A Lens to Pandemic Stay at Home Attitudes

Andrew Wentzel (iD), Lauren Levine, Vipul Dhariwal, Zahra Fatemi, Abari Bhattacharya (iD),
Barbara Di Eugenio (iD), Andrew Rojecki (iD), Elena Zheleva (iD), G.Elisabeta Marai (iD)
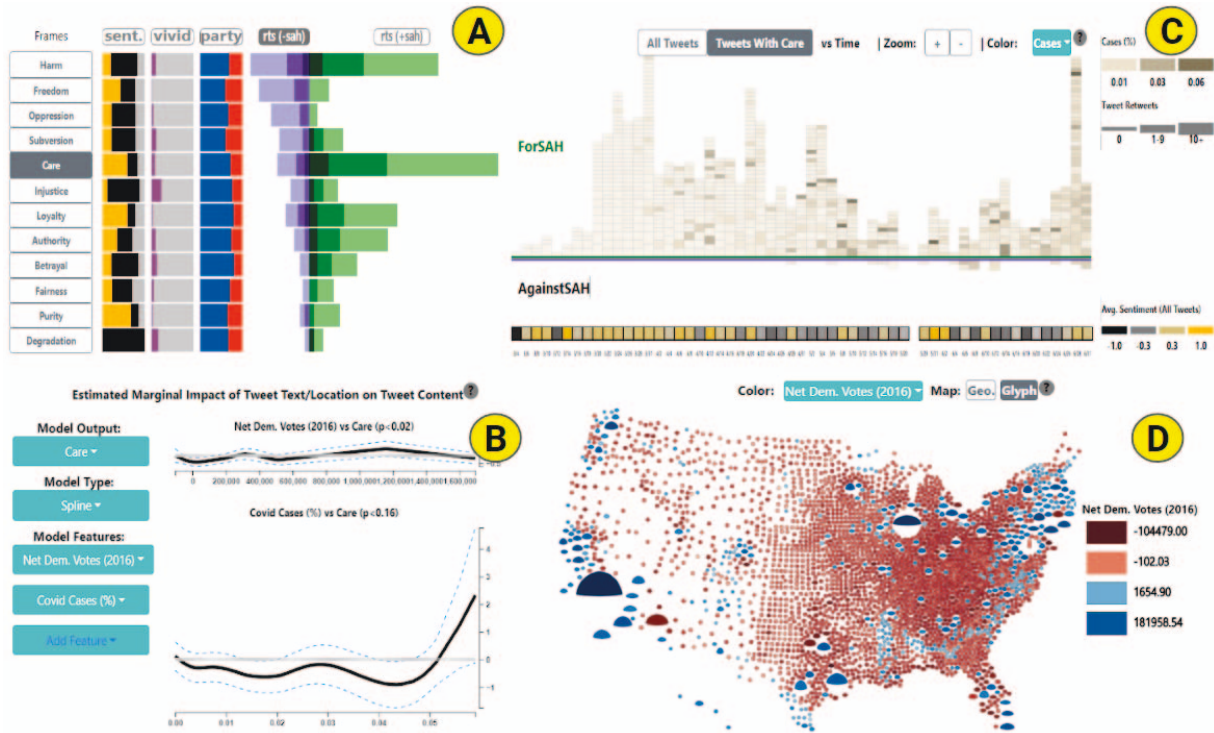
Figure 1: Analysis of moral frames (MF) in social media posts related to Stay-at-Home orders during the COVID-19 pandemic. (A) Tweet summary showing several MF tweet features, sorted here by tweets with a negative stance. (B) Inference panel showing partial dependence plots derives from generalized linear models (GAMS), here showing the relationship between county voting history and COVID-19 cases, and Care tweets. (C) Timeline of Care tweets, along with popularity, COVID-19 incidence, and sentiment. Note the negative spikes in April/June, after BLM protests, and the more negative sentiment (lower bar). A tooltip shows the text of a tweet supporting SAH orders. (D) Glyph-map of counties showing political party (color) and population (width) vs. Care tweets (height); major cities and rural white areas stand out.

