
Non-Pharmaceutical Intervention Discovery with Topic Modeling

Jonathan Smith¹ Borna Ghotbi^{2,3} Seungeun Yi¹ Mah Parsa^{2,4}

Abstract

We consider the task of discovering categories of non-pharmaceutical interventions during the evolving COVID-19 pandemic. We explore topic modeling on two corpora with national and international scope. These models discover existing categories when compared with human intervention labels while reduced human effort needed.

1. Introduction

The COVID-19 pandemic has seen varying government responses in countries around the world. More than 2.24 million confirmed cases and 152.5 thousand deaths worldwide (January 1, 2020 - April 19, 2020) (WHO, 2020) are attributed to novel coronavirus (SARS-CoV-2). Due to the high transmission rate, non-pharmaceutical interventions (NPIs) play a critical role in stemming the spread of the virus (Ferguson et al., 2020). Studies assess the effectiveness of interventions in "flattening the curve" (Koo et al., 2020; Tuite et al., 2020), using individual or compartmental models (Biswas et al., 2014) to simulate scenarios for fixed parameters. Broad categories encompassing physical distancing, business closures, and travel restrictions, are used and tailored to the specific country or region the model focuses on (Wang et al., 2020). More specific interventions are less likely to be tracked, leaving the effects of increased funding, stratified closure, and administrative flexibility unestimated. Increasing the granularity of NPIs recorded is limited by manual labor involved in intervention identification and evaluation. This is seen in the categorizations used by four projects currently compiling data (McCoy et al., 2020; HIT COVID, 2020; ACAPS, 2020; Hale et al., 2020) which record 63, 23, 35, and 13 NPI types respectively.

In this study we leverage the availability of global text data on intervention announcements to automate the task of intervention category discovery, which will speed up the identification of interventions with strong effects on transmission.

¹Layer 6 AI, Canada ²Vector Institute, Canada ³University of British Columbia, Canada ⁴University of Toronto, Canada. Correspondence to: Jonathan Smith <jonathan@layer6.ai>.

Table 1. Top 10 words per topic for 3 topics in D_{CAN} , using LDA with $K = 25$.

| Topic | Words |
|-------|---|
| 0 | businesses, financial, program, government, business, relief, time, covid, announced, fees |
| 11 | closed, closure, facilities, parks, public, closures, visitors, centres, please_visit, notice |
| 20 | park, leash, couple, equipment, play, enjoy, neighbourhood, staying, bylaws, however |

We propose a method for using topic modeling to discover intervention categorizations. This method uses two well-known topic modeling algorithms, Latent Dirichlet Allocation (LDA) and Hierarchical Dirichlet Processes (HDP), combined with the unique structure of the NPI category discovery problem. We introduce evaluation methods for comparing unsupervised discovery models with existing intervention labels and demonstrate that our approach discovers known categories of NPIs as seen in the selected topics in Table 1. Our method has shown useful performance on datasets with global scope, suggests that better categorizations can be learned by pooling knowledge across countries.

2. Datasets, Method, Results

2.1. Datasets

This paper uses two datasets: 1) the Oxford COVID-19 Government Response Tracker v4.0 (Hale et al., 2020) (D_{OXF}), which contains 13 intervention categories from 156 countries, and 2) CAN-NPI dataset (McCoy et al., 2020) (D_{CAN}), which contains 63 intervention categories found in Canada. The former provides notes for each change to the status of an intervention citing public announcement texts, quotes, email interactions with public health authorities, and links. In contrast, the latter provides more standardized full-text articles for each intervention tagged at the provincial, territorial, and municipal level in a single country. The D_{CAN} dataset also links all eligible interventions from the Oxford intervention categories. For consistency of results, the period of January 1, 2020 - April 19, 2020, is used for both of the datasets. Table 2 compares dataset size and composition. It is also essential that international data is

evenly sampled from across regions. In D_{OXF} , countries have similar distributions of average note counts in Asia (31), Europe (28), Africa (29), and the Americas (29).

