Changes in emotional content of forum messages precede user action on subject

Casey Doyle[1][0000-0002-0885-4954] and Daniel Krofcheck, Matthew Sweitzer, Asmeret Naugle, Asael Sorensen

Sandia National Laboratories, Albuquerque, New Mexico, USA
[1]cldoyle@sandia.gov

Abstract. While most online discussion remains confined to textual conversation, occasionally those participating escalate further and take action. This escalation indicates a change in the level of engagement from those that simply observe or discuss, and is a valuable transition to identify. We study emotional signals within textual conversations before, during, and after actions in two different contexts (at an individual level in an online game and in aggregate in a stock trading forum) and show that action is often preceded by a change in emotional context of the conversation. While we find only small changes in quantity of emotional content, the types of emotions expressed in emotional posts fluctuates greatly before action occurs, then can revert in the opposite direction once the period of action is over. Anger in particular exhibits large, fast swings during periods of action state transition in both datasets tested. These changes in emotional context relative to taking action provides an indicator of the altered mental state and activation that precedes action, evidencing the link between emotion and action in online forums.

Keywords: emotion, emotion classification, actions, action readiness, natural language processing, data mining, social media

1 Introduction

The mental process of making decisions and taking actions is complex, and a major area of study in psychology, philosophy, and cognition. Humans are constantly taking in new information and making decisions based on their current mental state. With the increasing prevalence of online interaction, however, information quantity has exploded and its effects on mental states and behaviors is still being explored. Behavior online is not always the same as offline, and the stimulus experienced can vary wildly. While studies suggest a link between offline personality traits and online behavior [12], there are also indications that online engagement does not always correlate to offline behavior [3]. Additionally, the way people treat social capital online and offline is distinct [29] which may lead to different risk calculations and actions. In order to better understand the dynamics of online behavior, we must ask which offline behavior models hold, and how the interaction between online and offline processes intersects. One key facet of this is understanding how individuals change from being an observer or participant in a conversation to being an actor, even if the relevant actions also occur online. Such a transition marks a concrete change in the level of engagement, and tracking the connection between discussion and action can illuminate some of the mental processes inherent to social participation.

In this paper, we investigate this transition from the perspective of emotional engagement in a subject. The connection between emotions and action readiness in psychology and cognitive science has been a subject of interest for decades [9, 24, 28]. Many studies present evidence for causal relationships between emotional state and actions [20] as well as explanatory links highlighting the deep intertwining of the two [7, 18]. By examining the link between online emotional expressions and action, we seek to provide insight into quantifiable action readiness and engagement escalation in this domain, extending prior studies in the area. To help quantify the effects studied, we present a data-centric view of online behavior, studying emotional signals within online posts correlated with actions taken outside of the discussion. Despite challenges presented by online anonymity, such analysis of action in social media is not an uncommon study. Most commonly, these occur as in-domain studies, where the challenges of anonymity can be overcome by applying machine learning approaches to forecasting of behaviors within a single social media outlet [19, 26]. By utilizing a single comprehensive dataset, information on interaction and action is available for each user in such studies [5, 22]. The models produced in these studies often serve to predict further interactions, social network evolution, and non-text actions such as liking or following other users to great effect. The single data sources and fully trained models produced, however, offer limited insight into the mental dynamics surrounding action and their relation to existing understanding of cognitive engagement. Other studies seek to extend action prediction to cross-domain action by correlating aggregate discussion features to independent events taking a generalized view of how conversational changes precede action [10]. While the challenges related to the lack of direct user-to-action information in this sort of analysis are obvious, by looking at large scale events such as protests [27], pandemic compliance [23], and political action [15] efforts have succeeded in predicting real-world events through models trained on preceding discussion. The aggregated data and predictive modeling utilized in these studies, however, can cause a lack of explainability that limits the ability to connect the conclusions to more traditional cognitive studies.

We seek to augment these prior predictive efforts by bridging the data science perspective of action prediction and the cognitive perspective of action motivation. By studying the specific emotional context of messages surrounding times of action in online communities, we seek to increase explainability of predictions by grounding them in connections to understood cognitive processes. Mechanisms of action priming and intention are core cognitive concepts that offer a wealth of insight into how individuals interact with their environments, and identification of whether the emotional signals contained in discussion can inform on mental state and readiness to engage further provides a new tool to understand how discourse and engagement evolve.