ABSTRACT

We describe the design process and the challenges we met during a rapid multi-disciplinary pandemic project related to stay-at-home orders and social media moral frames. Unlike our typical design experience, we had to handle a steeper learning curve, emerging and continually changing datasets, as well as under-specified design requirements, persistent low visual literacy, and an extremely fast turnaround for new data ingestion, prototyping, testing and deployment. We describe the lessons learned through this experience.

1 INTRODUCTION

In early 2020, in light of the emergence of the global COVID-19 pandemic, the U.S. National Science Foundation (NSF) began accepting proposals for non-medical, non-clinical research that could be used right away to explore how to model and understand the spread of COVID-19, how to inform and educate about the science of virus transmission and prevention, and how to encourage the development of procedures and actions to address the pandemic [9]. NSF invited researchers to apply for funding through the Rapid Response Research (RAPID) mechanism. This mechanism enables the NSF to receive and review proposals that have a high priority in terms of the availability of or access to data, facilities, or specialized equipment, as well as quick-response research on natural or anthropogenic disasters and similar unforeseen events. Requests for RAPID proposals could be for up to $200K and up to one year in duration.

Our collaborative idea was for a RAPID project analyzing people's moral take on and consequent response to government mandated stay-at-home (SAH) orders at the start of the pandemic, leading to potential insights into effective governmental messaging of such orders. Shared values, beliefs, and understandings were often reflected at the time in social media posts, therefore, we were interested in identifying the underlying values expressed in such posts, and how they related to attitudes with respect to SAH messaging. Concretely, the funded project aimed to analyze these values and attitudes using Moral Foundations (MF) theory, a psychological model for describing the different dimensions underlying social discourse, such as Care or Liberty or Loyalty. Building from the strengths of the team, the project focused on applying MF theory, in collaboration with social-science researchers, Causal Inference (CI)

researchers, and Natural Language Processing (NLP) researchers to help build and analyze a MF-annotated tweet corpus related to SAH orders. Additionally, the project aimed to leverage geospatial data to analyze datasets at multiple levels of spatial aggregation, and to compare temporal and spatial differences to enhance participation and promote positive public health outcomes.

Analyzing social media from a moral-framing perspective posed a number of challenges. First, while collaborators had clear insights into MF theory, applying these frames to social media was a difficult and poorly defined task. an example, a tweet stating "We are not staying home!" could be an expression of the "Liberty" frame in opposition to government lockdowns if it was tweeted at the end of March 2020. However, if the tweet occurred two months later during a major civil movement protest, it could be a tweet about solidarity with other protesters and thus would be an expression of the "Loyalty" frame. Therefore, meaningful analysis of the data requires understanding the tweeter's situational context, such as time, location, and major recent events. Second, since developing a corpus to analyze was part of the proposed work, design requirements had to constantly be updated as the expected scope of the data and resulting expectations changed week-to-week. We also found that annotating tweets with moral frames is a difficult task that requires trained experts. These issues resulted in rapid changes in the expected number of tweets that were both annotated and geotagged, the ability to tie multiple tweets to individual users, and the kinds of textual features that were deemed relevant. This resulted in us moving to an agile design approach, with rapid prototyping cycles that constantly changed. Finally, there were many design challenges that were imposed by the collaboration itself. For example, we found that our collaborators, who had a diverse range of backgrounds and experience levels, tended to have limited visual literacy. While the resulting design process was difficult, the collaboration did result in usable insights and results, with results being published at conferences [15, 38].

2 RELATED WORK AND BACKGROUND

Our team leveraged Moral Foundation Theory (MFT) for this project. MFT is a psychological model for describing the different dimensions underlying social discourse. MFT is widely used to study how values differ between individuals, and has been used to explain differences in political affiliation and belief systems [26]. Our version of MFT considered 6 foundation frames with opposing pairwise "virtues" and "vices" qualities, such as "Care" (virtue), which pertains to "the need to help or protect oneself or others", and the corresponding "Harm" (vice), which deals with "fear of damage or destruction to oneself or others" [17].

Recent visual analysis studies had looked at attitudes on social media related to public health, including attitudes towards the coronavirus [1], public health interventions [12, 21], climate change [11], and popular topics of discourse on social media [5, 6, 22]. Other systems have looked at information spread on social media, journalism, and misinformation [10, 16, 24, 35, 37], or focused on real-time information spread without geolocation [3, 8, 25]. Systems that focus on event detection [2, 13] tend to not analyze stance or moral foundations. Systems have also focused on temporal progression of topics [4, 7, 14, 20, 27, 43, 44], but none of these systems tie discourse to demographics or MFT framing.

