emotionVis: Designing an Emotion Text Inference Tool for Visual Analytics

Chris Zimmerman(✉), Mari-Klara Stein, Daniel Hardt, Christian Danielsen, and Ravi Vatrapu

Department of IT Management, Copenhagen Business School, Frederiksberg, Denmark
{cz.itm,mst.itm,dh.itm,cdfd.itm,rv.itm}@cbs.dk

Abstract. With increasingly high volumes of conversations across social media, the rapid detection of emotions is of significant strategic value to industry practitioners. Summarizing large volumes of text with computational linguistics and visual analytics allows for several new possibilities from general trend detection to specific applications in marketing practice, such as monitoring product launches, campaigns and public relations milestones. After collecting 1.6 million user-tagged feelings from 12 million online posts that mention emotions, we utilized machine learning techniques towards building an automatic 'feelings meter'; a tool for both researchers and practitioners to automatically detect emotional dimensions from text. Following several iterations, the test version has now taken shape as emotionVis, a dashboard prototype for inferring emotions from text while presenting the results for visual analysis.

1 Introduction

Emotion has become a popular topic in information systems research [1–6] as affective relationships are increasingly recognized as central to technology mediated interactions in general [7] and social networks in particular [8]. Emotions as "projections/displays of feeling" [9, 10] are the most visible layer of affect – explicitly shared in facial expressions, verbal and written communication. People go to social media in search of intensity, sensations and impressions that create affective jolts [8]. The resulting strong feelings are then often publicly shared (e.g., people voluntarily telling others 'how they are feeling' on Facebook). Due to this search for intensity, social media often invite intentionally provocative content and trolling behavior, rather than rational argumentation in a Habermasian public sphere [11].

While the uses and gratifications of social media include affective experiences, organizations seek to understand how their customers feel and which emotions they express online in relation to their brand. Such insights are used for improvements in marketing strategy, customer service, and the discovery of new business opportunities [12–15]. Opinions from the crowd are increasingly discovered with the help of automatic sentiment analysis. Yet state-of-the-art tools offer sentiment detection in the form of a polarity score, lacking context to what form of positivity or negativity was expressed. These results can potentially be misleading for practitioners [16]. Many research prototypes have overlooked fundamentals of information visualization by their own admission [17]; others have only attempted to pursue one angle of analysis [18]. Few research prototypes have

J. Parsons et al. (Eds.): DESRIST 2016, LNCS 9661, pp. 238–244, 2016.
DOI: 10.1007/978-3-319-39294-3_22

attempted multi-category emotion detection [19, 20]. If attempted, studies typically re-apply existing emotion classes from offline research (Ekman-6, Plutchik-8 or Scherer's emotion wheel) while leveraging generalized emotion lexicons (ANEW, OlympLex) that are not necessarily adept to online spaces. The dashboard prototype demonstrated herein differs in three fundamental ways from other existing tools. First, our training data set benefits from *specificity;* providing granularity of up to 143 discrete emotion types. Second, the tool offers *multi-dimensionality* by including both the valence and the arousal dimensions. Third, we offer a re-alignment in classification towards how people exhibit *emotions online*. Our tool reflects current emotion discourse online (particularly Facebook) which appears to be skewed towards high arousal and positive emotions (e.g., excitement) [15].

The artifact instantiates Action Design Research design principles from an ongoing engaged scholarship within a social media marketing agency [21]. Development cycles have been informed by research prototypes [22–24] and considers existing commercial tools such as Talkwalker, Topsy and Radian6 [25–27]. The emotionVis tool is accessible at cssl.cbs.dk/software/emotionVis (includes demo video).

2 Methodology: Artifact Design

Development began with data collection of 12 million public posts from Facebook that mentioned feelings. Of these, we used 1,618,499 posts, where the user tagged their text with one of 143 Facebook 'feelings tags' (feature since 2013) (Fig. 1). This provided us with a unique training set for our tool to detect emotion attributes from text. Standard approaches to text classification [28] involve the following steps: (1) Preprocessing: text is tokenized (so that words can be separately identified); (2) Feature Extraction: identifying word (unigrams, bigrams and trigrams) sequences; (3) Classification: a supervised machine learning algorithm is selected, which determines which combinations of features best predict the classification of interest. We followed the above steps to train four separate classifiers on our set of emotion-categorized posts.

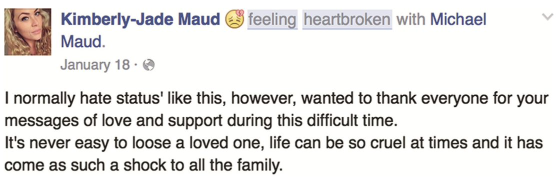

Fig. 1. Example of Facebook feeling tag used in training.