2 Data and Methods

Advances in natural language processing have allowed for great leaps in understanding of how to extract emotional signals out of unlabeled text [25], including characterizing temporal dynamics of emotional content [17]. Modern transformer architecture and large language models have taken this capability even further, resulting in accurate emotional classification models across many use cases [1]. We use one such model ('Emotion English DistilRoBERTa-base') [14] for its ability to classify accurately across datasets without custom dictionaries and annotations, and RoBERTa is among the highest performing and validated models at such tasks [2]. Additionally, the model provides quantification of well-understood emotional vectors within the text, based on Ekman's six basic emotions categorization [8]. As the data basis for this analysis, we compare posts and actions first from an online forum and corresponding MMORPG (*GameX*), then from an economic Reddit forum and corresponding stock market trade volume (wallstreetbets and gamestop).

First, we study forum posts by users in an online game (referred to here as *GameX*) [16] and related actions made by the posters in the game. *GameX* is an online roleplaying game with rich player interaction including trading, fighting, and resource management. The game has a strong economic system in it, with trading being a central part of the game play mechanics. We focus our analysis on these trading actions and relevant trade-topic posts in related forums. This data is uniquely well suited to the study at hand because while the forum centers around the game, they are distinct environments yet are connected both in subject and by user(posts made by players are connected to actions via linked accounts). We process the data using topic modeling to ensure subject relevance, then temporal binning based on proximity to relevant actions. Specifically, we use BERTopic [13] to select only posts with topics related to trade based on a corpus of words surrounding in-game trading mechanics. We then bin them based on whether the user that posted has traded (1) the day before (but not after), (2) the day after (but not before), (3) both the day before *and* the day after, (4) did not trade within a day of the post, or (5) never undertook a trade action. Then, applying the JHartmann Roberta emotion classification model [14] to the comments we obtain scores for Ekman's six basic emotion vectors [8]. Finally, for each post within a temporal category, we compare the emotional content of posts within the category to all those outside of the category for each user.

To extend our study further into cross-domain applications, we also perform an aggregate study on the GameStop short squeeze of 2021, in which posters on Reddit were highly involved in the stock dynamics [4]. While the anonymous nature of Reddit does not allow a one-to-one comparison of user posts and actions, this exemplar does provide a quantifiable exogenous signal (the volume of shares traded) that can be compared to the emotionality of posts, making it a

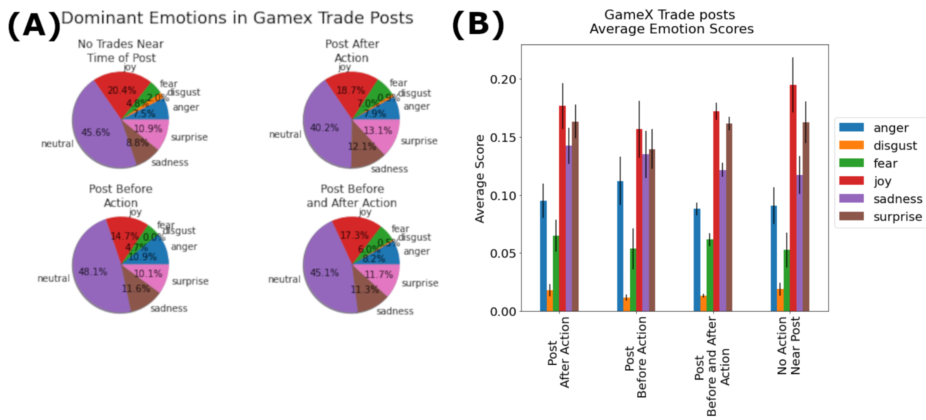

Fig. 1. Results of emotional analysis of posts in the GameX forums relative to trade actions taken by the poster. (A) The dominant (highest scored) emotions represented in the temporal category. (B) The average scores for each non-neutral emotion in each temporal category among.

useful complement *GameX* data. To make the study as similar as possible, we repeat many of the same steps as above. We scrape data from the wallstreetbets subreddit during the short squeeze, apply a topic model to filter relevance, then bin based on proximity to events. In this case, the events consist of the days on which the GameStop stock had a day-over-day difference in trading volume (buying and selling) in the 99^{th} quantile of all days between September of 2019 and April of 2023. This yields the days of January 13, 14, 22, 25, 26, 27 and February 24, 25, 26 of 2021. The temporal bins here align to those in the *GameX* study, with posts being (1) the day after a spike, (2) day before a spike, (3) on the final day of a mult-iday spike, (4) in the middle but before the end of a multi-day spike, or (5) not adjacent to a spike. As before the emotional classification model is used to compute emotional vectors, and the in-category means are compared to the out-of-category means to quantify the differences in emotional content in relation to the volume trade spikes.