3 METHODS

Our group had to rely on remote collaboration due to the pandemic lockdowns, with weekly meetings, and extensive use of collaborative tools leveraging our laboratory's expertise [32]. The core group spanned four research labs, evenly distributed: communications, NLP, causal inference in social media, and visual computing. Per our initial agreement, all group members are listed as co-authors. Our design process relied heavily on rapid prototyping using an activity-centered approach [31] and an agile framework, a combination
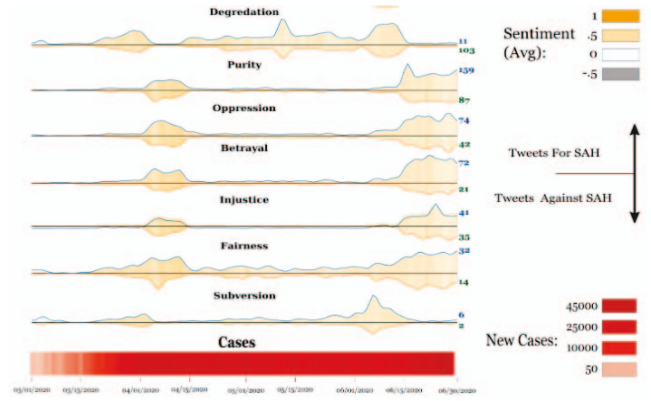

Figure 2: Design of timelines using aggregate tweets and small multiples for each tweets for each moral frame used on a larger non-geotagged dataset.

which had been very successful in our work with domain experts across the disciplines [29,33,34,42]. The project featured two stages, foraging for data and features, and hypothesis testing, which focused on establishing the cause of trends. Our collaborators' research activities were focused on identifying and explaining SAH-related phenomena and how Moral Framing factors into them. In short, we used MF as a lens to create a SAH-relevant dataset, and then to analyze and describe larger social trends. This agile, distributed design was similar to the distributed prototyping methods described by Losev et al [28]. However, our process focused more heavily on rapid prototypes without interaction design due to a larger emphasis on the data foraging stage.

3.1 Overall Design and Data Sources

The overall top-design emerged gradually and organically, and featured multiple coordinated views (Fig. 1). We gradually added data sources to our system, starting with a continually updated and expanded set of manually labeled geotagged U.S. tweets that expressed a moral frame and stance regarding stay-at-home (SAH) orders between March and May 2020. Tweets were manually labelled for SAH relevance, moral framing, and stance by expert annotators. Non-relevant tweets or tweets with no moral frames were excluded. The resulting corpus is described by Fatemi et al. [15]. During processing, each tweet was mapped to one of the 3113 U.S. counties, not including Alaska as we did not have election data for those areas. We then integrated several geospatial datasets: 2018 census information [36, 41], voting results from the 2016 US presidential election results [36], results from a New York Times survey on estimated mask usage [40], and daily COVID-19 cases and deaths for each county [23]. We also gradually extracted tweet features such as stance (pro SAH or against SAH), virality, sentiment, and vividness, and aggregated each moral frame stance within each county, resulting, over the length of the project, in 14 different continuous values for each county.