The first classifier detects individual emotions from the inputted text, leveraging 28 'Facebook feelings' with the most volume within our training set. The second classifier groups these feelings into 6 core emotions. It is common in emotions research to group discrete emotions into a smaller number of 'core' categories, such as joy, anger, sadness, fear and excitement [9, 29]. Our tool includes Joy, Sadness, Anger, Fear, Excitement and Empowerment, reflecting a wider range of high arousal, positive emotions – in line with the kinds of intense affective experiences

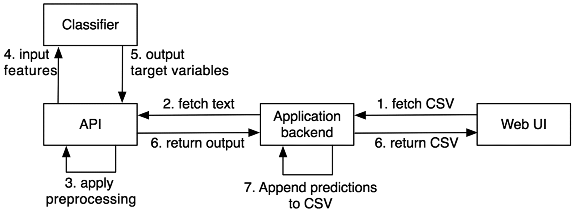

Fig. 2. Schematic diagram of the prototype.

people seek from and express on social media [8]. The third and fourth classifiers detect levels of valence (pos-neg sentiment) and levels of arousal from the inputted text. The resulting prototype consists of a backend and frontend (Fig. 2). The backend includes two Python Flask applications: an API interacting with the classifier and a web app interacting with the API. This app receives a CSV file from the user, extracts necessary data, sends it to the emotionVis API, adds classification scores to the CSV file, and returns the data to the user, along with computing the chart data for the dashboards. The front-end (UI) is made in HTML5, CSS and JavaScript. Bootstrap is used for the layout, while the D3.js and NVD3 visualization libraries are leveraged for the charts.

Action Design Research (ADR): The primary design consideration was to build a tool with value to Marketing and Social Media practitioners. ADR was used as a practical way of eliciting needs from the industry to inform the design [21]. Current development is taking place within a research environment while simultaneously testing the tool within a social media marketing agency. A survey was conducted among 30 practitioners at the agency to gauge which of the 143 feeling types were more relevant when making decisions. Early test versions of the tool were utilized by practitioners on a brand campaign by the electronics manufacturer Bang and Olufsen in 2015. In line with Sein et al.'s dual mission of making theoretical contributions while solving practitioner problems [4], we also see our tool as being useful for researchers who use similar datasets (post-level csv) from social channels (Facebook, Twitter, etc.). Design of the interface, thus, had these various users in mind.

Dashboard Affordances: The tool is intended to augment social media data that researchers or practitioners have at hand (e.g., post-level data exported from commercial tools as .csv files). Such data often provide details about the "who" and "when", which can be enhanced with "what feelings were felt" and topical clues as to why. The dashboard seeks to facilitate several affordances:

Emotional Alignment – A social media manager can compare the emotionality of postings made by the organization (e.g., happy or excited) with the emotionality of the comment chain generated in reaction to the published post.

Emotional Reverberation, Resilience, and Shifting – As emotions circulate on social media, they reverberate and intensify (e.g., a firestorm) or diffuse and wane. In downturns, brands may seek to rebound from negative emotional discourse. Sometimes the same objects can also elicit different emotions in different situations [30]. Our tool enables the tracing of 'emotional trajectories' within zoom-able time series/area charts [20]. In addition, as affect circulates, it creates an archive of expressed emotions [31]. Our tool affords tracing the *accumulation of emotions* over time.

Emotional Stickiness – As affect circulates, particular feelings can get 'stuck' to certain bodies, spaces, situations, etc. [8]. E.g., in a case of Marius the Giraffe, negative feelings got "stuck" to the scientific director of Copenhagen Zoo, who was the public face of the giraffe's controversial execution [32]. Term frequencies within emotion categories can unveil changing attachments between emotions and topics within text.

Design of the Prototype: The user begins by importing their own data (Fig. 3). Users are free to choose and indicate header names and the form of text file they are uploading (1). The layout design facilitates thematic progression in data exploration. Users first glance at the overall distributions of core emotions, followed by the breakdown of these groups into discrete emotions, and ultimately exploring individual posts and actors. Chart sophistication progresses at the same time (from bar charts to dual-axis to sunburst). As with other emotion-based tools [20], the color scheme has been carefully considered, while respecting existing connotations that humans identify with color. Two forms of bar chart show the **distribution of emotions** (2–3). A horizontal stacked bar chart shows a part-to-whole composition of core emotions. The ordering of emotions in this band (from most positive to negative) corresponds to the order of 28 sub-level emotions in the bar chart immediately below. The next layers illustrate **volumes over time** (4–5). Conversation volume is represented by gray bars in the background. This is overlaid with two lines that represent arousal and valence levels. A time slider encourages users to drill-down to specific windows in time.