3 Results

First, we take the case of *GameX*. After filtering comments for those relevant to trade and then binning them based on proximity to action, we can describe the emotional landscape via the dominant emotion (defined as the emotional category with the highest score on any given comment) in each category. Figure 1(A) shows that there is little relation between the quantity of emotional vs neutral posts and action proximity. There is a slight increase in emotionally dominant posts shortly after trades, but it is small relative to the overall proportion of posts that remain neutral. There is a noticeable effect of joy decreasing around action, and indeed being at its lowest point *before* action is taken, rebounding

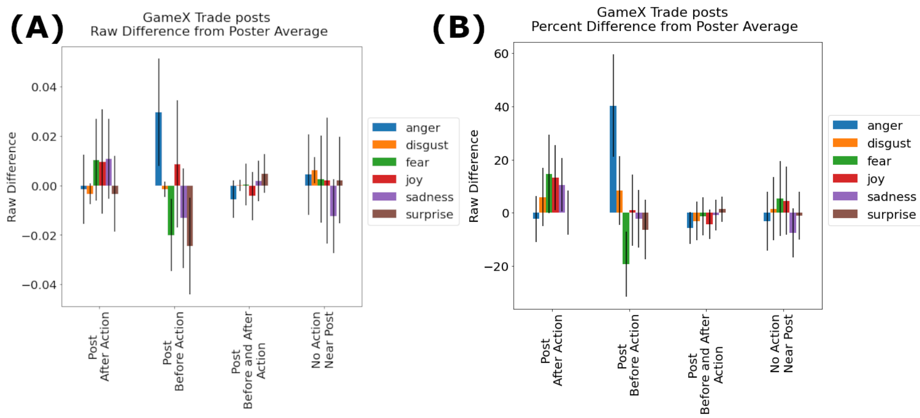

Fig. 2. Relative non-neutral emotional scores based on analysis of posts in the GameX forums relative to actions taken by the poster. (A) The raw difference in scores between the poster during that temporal category and the average post from that user. (B) The percent difference relative to the average score of posts from that user.

after actions are completed. Similarly, we see an increase in anger in posts that occur directly before trade actions, which also rebounds after the action is completed. The average emotional scores of each category shown in Figure 1(B) tells a similar story; joy is highest when removed from action and lowest right before action, while anger is highest before action. The general culture of the forums can be seen as well, with joy, surprise, and sadness dominating in general, only rivaled by anger right before trade actions are performed.

This provides the first evidence that emotional content is different prior to action than it is at other times, and the observable rebounding of emotions after actions also aligns with bidirectional theories of emotion and action [18]. Before people act, they undergo an emotional shift that aligns with the priming for action; they have changed states and are ready to engage further. After action, the outcome of their action, as well as generalized emotional feedback from acting, affects emotional state and discussions on the matter before regression to the baseline emotional state. Indeed, here emotional content is most different from the baseline *before* action, while the emotional signatures in between actions and after actions are well aligned with a regression back to the no action case.

Further, by dividing up posts by poster, we can better understand individual emotional variations. By creating averages specific to each poster and defining the temporal deviation from those individual averages, we can describe behavioral changes in a more focused way and present them as the average change from each individual's baseline behavior, shown in Figure 2. Overall, these results reflect those shown in Figure 1, although with a more striking difference between the 'post before action' category and all other categories. The most obvious result is that posts prior to taking a trade action are *much* angrier, but there is also a stark drop in fear before trading and a decrease in surprise. A