3.2 Four Custom Views

One view was dedicated to data summarization, and supported foraging activities at the start of the project. It gives an overview of the tweet features of interest, such as sentiment or vividness, broken down by Moral Frame. Due to limited visual literacy among the group at the start of the project, the view made extensive use of rotated, colored bar charts. To alleviate the cognitive load due to the extensive use of color, the colors were mapped as intuitively as possible (e.g., gray for not vivid, purple for vivid) based on perceptual principles and common media interpretation.
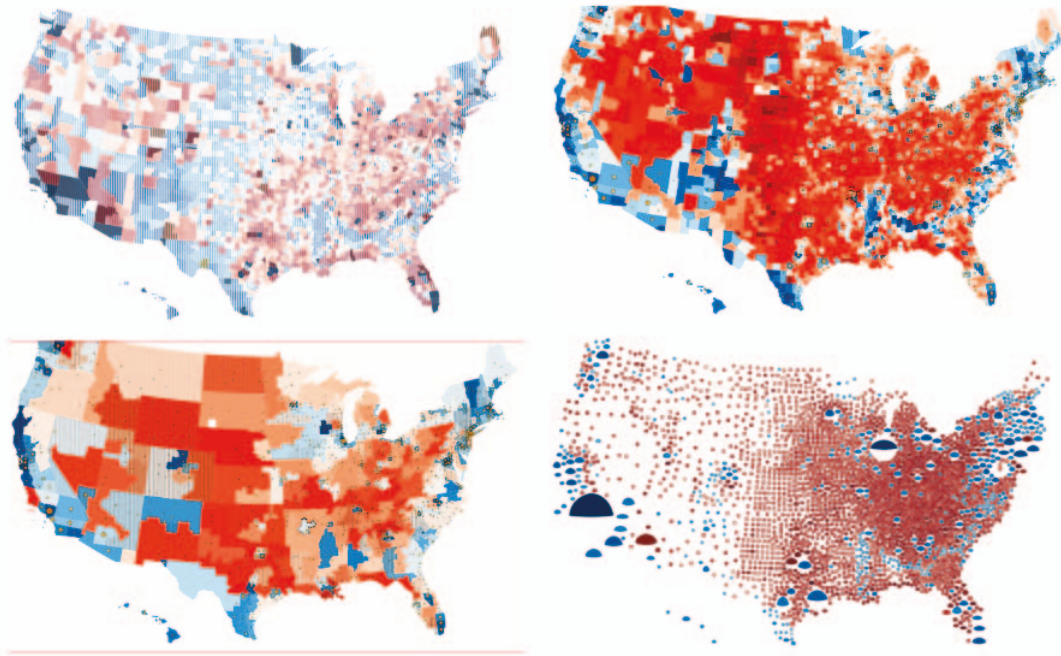
Figure 3: Examples of the progression of map design through the design process. (Top-left) Choropleth map using texture weaving [18] to encode two features. (Top-right) Choropleth map with a single color and glyphs showing tweet distribution. (Bottom-left) Choropleth map using textures and glyphs. Areas are aggregated by the intersection of district voting maps and counties for demographic data to make each section more evenly distributed in terms of population. Glyphs show tweets aggregated at the county level. (Bottom-right) Glyph-based map where shape encodes tweet features and color encodes demographics. All maps are zoomable

A custom Timeline panel (Fig. 2) showed the temporal distribution of the tweets, along with changes in COVID-19 rates, overall sentiment, and tweet popularity. The timeline panel grew with the project, with additional features being mapped to it. Again due to limited visual literacy, the timeline used basic cues such as layout, height, and color (hue and saturation) to encode time-varying multi-dimensional data. Because of scalability issues with this rich custom encoding, later applications to larger datasets used sparklines as an alternative.

Later on during the foraging stage, a Geospatial map panel showed the geographic distribution of tweets with a given Moral Frame, overlaid with demographic data. The design of this panel featured several costly trial-and-error design cycles (Fig. 3), where texture-blending based approaches which worked on other scientific datasets [18], and which seemed acceptable during lo-fi prototyping, backfired within our group in the hi-fi prototyping stage. We gradually developed an alternative, far more successful glyph encoding.

Finally, during the hypothesis testing stage, providing support for inference analysis became important. An Inference panel was created (Fig. 1), which allows for building predictive models within the front-end of the system, to support inferences about how different factors influence average tweet features.

4 EVALUATION

Our system aimed to support the development of insights related to SAH policy application in the U.S. We report here one of the case studies that our team performed over the course of several months, with results published and presented in several venues [15, 38]. The case study (Fig. 1) was performed remotely, with the team piloting the investigation via Zoom meetings, and the lead author operating the front-end of the system accordingly. Collaborators were also given independent access to the front end, which they used for additional analyses between meetings. For brevity, we distill the

main insights here:

Major Frames The dominant frames were Care, followed by Harm, which share the same foundation. Care tweets were overwhelmingly in support of SAH orders, while Harm were more mixed. Tweets with Harm also increased after the first month of Lockdowns, around times when SAH orders were removed, and had lower sentiment overall, which may reflect pandemic fatigue. Care tweets were significantly correlated with areas that self-reported higher mask use. An important implication of this analysis is that Care-targeted messaging of SAH orders helps.

