

Supplementary Materials and References

Unmasking the conversation on masks: Natural language processing for topical sentiment analysis of COVID-19 Twitter discourse

Abraham Sanders, Rachael White, Lauren S. Severson,
Rufeng Ma, Richard McQueen, Haniel C. Alcânata Paulo,
Yucheng Zhang, John S. Erickson, Kristin P. Bennett,
Rensselaer Polytechnic Institute, Troy, New York

This supplement contains additional details on the composition of our analysis pipeline and computational methods. Interpretations are included for each of the fifteen clusters generated.

1 Analysis Pipeline

1.1 Embedding & Sentiment Scoring

We use the universal-sentence-encoder-large v5 implementation from TensorFlow Hub. The pre-trained model can be obtained at <https://tfhub.dev/google/universal-sentence-encoder-large/5>.

1.2 Clustering & Subclustering

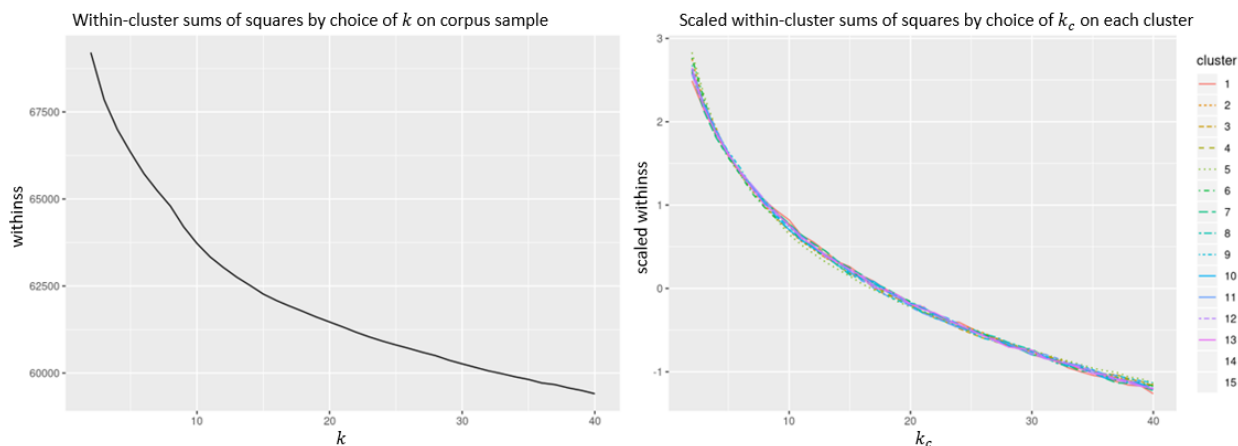


Figure 1: The curvature of the k-means objective function begins to flatten out at $k = 15$ for clusters and $k_c = 15$ for subclusters.

We use the Hartigan-Wong¹ k-means algorithm included in the ‘stats’ package in R² with 25 random starts for a maximum of 200 iterations. To find suitable choices for k and k_c we use the elbow method, where the within-cluster sums of squares objective function is measured over a range of choices for k and k_c in an attempt to find the point which strikes a balance between minimization of the objective and avoiding over-clustering. We chose 15 for both k and k_c . Figure 1 contains the objective function measurements. The measurements for k_c on the right are normalized to zero mean and unit variance so the curves can be easily compared.

We use the Barnes-Hut implementation³ of t-SNE included in the ‘Rtsne’ package⁴ for R. An initial PCA step is done and then t-SNE is run with $\theta = 0.5$. When running on the entire corpus sample, we chose a perplexity of 70 for 2000 maximum iterations. When running on individual clusters, we chose a perplexity of 25 for 750 maximum iterations.

We use the Plotly⁵ package for cluster and subcluster visualization.