A view displaying the **breakdown of a conversation (by core emotions) over time** follows. Individual emotions can be removed to isolate a stream of interest over time. A **post-level visualization** maps all posts within a conversation onto a two-dimensional (arousal and valence) scatterplot (6). This allows the user to see the emotional 'footprint' of the entire conversation. The practitioner can also see where their own published posts lie within the footprint – in comparison to the crowd – as visual feedback on emotional alignment. The sunburst diagram facilitates a **breakdown of core emotions** (7). One can see, for example, that a large portion of anger detected may have originated from annoyance, rather than disgust. This granularity elaborates on the unique footprint a conversation may hold. The last series of charts **rank people** in the conversation who express the highest average levels of each core emotion (8). These serve as actionable opportunities for social media managers to engage with people flagged as having particular emotions towards the brand.

3 Limitations and Future Perspectives

With the latest iteration of this prototype, opportunities for action research are expected to widen as practitioners begin to use the tool. EmotionVis has been tasked for reporting within a social media audit of a global hotel chain, as well as conversation monitoring

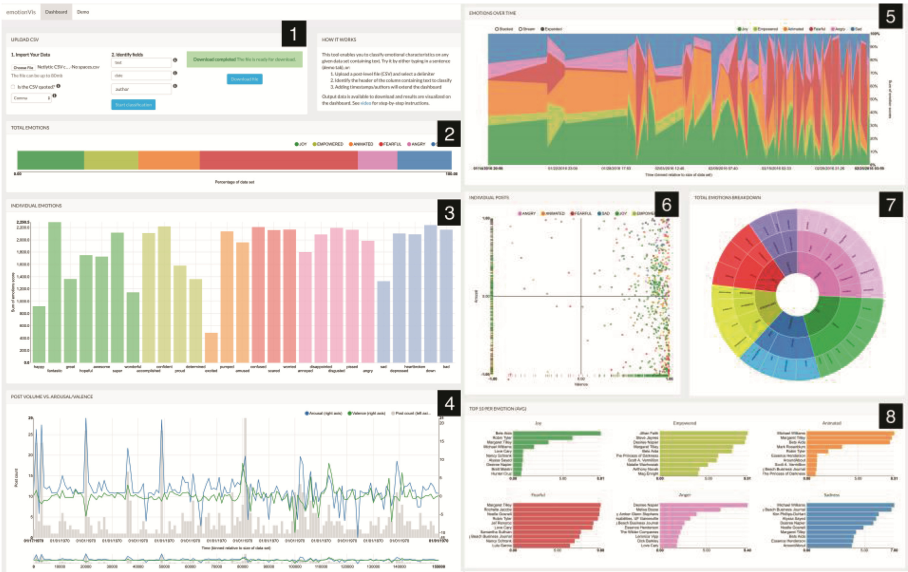

Fig. 3. The User Interface (UI) presents data for different affordances via specific visualizations (Color figure online).

during the European football championships (Euro 2016) by the tournament's main sponsor. Such applications will provide valuable future direction in the development of the prototype. Once the tool is directly connected to APIs from social media channels, it will also provide a real-time interface benefiting from consistent display in the dashboard. Currently the visualizations are also serving designers as self-assessment instruments to fine-tune the classifier. In future work, we will systematically assess the accuracy of detecting different emotions in test data.

References

1. Beaudry, A., Pinsonneault, A.: The other side of acceptance: studying the direct and indirect effects of emotions on information technology use. MIS Q. **34**, 689–710 (2010)
2. McGrath, K.: Affection not affliction: the role of emotions in information systems and organizational change. Inf. Organ. **16**, 277–303 (2006)
3. De Guinea, A., Markus, M.: Why break the habit of a lifetime? Rethinking the roles of intention, habit, and emotion in continuing information technology use. Mis Q. **33**, 433–444 (2009)
4. Stein, M., Newell, S., Wagner, E., Galliers, R.: Felt quality of sociomaterial relations: introducing emotions into sociomaterial theorizing. Inf. Organ. **24**, 156–175 (2014)
5. Stein, M., Newell, S., Wagner, E., Galliers, R.: Coping with information technology: mixed emotions, vacillation, and nonconforming use patterns. Mis Q. **39**, 367–392 (2015)
6. Zhang, L., Jiang, M., Farid, D., Hossain, M.: Intelligent facial emotion recognition and semantic-based topic detection for a humanoid robot. Expert Syst. Appl. **40**, 5160–5168 (2013)