Political Polarization The frames most associated with Democratic areas were Betrayal and Loyalty, which share a foundation, while the libertarian frames Subversion and Freedom were most associated with Republican areas. Betrayal tweets all happened in the start of May, largely in response to unmasked protests, and the mandated ending of SAH orders in major left-leaning cities located within right-leaning states. Notably, no Betrayal tweets came from Chicago, which may be due to the Mayor's continuation of SAH orders. Subversion and Freedom tweets increased after the first renewal of SAH orders in April. An important implication is that higher-granularity political analysis would help with targeted SAH messaging, for example including a Libertarian perspective in addition to Republican or Democrat views.

Spatial-Temporal Trends Tweets regarding SAH orders peaked at the start of the pandemic, with recurrences focused on major cities when SAH orders were renewed or ended. We also noted a large dip at the end of May, likely due to the George Floyd protests occupying the public zeitgeist at this time, with a short increase at the start of June in response to updates in public mask mandates in several cities. An important finding was that rural areas were underrepresented in the social media dataset.

Qualitative feedback from our collaborators across the disciplines was enthusiastic: *"Excellent for data exploration"*, *"Great for in-*

vestigating anomalies in the data.", *"Amazing work, the encodings work well together."*, *"This interface does very nicely with the data and domain knowledge we've been given."*, *"It's nice, pleasant to look at. And extremely informative."* In terms of ratings, components were rated favorably, with two requests for the additional on-demand glyph pictorial explanation. Experts found the system met these goals well: understanding the MF distribution, temporal MF distribution, MF political context, and verifying MF hypotheses. Most also found the system useful in analyzing MF tweet features, the corpus, exploring sentiment and topic relationships, and MF geographical distribution. The Causal inference and Communications specialists found the system very useful for understanding the emerging corpus, whereas some of the NLP experts continued to rely for corpus analysis tasks on their standard approach (LDA, topic clustering).

5 DISCUSSION

Our original design for this project was to create a large-scale analysis tool for analyzing networks of twitter stance and how different moral-framing affected the propagation of certain posts. In reality, we found that there were many difficulties with assessing a developing, real time problem while also incorporating a rich set of features in a collaborative setting. Despite this we did manage to create a unique system despite issues related to availability of data and issues that arose from a short-term collaboration with changing goals. Although our project did not seek to advance Moral Framing theory, our results point to a need for careful consideration of how the MF data is collected, for example using more granular political affiliation measures.

Generalizability and Scalability In response to collaborator and reviewer requests, we expanded our system to be usable for other datasets. Specifically, we relied on an earlier dataset of tweets related to the Black Lives Matter Movement [19], where we used hashtag content to estimate stance. The resulting dataset was slightly larger and covered a longer time period, with 1901 geotagged tweets spanning 2 years. The main challenge was adapting the timeline to not encode COVID-19 rates in the timeline, and to use textual features instead. Our results suggest that our approach generalizes well, with the main issue being the availability of geotagged data with sufficient quality annotation. In terms of scalability, we mainly faced issues with adapting the timeline to larger tweets. In a separate analysis using a non-geotagged dataset, we relied instead on a design that relied on small multiples, and shown in Fig. 2.

Our collaboration was a learning process in producing rapid, evolving designs, leading to several design lessons, beyond visual scaffolding [30] for improved visual literacy:

1. Minimize assumptions about the data when designing during the data collection process. This allows us to gradually specialize designs while keeping earlier progress in the event that the data changes. For example, initially, collaborators' main interest was in the relationship between moral framing and political affiliation. However, data availability made collecting multiple relevant tweets from each individual difficult. Thus we had to rely on location to infer demographics for tweets from each region, and reason about moral framing on an aggregate scale. Given the necessity of context, we further found that automatic annotation of moral frames performed too poorly. As a result, obtaining a pool of geotagged tweets that expressed a moral stance regarding stay at home orders was difficult, and yielded a much smaller dataset than what was originally expected at the start of the design process.