1.3 Cluster & Subcluster Labeling

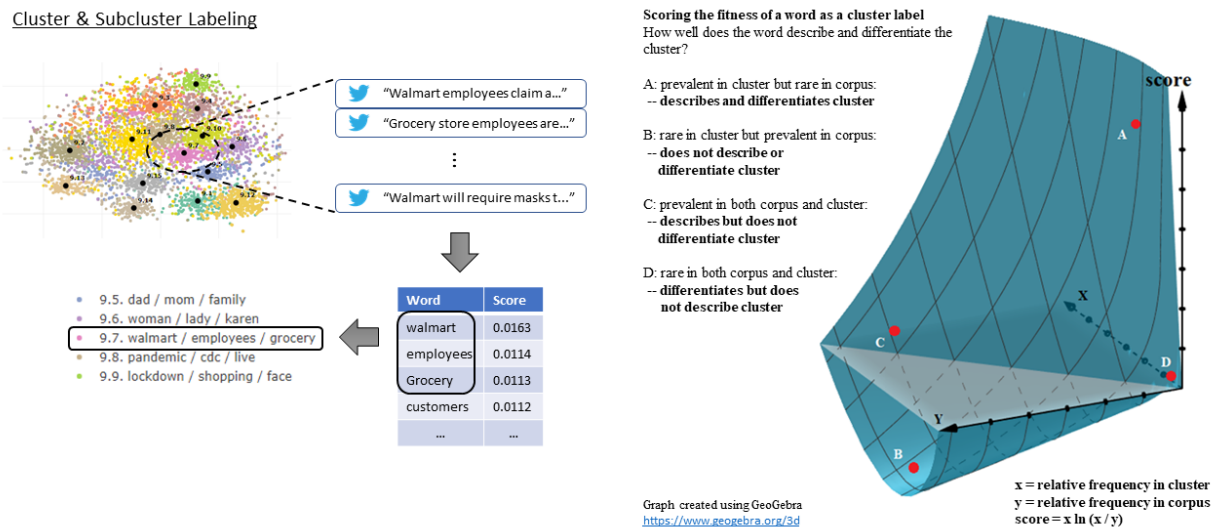


Figure 2: The figure on the left illustrates the labeling process for each cluster and subcluster, with words being ranked according to their score. The figure on the right shows a plot for the score values across the cluster and corpus probability domains, as well as interpretations of different score regions.

1.4 Cluster & Subcluster Summarization

DistilBART is based on Facebook’s BART model, a sequence-to-sequence architecture combining a BERT⁶-like bi-directional transformer encoder with a GPT⁷-like transformer decoder. The instance we use is configured with 12 encoder layers, 6 decoder layers, and 16 attention heads in each encoder and decoder layer. We decode with a sampling temperature of 1.0 and a beam width of 6 for a maximum of 60 tokens. The fine-tuned model instance and its complete configuration can be obtained at <https://huggingface.co/sshleifer/distilbart-xsum-12-6>.

We experimented with two additional strategies for selecting the input for cluster summarization - concatenating tweets closest to the cluster centers and concatenating the single closest tweets to each child subcluster center. We found by manually comparing the generated summaries to a sample of tweets from each cluster that these strategies are more prone to misrepresentation of the cluster theme than the strategy of concatenating the model-generated summaries of its subclusters. We repeated our experiments for each summarization strategy using two DistilBART model instances - one fine-tuned on the xsum task and one fine-tuned on the cnn/dailymail summarization task.^{8,9} We used the same inference hyperparameters specified above in all experiments. We found that the models fine-tuned on the xsum task yielded more concise summaries while those trained on the cnn/dailymail dataset yielded more verbose summaries which tended to quote original tweet text from the corpus directly. A full listing of the experimental results can be accessed at <https://therensselaeridea.github.io/COVID-masks-nlp/analysis/summarization.html>, and the model instance fine-tuned on the cnn/dailymail dataset can be obtained at <https://huggingface.co/sshleifer/distilbart-cnn-12-6>.