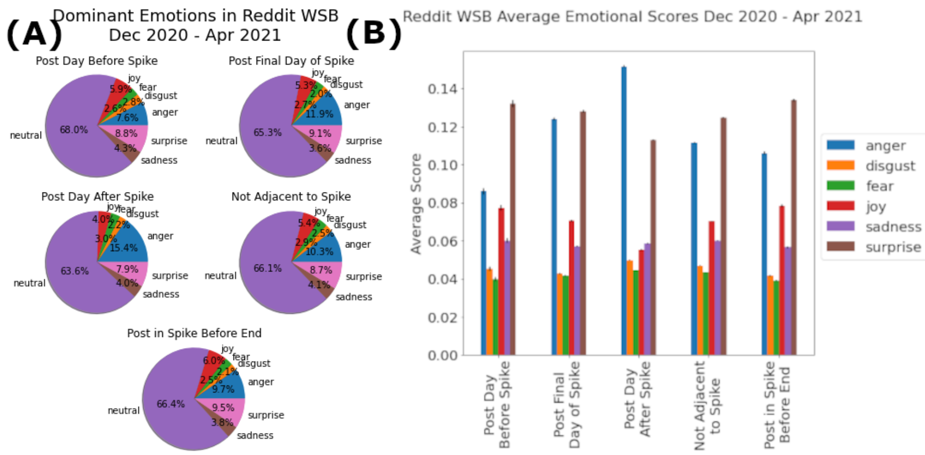

Fig. 3. Results of emotional analysis of posts in the Reddit Wallstreetbets forum during the period of December 2020 to April 2021 relative to spikes in the volume traded of the GameStop stock. (A) The dominant (highest scored) emotions represented in the temporal category. (B) The average scores for each non-neutral emotion in each temporal category.

similar post-action rebounding effect is apparent, as fear increases well beyond the average after trading, before settling back close to the average fear during no trade actions. Similarly, a disproportionate increase in joy can be seen post-action compared to other categories.

Continuing this analysis with the Reddit Wallstreetbets dataset, some overall trends remain despite a clear difference in setting and culture. Figure 3(A) shows that posts on Reddit are on average more neutral than those on the *GameX* forums, and we see very little difference in the proportion of neutrally dominant posts regardless of time category. Interestingly, the small effect observed does echo that of the *GameX* forums; posts are most neutral the day before periods of action (here defined by trade volume spikes) and most emotional the day after periods of action. Again the baseline (not being adjacent to a spike) is in between the two, so there is a repeated though small effect of lower emotional content before periods of action. Within the non-neutral dominant emotions, the clear observable effect is the decrease in anger before spikes and increase after spikes. Interestingly, anger was also the primary emotion shift in the *GameX* forums, but in the opposite direction. Figure 3(B) echoes this, showing that anger and surprise dominate non-neutral emotional content, with joy only rivaling anger the day before spikes. Anger sees large shifts based on temporal category, followed by surprise which is at its minimum the day after spikes and maximum before and during spikes. This emotional change observed (and the general emotional culture of the boards) seem to be opposed; in the GameX data joy is the primary dominant emotion and actions are preceded by a spike in anger, while in the

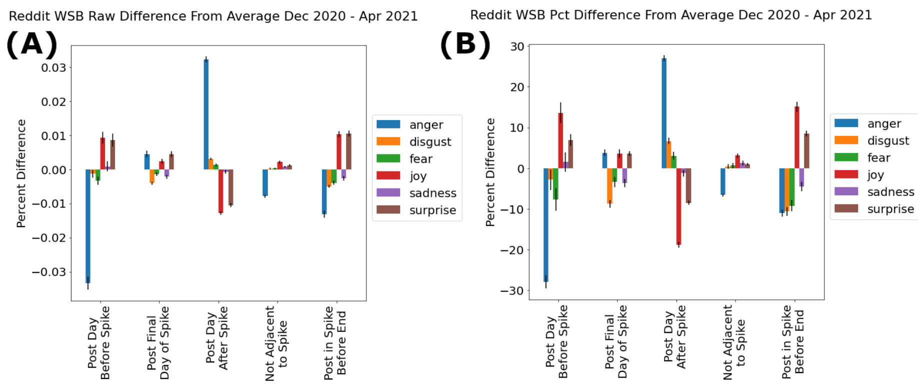

Fig. 4. Relative non-neutral emotional scores of posts in the Reddit Wallstreetbets forum during the period of December 2020 to April 2021 relative to spikes in the volume traded of the GameStop stock. (A) The raw difference in scores between the poster during that temporal category and the average. (B) The percent difference relative to the average score.

Reddit data anger is the dominant emotion and action is preceded by a drop in anger and spike in joy.

To echo the analysis performed above, we analyze the Reddit data via the change in emotional scores by time category from average in Figure 4. As before anger is the primary difference between timing categories, dominated by severe drops in anger the day before events. This drop in anger is accompanied by increases in joy and surprise, which continues throughout the high trading activity period. During multi day trading spikes, anger is still lower than normal but much higher than the day before trading, while the increase in joy and surprise grow and other negative emotions such as fear and disgust begin to drop. On the final day of trading spikes (including single-day spikes), we start to see the beginning of the rebound process. Anger is already higher than average at this point, and while the effects on joy, surprise, fear, and disgust are still there the effects have largely shrunken back to near baseline. Then, on the day after the spike, the effects are reversed. Anger spikes to extreme levels 30% higher than normal, while joy and surprise decrease rapidly. Fear and disgust also reverse, and are slightly higher than average. Overall, many of the effects observed on the individual level are echoed, but with an increased emphasis on the full emotional cycle. We can see a clear trend where *before* action takes place emotions enter a state of flux, changing to a new state that contains a consistent signature throughout the period of activity, then begins to reverse and reaches a rebound state the day after activity. Additionally, we see indications that different emotions operate on different timescales, as anger is fast to change (reaching its minimum before action begins then receding back closer to the mean during action) while other emotions are slower moving (disgust minimizing and joy maximizing during action). These slower moving emotions also remain in flux longer, not peaking in their rebound until the day after action.