2. Collaboration issues are exacerbated by tight project timelines. We found that in the early stages our collaborators were unable to articulate what they wanted out of the designs: their earliest requirements were a basic COVID-19 dashboard with

political associations overlaid on the map. These findings were similar to those described by Sondag et al. [39]. However, we found that true requirements were more easily assessed by analyzing what points were discussed during meetings and what researchers were most focused on in their analysis. Collaborators also had insufficiently defined goals, and we found that there was not enough time to allow proper collaboration practices to mature (for example, a social science collaborator published project joint results without crediting the whole team). If time allows it, we would recommend relying on "ethnographic" methods rather than interviews for requirement gathering. We would also recommend repeatedly revisiting and enforcing collaborative practices during group meetings.

3. Keep designs generalizable. Because our study changed so rapidly, by the time the data collection and design process was finished, the original topic was considered "obsolete", with different topics taking over the online debate. In addition, our original data came from Twitter, which now has stricter limits on the amount of data scraping that can be used. In turn, our work needs to adapt to different topics, as well as different data sources to stay relevant. Our end design generalized well across case studies and datasets.

4. Dissemination issues stemming from data. Disappointingly, during the publication process we met resistance due to a misalignment of review expectation and design practice that couldn't be rectified within the project timeline. Most of these issues resulted from the unexpected difficulty in gathering a dataset of sufficient size and quality, largely due to the fact that automatic annotation proved to be ineffective, and the project timeline didn't provide sufficient time to manually collect an annotated sufficiently sized dataset. This was a common review issue in both CI/NLP-centered and vis-centric venues. As a result, several reviewers either claimed we had an insufficient dataset when using manual data, or we had an insufficient algorithm when using automatic annotation.

6 CONCLUSION

In conclusion, our project was successful in producing insights, and thus in achieving its proposed goals. It was also satisfying to be able to contribute to a better understanding of the pandemic policies, in particular at a time when vaccines were not yet available. At the same time, it was a difficult, extremely intensive experience, leading to several cases of burnout and dissatisfaction on the team, which were then aggravated by repeated difficulties in getting the methods published. If at all possible, we would rather embark on projects where (most of) the data has been already collected, and (most of) the requirements can be safely established during requirements engineering.

REFERENCES

[1] L. M. Aiello, D. Quercia, K. Zhou, M. Constantinides, S. Šćepanović, and S. Joglekar. How epidemic psychology works on twitter: Evolution of responses to the covid-19 pandemic in the us. *Humanit Soc Sci Commun*, 8, 2021. 2

[2] H. Bosch et al. Scatterblogs2: Real-time monitoring of microblog messages through user-guided filtering. *Trans. Vis. Comp. Graph.*, 19(12):2022–2031, 2013. doi: 10.1109/TVCG.2013.186 2

[3] N. Cao et al. Whisper: Tracing the spatiotemporal process of information diffusion in real time. *Trans. Vis. Comp. Graph.*, 18(12):2649–2658, 2012. doi: 10.1109/TVCG.2012.291 2

[4] N. Cao et al. Targetvue: Visual analysis of anomalous user behaviors in online communication systems. *Trans. Vis. Comp. Graph.*, 22(1):280–289, 2016. doi: 10.1109/TVCG.2015.2467196 2

[5] R. Chandrasekaran et al. Topics, trends, and sentiments of tweets about the covid-19 pandemic: Temporal infoveillance study. *J. Med. Int. Res.*, 22(10):e22624, 2020. doi: 10.2196/22624 2

[6] C.-H. Chang, M. Monselise, and C. C. Yang. What are people concerned about during the pandemic? detecting evolving topics about covid-19 from twitter. *J. Health. Info. Res.*, 5(1):70–97, 2021. doi: 10.1007/s41666-020-00083-3 2

[7] S. Chen et al. Co-bridges: Pair-wise visual connection and comparison for multi-item data streams. *Trans. Vis. Comp. Graph.*, 2020. doi: 10.1109/TVCG.2020.3030411 2

[8] S. Chen et al. R-map: A map metaphor for visualizing information reposting process in social media. *Trans. Vis. Comp. Graph.*, 26(1):1204–1214, 2020. doi: 10.1109/TVCG.2019.2934263 2

[9] F. Crim. Dear colleague letter on the coronavirus disease 2019 (covid-19). https://www.nsf.gov/pubs/2020/nsf20052/nsf20052.jsp, 2020. 1