2 Statistical Methods

The divisiveness score is based on the Sarle’s Bimodality Coefficient¹⁰ (BC) with added correction in order to factor in uncertainty from small sample sizes. Letting γ and κ be the sample skewness and kurtosis of the discretized sentiment and assuming negligible covariance between both statistics, we can approximate the BC mean and variance as

$$\mathbb{E}[\text{BC}] = \frac{\gamma^2 + 1}{\kappa}$$
$$\text{Var}[\text{BC}] = \frac{1}{\kappa^2} \text{Var}[\gamma^2] + \frac{(\gamma^2 + 1)^2}{\kappa^4} \text{Var}[\kappa]$$

¹¹ where $\text{Var}[\gamma^2]$ and $\text{Var}[\kappa]$ have known formulas.¹² The uniform distribution has a known BC of $5/9$ ¹⁰ and we reason that for small samples, where $\text{Var}[\text{BC}]$ is large, we cannot ensure with a high degree of certainty that $\text{BC} \neq 5/9$, which motivates our sample correction through comparison to the standard normal distribution. We take $|z| = |5/9 - \mathbb{E}[\text{BC}]| / \sqrt{\text{Var}[\text{BC}]}$ and let $w = \text{erf}(|z|/\sqrt{2})$ since $|z|$ is belongs to a standard normal truncated on \mathbb{R}^+ , so that the corrected bimodality coefficient BCc is given by the weighted average

$$\text{BCc} = w\text{BC} + (1 - w)\frac{5}{9}$$

As such, for samples where $5/9$ is likely to be sampled from the distribution of BC the value of BCc will be closer to $5/9$, while for samples where such sampling is unlikely the value of BCc will be closer to the sample BC.

Finally, in order improve visual interpretability, we define the divisiveness score as

$$\text{Div} = \text{logit}(\text{BCc}) - \text{logit}(5/9)$$

which ensures the previously mentioned properties with respect to values greater than, less than and equal to zero.

3 All Cluster Interpretations

• Cluster 1: trump / president / realdonaldtrump

- **DistilBart summary:** *People have been reacting to news that President Donald Trump has refused to wear a face mask in public to protect himself from the deadly coronavirus pandemic.*
- **Interpretation:** This cluster makes apparent a trend of Twitter users expressing a spectrum of attitudes towards the standing U.S. president, Donald Trump, during the first five months of 2020. Opinions specifically revolve around Trump’s handling of the COVID-19 pandemic in the United States. Distinctly, there exists an evident theme of frustration arising from observations that Trump has refused to wear a mask in public appearances, despite statements from public health officials encouraging the action. It should be noted that, in complement, a sizeable discussion thread of a more positive and supporting nature also exists in relation to President Trump. A major theme observed here among the pro-Trump tweets is the impression that the media is biased against the president, and that this in turn fosters a public motive to exaggerate the virus. The anti-Trump tweets in this cluster are mostly focused on the president’s long refusal to wear a face mask, although this finding is predictable given the nature of the data set from which the tweets are drawn.

• Cluster 2: vaccine / flu / stop

- **DistilBART summary:** *Following the news that people in the US are being urged to wear face-covering masks to prevent the spread of a new virus that has killed more than 4,000 people in China.*
- **Interpretation:** Cluster 2, “vaccine / flu / stop”, is a grim cluster in terms of its overall sentiment, and is distinctly polemical in its semantics. It is found that the majority of tweets sampled from this cluster are pro-mask tweets complaining about individuals who don’t wear masks. The dominant attitude towards

masks observed among the tweets sampled for inspection is positive, despite the overall negative sentimentality computed for the cluster as a whole. In contrast with the more semantically upbeat "face / hands / stay" cluster (Cluster 15), this aggregation contains an apparent host of tweets related to death and dying. The social nature of disease is a major motif (i.e. "Your actions affect all of us.").

- **Cluster 3: lockdown / social / distancing**

- **DistilBART summary:** *Following the news that the US government has ordered people to wear face masks in public to prevent the spread of the deadly Covid-19 coronavirus, people across the world have been reacting to the news on social media.*
- **Interpretation:** Cluster 3 gives an indication of the societal turbulence relating to and arising from mask mandates, social distancing enforcement, and similar lockdown-related occurrences globally. Paradoxically, the overall average sentiment of -0.0941 computed for this cluster is borderline neutral. Individual topics manifesting in this representation are observed to vary greatly, but the concerns represented in the tweets sampled appear to be, at minimum, tangentially centered around the themes of imprisonment, isolation, and quarantine. A strong racial emphasis is evident, with discourse notably focusing around protests of the Black Lives Matter movement, an international phenomenon co-occurring with the coronavirus pandemic mid-year. Several subclusters of Cluster 3 entertain conversations about international responses to the virus, notably around the idea that mask-wearing to prevent the spread of disease agents is a long-standing cultural norm in some regions. In keeping with the slightly negative overall computed sentiment for this cluster, many of the tweets seem to carry a sarcastic tone and a strong indication of resentment towards perceived hypocrisy surrounding mask-usage.