4 Discussion and Conclusion

Action readiness is determined by cognitive state, restricted by ability and desire to act. Understanding the mental state of action readiness is a longstanding field of study, and is closely tied to emotions and their role in motivating or explaining imminent behavior. While these connections are well known in the context of general action, they not been fully explored in the context of online interactions and personas. We utilize modern emotional classification techniques that allow for the identification of emotions across a broad range of environments to analyze the connection between emotions and actions in two different online datasets. We show that on an individual level, emotional posts in a forum about an online game exhibit large changes in emotional content right before action, and then reverse before returning to baseline. This effect is shown even more clearly in the second dataset, a Reddit economic forum engaging in the GameStop short squeeze of 2021, where again there are large changes in emotional content the day before trading activity spikes, followed by a reversal and return to baseline. In particular, in both datasets it is *anger* that shows the most volatility related to action and changes the fastest, although the direction of the anger fluctuations are opposite in the two datasets (in the game data anger spikes before action, while in the Reddit data anger drops then spikes after), a difference that is possibly tied to board culture and anger dominance or related to the context of the discussion and action. Indeed, while prior work related anger to video game activity largely focuses on general correlations between anger and violent video games [6], emotions including anger have shown motivational effects for people to continue playing games [21]. From the perspective of stock trading, positive emotions have been shown to relate to anomalies in stock trading and underestimation of risk [11]. Despite the apparent specificity of the directionality of change, our work indicates that the emotional transition process for individuals undertaking actions is observable across platforms as a generalized fluctuation in expressed emotion. This supports the concept of emotion as a core component of action readiness, and indicates a strong signal for inclusion in future predictive modeling. Further work is warranted on the role of activity specification and domain culture as while an observable change is consistent across datasets the directionality varies. Additionally, further study into higher resolution timeseries of emotional content is warranted. In all, the presence and observable connection between emotions and actions in online settings demonstrates a link that has yet to be fully explored, while modern techniques in emotion detection and classification have primed the field for expansion in this direction that could greatly benefit understanding of action prediction, mental states surrounding action, and the interaction between online and offline behavior.

Acknowledgments. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC

(NTESS), a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration (DOE/NNSA) under contract DE-NA0003525. This written work is authored by an employee of NTESS. The employee, not NTESS, owns the right, title and interest in and to the written work and is responsible for its contents. Any subjective views or opinions that might be expressed in the written work do not necessarily represent the views of the U.S. Government. The publisher acknowledges that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this written work or allow others to do so, for U.S. Government purposes. The DOE will provide public access to results of federally sponsored research in accordance with the DOE Public Access Plan.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Acheampong, F.A., Nunoo-Mensah, H., Chen, W.: Transformer models for text-based emotion detection: a review of bert-based approaches. Artificial Intelligence Review pp. 1–41 (2021)
2. Adoma, A.F., Henry, N.M., Chen, W.: Comparative analyses of bert, roberta, distilbert, and xlnet for text-based emotion recognition. In: 2020 17th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP). pp. 117–121. IEEE (2020)
3. Alhabash, S., McAlister, A.R., Lou, C., Hagerstrom, A.: From clicks to behaviors: The mediating effect of intentions to like, share, and comment on the relationship between message evaluations and offline behavioral intentions. Journal of Interactive Advertising **15**(2), 82–96 (2015)
4. Anand, A., Pathak, J.: The role of reddit in the gamestop short squeeze. Economics Letters **211**, 110249 (2022)
5. De Choudhury, M., Sundaram, H., John, A., Seligmann, D.D.: Social synchrony: Predicting mimicry of user actions in online social media. In: 2009 International conference on computational science and engineering. vol. 4, pp. 151–158. IEEE (2009)
6. Demirok, M., Ozdamli, F., Hursen, C., Ozcinar, Z., Kutguner, M., Uzunboylu, H.: The relationship of computer games and reported anger in young people. Australian Journal of Guidance and Counselling **22**(1), 33–43 (2012)
7. Döring, S.A.: Explaining action by emotion. The Philosophical Quarterly **53**(211), 214–230 (2003)
8. Ekman, P., et al.: Basic emotions. Handbook of cognition and emotion **98**(45-60), 16 (1999)
9. Frijda, N.H., Kuipers, P., Ter Schure, E.: Relations among emotion, appraisal, and emotional action readiness. Journal of personality and social psychology **57**(2), 212 (1989)
10. Glenski, M., Weninger, T.: Predicting user-interactions on reddit. In: Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017. pp. 609–612 (2017)
11. Goodell, J.W., Kumar, S., Rao, P., Verma, S.: Emotions and stock market anomalies: A systematic review. Journal of Behavioral and Experimental Finance **37**, 100722 (2023)