[10] A. Dant and J. Richards. Behind the rumours: How we built our twitter riots interactive. *The Guardian*, 8, 2011. doi: uk/interactive/2011/dec/07/london-riots-twitter 2

[11] N. Diakopoulos, A. X. Zhang, D. Elgesem, and A. Salway. Identifying and analyzing moral evaluation frames in climate change blog discourse. In *8th AAAI Conf. on Weblogs & Soc. Med.*, 2014. 2

[12] C. Doogan et al. Public perceptions and attitudes toward covid-19 nonpharmaceutical interventions across six countries: a topic modeling analysis of twitter data. *J. Med. Int. Res.*, 22:e21419, 2020. doi: 10.2196/21419 2

[13] W. Dou et al. Leadline: Interactive visual analysis of text data through event identification and exploration. In *Trans. Vis. Comp. Graph.*, pp. 93–102, 2012. doi: 10.1109/VAST.2012.6400485 2

[14] M. Dörk, D. Gruen, C. Williamson, and S. Carpendale. A visual backchannel for large-scale events. *Trans. Vis. Comp. Graph.*, 16(6):1129–1138, 2010. doi: 10.1109/TVCG.2010.129 2

[15] Z. Fatemi, A. Bhattacharya, E. Zheleva, B. Di Eugenio, V. Dhariwal, L. Levine, A. Rojecki, A. Wentzel, and G. Marai. Understanding stay-at-home attitudes through framing analysis of tweets. *2022 Int. Conf. on Data Science and Adv. Anal.*, 8 2022. 2, 3

[16] Google. Google+ ripples. https://services.google.com/fh/files/misc/ripples.pdf, 2012. 2

[17] J. Graham et al. Moral foundations theory: The pragmatic validity of moral pluralism. In *Adv. in Exp. Soc. Psyc.*, vol. 47, pp. 55–130. Elsevier, 2013. doi: 10.1016/B978-0-12-407236-7.00002-4 2

[18] H. Hagh-Shenas, S. Kim, V. Interrante, and C. Healey. Weaving versus blending: a quantitative assessment of the information carrying capacities of two alternative methods for conveying multivariate data with color. *Trans. Vis. and Comp. Graph.*, 13(6):1270–1277, 2007. doi: 10.1109/TVCG.2007.70623 3

[19] J. Hoover et al. Moral foundations twitter corpus: A collection of 35k tweets annotated for moral sentiment. *Soc. Psy. and Pers. Sci.*, 11(8):1057–1071, 2020. doi: 10.1177/1948550619876629 4

[20] M. Hu, K. Wongsuphasawat, and J. Stasko. Visualizing social media content with sententree. *Trans. Vis. Comp. Graph.*, 23(1):621–630, 2017. doi: 10.1109/TVCG.2016.2598590 2

[21] T. Hu et al. Revealing public opinion towards covid-19 vaccines with twitter data in the united states: a spatiotemporal perspective. *J. Med. Int. Res.*, 26(7), 2021. doi: 10.2196/30854 2

[22] H. Jang et al. Tracking covid-19 discourse on twitter in north america: Infodemiology study using topic modeling and aspect-based sentiment analysis. *J. Med. Int. Res.*, 23(2):e25431, 2021. doi: 10.2196/25431 2

[23] John Hopkins University Center for Systems Science and Engineering. https://github.com/CSSEGISandData/COVID-19/tree/master/csse_covid_19_data, 2020. 2

[24] A. Karduni, I. Cho, R. Wesslen, S. Santhanam, S. Volkova, D. L. Arendt, S. Shaikh, and W. Dou. Vulnerable to misinformation? verifi! ACM, 2019. doi: 10.1145/3301275.3302320 2

[25] J. Knittel, S. Koch, T. Tang, W. Chen, Y. Wu, S. Liu, and T. Ertl. Real-time visual analysis of high-volume social media posts. *Trans. Vis. Comp. Graph.*, 28(1):879–889, 2021. doi: 10.48550/arXiv.2108.03052 2

[26] S. Koleva et al. Tracing the threads: How five moral concerns (especially purity) help explain culture war attitudes. *Res. in Pers.*, 46(2):184–194, 2012. doi: 10.1016/j.jrp.2012.01.006 2

