. The Word Tree, an Interactive Visual Concordance. *IEEE Transactions on Visualization and Computer Graphics*, 14, 6 (2008), 1221-1228.
48. Yang, J., Adamic, L., Ackerman, M., Wen, Z. and Lin, C.-Y. The way I talk to you: Sentiment expression in an organizational context. In *Proc. CHI 2012*, ACM Press (2012).