- **Cluster 4: cdc / public / risk**

- **DistilBART summary:** *People have been sharing their opinions on whether or not they should wear face masks in public to prevent the spread of the deadly coronavirus.*
- **Interpretation:** Cluster 4 documents the narrative of the interplay between CDC recommendations and the public response regarding the usage of masks as protection from COVID-19. In early April, the CDC began encouraging individuals to wear a mask or face covering when leaving home.⁷ Extensive examples of public dissent resulting from the the issued guidelines, specifically surrounding the question "Is it necessary to wear a mask?", can be noted within this cluster. We studied tweets with the highest and lowest sentiment scores, and found that the most negative showcase individuals expressing skepticism towards or discouraging mask-wearing. The tweets with higher-scoring sentiment trend towards encouraging others to wear a mask, and do so with gentler and friendlier wordings. This cluster is also distinctive in the frequency at which scientific language and terminology present within it. For example, references are made to vaccines, health research, and the scientific evidence lending support to the official consensus on how the coronavirus is transmitted. Extensive consideration is given to the topic of risk, both related to the virus and to failure to follow preventative measures.

- **Cluster 5: pandemic / global / middle**

- **DistilBART summary:** *Following a report by the US Centers for Disease Control and Prevention (CDC) that the coronavirus pandemic could be brought under control in four to eight weeks if people wore face masks in public.*
- **Interpretation:** This cluster gives voice to a variety of issues that people are worried about, such as that wearing a mask is not necessarily comfortable, and concern about other people who are not wearing a mask in public. Some tweets also merged a political slant into the concept of wearing a mask. For instance, when "pandemic" and "global" were mentioned, politics between nations about the coronavirus were compared. This cluster is highly related to government orders, including Trump government guidelines on COVID-19. The sentiment of this cluster is close to neutral but the most negative tweets showed their dissatisfaction using a global perspective.

- **Cluster 6: coronavirus / mandate / governor**

- **DistilBART summary:** *Following a spike in cases of the deadly coronavirus (COVID-19) in the US, social media users have been reacting to news that a governor in Oklahoma has tested positive for the virus.*
- **Interpretation:** Although the BART-generated summary for this cluster indicates its specificity to discussion about one political figure from Oklahoma, examination of tweets within Cluster 6 make clear that this cluster includes discussion about U.S. governors in a more general sense. The governmental efforts of individual states to respond to and contain the spread of COVID-19 are heavily represented in the conversation, with the cumulative suggestion being that the attitudes of political leaders hold considerable influence over a given state's trajectory for recovery from viral devastation. This cluster is about politics, with a specific emphasis on the enforcement of mask mandates within individual states in the U.S. The overall average sentiment score falls within the range we define as neutral.

- **Cluster 7: 19 / fucking / real**

- **DistilBART summary:** *Social media users have been sharing their reactions to the news that Covid 19 is back.*
- **Interpretation:** Cluster 7 speaks to the reality that some regions around the world have experienced a resurgence in COVID-19 infections after individuals returned to their daily routines preemptively, before the proper containment of the virus could be achieved in those regions. The tone of tweets overall is anxiety-ridden and highly tense, especially amongst those tweets explicitly dealing with individual experiences of infection with the virus. There is a haunting aura that arises from the apparent public impression of COVID-19 as a looming spectre. Among the tweets that had the most negative sentiment scores, we observe a trend of people sharing grievances in the form of stories about their family members or friends testing positive for the virus or suffering from COVID-19. Tweets of this strain urge others to wear a mask with high apparent emotion. The fact that this apparently impassioned cluster was calculated as having an overall neutral sentiment speaks to the polarity of the discourse at play.