[27] K. Kucher, R. M. Martins, C. Paradis, and A. Kerren. Stancevis prime: Visual analysis of sentiment and stance in social media texts. *J. Vis.*, 23(6):1015–1034, dec 2020. doi: 10.1007/s12650-020-00684-5 2

[28] T. Losev, S. Storteboom, S. Carpendale, and S. Knudsen. Distributed synchronous visualization design: Challenges and strategies. In *2020 IEEE Workshop on Evaluation and Beyond - Methodological Approaches to Visualization (BELIV)*, pp. 1–10, 2020. doi: 10.1109/BELIV51497.2020.00008 2

[29] T. Luciani et al. Details-first, show context, overview last: Supporting exploration of viscous fingers in large-scale ensemble simulations. *Trans. Vis. Comp. Graph.*, 2019. doi: 10.1109/TVCG.2018.2864849 2

[30] G. E. Marai. Visual Scaffolding in Integrated Spatial and Nonspatial Analysis. In E. Bertini and J. C. Roberts, eds., *EuroVis Work. on Vis. Anal. (EuroVA)*. The Eurographics Association, 2015. doi: 10.2312/eurova.20151097 4

[31] G. E. Marai. Activity-centered domain characterization for problem-driven scientific visualization. *Trans. Vis. and Comp. Graph.*, 24(1):913–922, 2017. doi: 10.1109/TVCG.2017.2744459 2

[32] G. E. Marai et al. Immersive analytics lessons from the electronic visualization laboratory: A 25-year perspective. *IEEE Comp. Graph. Applic.*, 39(3):54–66, 2019. 2

[33] A. Maries and A. e. a. Haque. Interactive exploration of stress tensors used in computational turbulent combustion. In *New Developments in the Visualization and Processing of Tensor Fields*, pp. 137–156. Springer, 2012. 2

[34] A. Maries, N. Mays, et al. GRACE: A Visual Comparison Framework for Integrated Spatial and Non-Spatial Geriatric Data. *IEEE Trans. Vis. Comp. Graph.*, 19(12):2916–2925, 2013. doi: 10.1109/tvcg.2013.161 2

[35] A. Mishra, S. Ginjpalli, and C. Bryan. News kaleidoscope: Visual investigation of coverage diversity in news event reporting. In *Pacific Vis.*, pp. 131–140. IEEE Computer Society, Los Alamitos, CA, USA, apr 2022. doi: 10.1109/PacificVis53943.2022.00022 2

[36] MIT Election Data and Science Lab. https://github.com/MEDSL/2018-elections-unoffical/blob/master/election-context-2018.md, 2020. 2

[37] P. Rodgers, L. Zhang, and A. Fish. General euler diagram generation. In *Diagrams*, pp. 13–27, 2008. doi: 10.1007/978-3-540-87730-1_6 2

[38] A. Rojecki et al. The moral imperatives of self-quarantining. *American Political Science Association Annual Meeting*, Oct 2021. 2, 3

[39] M. Sondag, C. Turkay, K. Xu, L. Matthews, S. Mohr, and D. Archambault. Visual analytics of contact tracing policy simulations during an emergency response. *Computer Graphics Forum*, 41(3):29–41, 2022. doi: 10.1111/cgf.14520 4

[40] The New York Times and Dynata. A detailed map of who is waring masks in the u.s. https://github.com/nytimes/covid-19-data/blob/master/mask-use, 2020. 2

[41] United States Census Bureau. https://www.census.gov/data/tables/time-series/demo/voting-and-registration/congressional-voting-tables.html, 2020. 2

[42] A. Wentzel and P. e. a. Hanula. Cohort-based T-SSIM visual computing for radiation therapy prediction and exploration. *Trans. Vis. Comp. Graph.*, 26(1):949–959, 2019. 2

[43] Y. Wu et al. Opinionflow: Visual analysis of opinion diffusion on social media. *Trans. Vis. Comp. Graph.*, 20(12):1763–1772, 2014. doi: 10.1109/TVCG.2014.2346920 2

[44] J. Zhao et al. Fluxflow: Visual analysis of anomalous information spreading on social media. *Trans. Vis. Comp. Graph.*, 20(12):1773–1782, 2014. doi: 10.1109/TVCG.2014.2346922 2