12. Gosling, S.D., Augustine, A.A., Vazire, S., Holtzman, N., Gaddis, S.: Manifestations of personality in online social networks: Self-reported facebook-related behaviors and observable profile information. Cyberpsychology, Behavior, and Social Networking **14**(9), 483–488 (2011)
13. Grootendorst, M.: Bertopic: Neural topic modeling with a class-based tf-idf procedure. arXiv preprint arXiv:2203.05794 (2022)
14. Hartmann, J.: Emotion english distilroberta-base. https://huggingface.co/j-hartmann/emotion-english-distilroberta-base/ (2022)
15. Korolov, R., Lu, D., Wang, J., Zhou, G., Bonial, C., Voss, C., Kaplan, L., Wallace, W., Han, J., Ji, H.: On predicting social unrest using social media. In: 2016 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM). pp. 89–95. IEEE (2016)
16. Lakkaraju, K., Whetzel, J.H., Lee, J., Bier, A.B., Cardona-Rivera, R.E., Bernstein, J.R.R.: Validating agent based models through virtual worlds. Tech. rep., Sandia National Lab.(SNL-NM), Albuquerque, NM (United States) (2014)
17. Liu, Z., Yang, C., Rüdian, S., Liu, S., Zhao, L., Wang, T.: Temporal emotion-aspect modeling for discovering what students are concerned about in online course forums. Interactive Learning Environments **27**(5-6), 598–627 (2019)
18. Lowe, R., Ziemke, T.: The feeling of action tendencies: on the emotional regulation of goal-directed behavior. Frontiers in psychology **2**, 346 (2011)
19. Phillips, L., Dowling, C., Shaffer, K., Hodas, N., Volkova, S.: Using social media to predict the future: a systematic literature review. arXiv preprint arXiv:1706.06134 (2017)
20. Scarantino, A.: Do emotions cause actions, and if so how? Emotion Review **9**(4), 326–334 (2017)
21. Schønau-Fog, H., Bjørner, T.: "sure, i would like to continue" a method for mapping the experience of engagement in video games. Bulletin of Science, Technology & Society **32**(5), 405–412 (2012)
22. Shrestha, P., Maharjan, S., Arendt, D., Volkova, S.: Learning from dynamic user interaction graphs to forecast diverse social behavior. In: Proceedings of the 28th ACM international conference on information and knowledge management. pp. 2033–2042 (2019)
23. Volkova, S., Ayton, E., Porterfield, K., Corley, C.D.: Forecasting influenza-like illness dynamics for military populations using neural networks and social media. PloS one **12**(12), e0188941 (2017)
24. Weiner, B.: Attribution, emotion, and action. Handbook of motivation and cognition: Foundations of social behavior p. 281–312 (1986)
25. Yadollahi, A., Shahraki, A.G., Zaiane, O.R.: Current state of text sentiment analysis from opinion to emotion mining. ACM Computing Surveys (CSUR) **50**(2), 1–33 (2017)
26. Yu, S., Kak, S.: A survey of prediction using social media. arXiv preprint arXiv:1203.1647 (2012)
27. Zhou, Y., Zhang, L., Liu, X., Zhang, Z., Bai, S., Zhu, T.: Predicting the trends of social events on chinese social media. Cyberpsychology, Behavior, and Social Networking **20**(9), 533–539 (2017)
28. Zhu, J., Thagard, P.: Emotion and action. Philosophical psychology **15**(1), 19–36 (2002)
29. de Zúñiga, H.G., Barnidge, M., Scherman, A.: Social media social capital, offline social capital, and citizenship: Exploring asymmetrical social capital effects. Political Communication **34**(1), 44–68 (2017)