- **Cluster 8: wearamask / maskwearing / masking**

- **DistilBART summary:** *Social media users have been reacting to news that a deadly coronavirus outbreak in the US has killed more than 1,000 people.*
- **Interpretation:** Looking past the fact that Cluster 9 falls within the sentiment range we define as neutral, this cluster marks the transition from the class of clusters presented with overall negative mean sentiment scores to those with overall positive. In this cluster, discussion streams give rise to a general narrative of people sharing life experiences during the coronavirus. The importance of wearing a mask is again emphasized, making this representation consistent with the semantic character of previous clusters. Many tweets inspected include media reports on new cases in each state. The most negative tweets seem to earnestly preach the necessity of taking preventative measures to combat the spread of coronavirus amongst minorities who are found to fall victim to the disease on a highly disproportionate basis, including those who are especially sensitive to the respiratory symptoms of the illness. The most positive tweets express happiness at apparent observations of others starting to embrace mask-wearing. Another event mentioned frequently in this cluster is the mask-wearing rules of the 1918 pandemic. Tweets of this strain speak to the idea that the division in public attitudes towards mask wearing occurring amidst the 2020 coronavirus pandemic is mirrored in previous health crises of similar scale.

- **Cluster 9: quarantine / gloves / store**

- **DistilBART summary:** *Social media users have been sharing their experiences of wearing face masks in the wake of the recent outbreak of the deadly coronavirus in the US.*

- **Interpretation:** In this cluster, people describe their experiences venturing into the public sphere to acquire basic necessities amidst COVID-19. Notably, in the subcluster asthma/breathing/breathe, those who have asthma or other breathing issues assert that they are wearing masks to prevent the spread of virus, and the voices of these tweets amplify the admonition that others who do not have breathing issues have no excuse not to do so. The most negative tweets focus on people who elect not to wear a mask while using public transportation, and the most positive point out silver linings of life in quarantine, such as having more time to interact with family while at home.

- **Cluster 10: corona / hero / gagging**

- **DistilBART summary:** *Social media users have been reacting to news that people in India are wearing face masks to protect themselves from the deadly corona virus.*
- **Interpretation:** In Cluster 10, medical workers are a group that is frequently discussed. One narrative that seems to be of particular interest When medical workers are taking care of patients in Europe and Asia, because of the virus rapidly spreading, medical systems in Europe and Asia collapsed quickly in that the number of patients exceeded hospital capacity. It is observed that because of the shortage of medical workers, individuals had to overload on work.

- **Cluster 11: droplets / spread / stop**

- **DistilBART summary:** *People have been sharing their opinions on whether or not they should wear face masks to protect themselves against the norovirus.*
- **Interpretation:** This cluster contains information about people discussing how COVID-19 is transmitted. There are notable mentions of “droplet” and virus size. Many tweets mentioned new studies that showed masks are the most effective way to stop coronavirus transmission by blocking respiratory droplets. This cluster is more about scientific studies on masks and the way the virus is spreading. The most positive tweets, interestingly, appear to be those that report such scientific results, even despite the emotion-masking presence of including statistics and terminologies.

- **Cluster 12: n95 / surgical / microns**

- **DistilBART summary:** *News that a shortage of N95 respirator masks in the US is causing a worldwide shortage has been shared on social media.*
- **Interpretation:** Discourse within Cluster 12 focuses on information about N95 masks and related forms of personal protective equipment (PPE). The evolution of the conversation around the accessibility of medical resources over the timeline of the pandemic is clearly represented. One notable stream of discussion points to the presence of a debate over how useful cloth masks are as guards against infectious agents in comparison to surgical masks. The shortage of respirators experienced by the medical community in the United States is also referenced, as is the concept that a change in tonality and meaning surrounding the suggested usage of N95 masks was observed from the U.S. CDC shortly after the pandemic infiltrated U.S. borders.