7. Vatrapu, R.K., Suthers, D.D.: Technological intersubjectivity in computer supported intercultural collaboration. In: Proceedings of the 2009 International Workshop on Intercultural Collaboration - IWIC 2009, pp. 155–164 (2009)
8. Hillis, K., Paasonen, S., Petit, M.: Networked Affect. MIT Press, Cambridge (2015)
9. Ekman, P., Friesen, W.: Universals and cultural differences in the judgments of facial expressions of emotion. J. Pers. Soc. Psychol. **53**, 712 (1987)
10. Shouse, E.: Feeling, emotion, affect. M/C Journal **8**(8), 26 (2005)
11. Robertson, S.P., Vatrapu, R.K., Medina, R.: Off the wall political discourse: facebook use in the 2008 U.S. presidential election. Inf. Polity **15**(1–2), 11–31 (2010)
12. Goh, K., Heng, C., Lin, Z.: Social media brand community and consumer behavior: quantifying the relative impact of user-and marketer-generated content. Inf. Syst. Res. **24**, 88–107 (2013)
13. Holsapple, C., Hsiao, S., Pakath, R.: Business social media analytics: definition, benefits, and challenges. Bus. Soc. Media Anal. **2010**, 1–12 (2014)
14. Kurniawati, K., Shanks, G., Bekmamedova, N.: The business impact of social media analytics. ECIS **13**, 13 (2013)
15. Zimmerman, C., Stein, M., Hardt, D., Vatrapu, R.: Emergence of things felt: harnessing the semantic space of facebook feeling tags (2015)
16. Marcus, A., Bernstein, M., Badar, O.: Twitinfo: aggregating and visualizing microblogs for event exploration. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pp. 227–236. ACM (2011)
17. Guzman, E., Bruegge, B.: Towards emotional awareness in software development teams. In: Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering, pp. 671–674. ACM (2013)
18. Wensel, A., Sood, S.: Vibes: visualizing changing emotional states in personal stories. In: Proceedings of the 2nd ACM International Workshop on Story Representation, Mechanism and Context, pp. 49–56. ACM (2008)
19. Pu, P.: EmotionWatch: visualizing fine-grained emotions in event-related tweets. In: Proceedings of the in the 8th International AAAI Conference on Weblogs and Social Media (2014)
20. Zhao, J., Gou, L., Wang, F., Zhou, M.: PEARL: an interactive visual analytic tool for understanding personal emotion style derived from social media. In: 2014 IEEE Conference on Visual Analytics Science and Technology, pp. 203–212. IEEE (2014)
21. Sein, M., Henfridsson, O., Purao, S., Rossi, M., Lindgren, R.: Action design research. MIS Q. **35**, 37–56 (2011)
22. Hussain, A., Vatrapu, R.: Design, development, and evaluation of the social data analytics tool. In: Tremblay, M.C., VanderMeer, D., Rothenberger, M., Gupta, A., Yoon, V. (eds.) DESRIST 2014. LNCS, vol. 8463, pp. 368–372. Springer, Heidelberg (2014)
23. Zimmerman, C., Vatrapu, R.: Designing the social newsroom: visual analytics for social business intelligence. In: Donnellan, B., Helfert, M., Kenneally, J., VanderMeer, D., Rothenberger, M., Winter, R. (eds.). LNCS, vol. 9073, pp. 386–390Springer, Heidelberg (2015)
24. Flesch, B., Vatrapu, R.: Social set visualizer: a set theoretical approach to big social data analytics of real-world events. In: 2015 IEEE International Conference on Big Data (Big Data), pp. 2418–2427. IEEE (2015)
25. Trendiction, Talkwalker. www.talkwalker.com. Accessed 13 April 2015
26. Apple (purchased and closed in 2014), Topsy. www.topsy.com/analytics
27. Salesforce, Radian6. www.radian6.com. Accessed 13 April 2015

28. Pang, B., Lee, L., Vaithyanathan, S.: Thumbs up? Sentiment classification using machine learning techniques. In: Proceedings Conference on Empirical Methods in Natural Language Processing 6–7 July 2002, Philadephia, Pennsylvania, USA, pp. 79–86 (2002)
29. Parrott, W.: Emotions in Social Psychology: Essential Readings. Psychology Press, Abingdon (2001)
30. Kuntsman, A.: Introduction: affective fabrics of digital cultures. In: Karatzogianni, A., Kuntsman, A. (eds.) Digital Cultures and the Politics of Emotion, pp. 1–17. Palgrave Macmillan, Basingstoke (2012)
31. Cvetkovich, A.: An Archive of Feelings: Trauma, Sexuality, and Lesbian Public Cultures. Duke University Press, Durham (2003)
32. Zimmerman, C., Chen, Y., Hardt, D., Vatrapu, R.: Marius, the giraffe: a comparative informatics case study of linguistic features of the social media discourse. In: Proceedings of the 5th ACM International Conference on Collaboration Across Boundaries: Culture, Distance and Technology, pp. 131–140. ACM (2014)