- **Cluster 13: coronavirus / face / make**

- **DistilBART summary:** *Following the news that the deadly coronavirus outbreak in Texas has been declared a public health emergency, people have been sharing their views on the virus on social media.*
- **Interpretation:** Tweets inspected from this cluster constitute an overall neutral conversation about the coronavirus as a disease and as a historical moment. Commonplace concerns about individual health matters such as testing and treatment accessibility, mask-wearing as a preventative measure, and the scale and pace of the virus as an infectious agent are expressed in abundance. Suggestions and tutorials for

homemade face masks make up one prominent branch of conversation. There are also notable cases of negatively-toned tweets. For example, people use the adjective “selfish” to imply that those who are not wearing a mask should experience guilt, rather than following a trend observed in previous clusters of individual tweets outright attacking those unwilling to wear a mask.

- **Cluster 14: face / facemask / coronavirus**

- **DistilBART summary:** *Following the news that the deadly Coronavirus-19 virus is spreading around the world, people have been sharing their views and advice on face masks on social media.*
- **Interpretation:** Cluster 14 takes on a distinctly commercial and artistic theme. Tweets inspected for this cluster tend to adopt a gentler tone than previous clusters in their advocacy of mask usage to combat COVID-19. The overall average sentiment score is medium-positive. Many tweets are clearly advertisements which include merchants promoting custom/personalized face mask products. Other positive tweets include individuals expressing their feelings on and experience with creating handmade masks.

- **Cluster 15: hands / wash / stay**

- **DistilBART summary:** *Social media users have been sharing their tips and advice on how to prevent the spread of the deadly coronavirus.*
- **Interpretation:** Our most positive cluster overall, “hand / wash / stay” is composed of distinct thrusts of tweets sharing tips on prevention measures for stopping the spread of COVID-19, as well as helpful tips for self-protection from the virus. There appears to be highly positive sentiment expressed towards masks and other PPE in general, and well-meaning admonitions such as “Wash your hands and socially distance!” are frequent. In contrast to other clusters we have explored, the Cluster 15 tweets surveyed contain comparatively little in the way of aggressive, sarcastic or antagonistic semantic content. As such, this cluster may be interpreted to be an echo of the official messaging of the CDC and similar organizations.

References

1. J. A. Hartigan and M. A. Wong, “Algorithm AS 136: A K-Means Clustering Algorithm,” *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, vol. 28, no. 1, pp. 100–108, 1979. [Online]. Available: <http://www.jstor.org/stable/2346830>
2. R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2019. [Online]. Available: <https://www.R-project.org/>
3. L. van der Maaten, “Accelerating t-sne using tree-based algorithms,” *Journal of Machine Learning Research*, vol. 15, pp. 3221–3245, 2014.
4. J. H. Krijthe, *Rtsne: T-Distributed Stochastic Neighbor Embedding using Barnes-Hut Implementation*, 2015, r package version 0.15. [Online]. Available: <https://github.com/jkrijthe/Rtsne>
5. P. T. Inc. (2015) Collaborative data science. Montreal, QC. [Online]. Available: <https://plot.ly>
6. J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
7. A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, “Improving language understanding by generative pre-training,” 2018.
8. A. See, P. J. Liu, and C. D. Manning, “Get to the point: Summarization with pointer-generator networks,” *CoRR*, vol. abs/1704.04368, 2017. [Online]. Available: <http://arxiv.org/abs/1704.04368>
9. K. M. Hermann, T. Kocisky, E. Grefenstette, L. Espeholt, W. Kay, M. Suleyman, and P. Blunsom, “Teaching machines to read and comprehend,” in *Advances in neural information processing systems*, 2015, pp. 1693–1701.

10. R. Pfister, K. A. Schwarz, M. Janczyk, R. Dale, and J. Freeman, "Good things peak in pairs: a note on the bimodality coefficient," *Frontiers in psychology*, vol. 4, p. 700, 2013.
11. H. Seltman, "Approximations for mean and variance of a ratio," *unpublished note*, 2012.
12. D. B. Wright and J. A. Herrington, "Problematic standard errors and confidence intervals for skewness and kurtosis," *Behavior Research Methods*, vol. 43, no. 1, pp. 8–17, 2011.