Table 2. Summary of datasets: The number of documents D , vocabulary size V , average words per document W_{mean} , and standard deviation in words per document W_{std} .

| Dataset | D | V | W_{mean} | W_{std} |
|-----------|-------|---------|------------|-----------|
| D_{CAN} | 1,290 | 574,417 | 248 | 204 |
| D_{OXF} | 3,900 | 114,743 | 17 | 16 |

2.2. Method

2.2.1. PREPROCESSING AND TOPIC MODELING

A corpus is constructed from each dataset. We preprocess the documents in each dataset to remove default english stopwords from NTLK (Bird et al., 2009) and geographic words (eg. province or country names) from data labels. Bigrams are constructed from word pairs that generally occur together (eg. “public health”) and are added as well.

We apply LDA (Blei, 2012) and HDP (Teh et al., 2004) using `gensim` (Řehůřek & Sojka, 2010) to generate topics for each corpus. Topic coherence, C_v , is calculated as a measure of topic interpretability using co-occurrence of words and cosine similarity between each top word and the total of all top words (Röder et al., 2015; Syed & Spruit, 2017). LDA is run to generate a specific number of topics, K , then coherence is calculated for each K . For LDA, the maximum coherence is associated with the near-optimal number of topics found from a corpus. HDP predetermines the number of topics and calculates a single coherence over all of them. The online nature of these algorithms allows for topics that can be easily extended when more data is gathered, helping researchers respond quickly in a fast-moving pandemic.

2.2.2. SIMILARITY EVALUATION

We evaluate the suitability of a topic as an intervention category by comparing with existing labels from subject matter experts (SMEs). Mean cosine similarity (S) between weighted topic occurrence per document ($\mathbf{T} \in [0, 1]^{D \times T}$) and label occurrence ($\mathbf{L} \in [0, 1]^{D \times C}$) for D documents, C labels, and T topics is one metric. This is calculated as:

$$\frac{1}{CT} \sum_{i=1}^T \sum_{j=1}^C (\mathbf{T}^\top \mathbf{L})_{i,j}$$

A higher similarity shows that topics are similar in structure to given labels. We evaluate individual label discovery by assuming that each topic can correspond to at least one intervention category. To determine the coverage (Cov) we

Table 3. Evaluation of topic coherence (C_v), mean cosine similarity between topics and intervention categories (S), and coverage ratio of top 1 topic selection over intervention categories $Cov(\%)$ for the top K topics. Two topic models (TM) are compared.

| Dataset | TM | K | C_v | S | $Cov(\%)$ |
|-----------|-----|-----|--------------|--------------|-------------|
| D_{CAN} | LDA | 10 | 0.435 | 0.029 | 0.17 |
| | LDA | 25 | 0.451 | 0.023 | 0.33 |
| | LDA | 50 | 0.427 | 0.023 | 0.44 |
| | LDA | 100 | 0.379 | 0.018 | 0.63 |
| D_{OXF} | HDP | 10 | 0.608 | 0.029 | 0.17 |
| | HDP | 25 | 0.608 | 0.025 | 0.37 |
| | HDP | 50 | 0.608 | 0.019 | 0.59 |
| | HDP | 100 | 0.608 | 0.013 | 0.68 |
| D_{OXF} | LDA | 10 | 0.505 | 0.010 | 0.64 |
| | LDA | 25 | 0.427 | 0.010 | 0.86 |
| | LDA | 50 | 0.436 | 0.010 | 0.93 |
| | LDA | 100 | 0.373 | 0.008 | 0.93 |
| D_{OXF} | HDP | 10 | 0.683 | 0.141 | 0.64 |
| | HDP | 25 | 0.683 | 0.137 | 0.93 |
| | HDP | 50 | 0.683 | 0.128 | 0.93 |
| | HDP | 100 | 0.683 | 0.114 | 1.00 |

select the topic with the highest similarity for each label and count the number of distinct topics used to cover all labels. Coverage ratio ($Cov(\%)$) is found with Cov/C . Coverage ratio combined with high similarity indicates the discovery of label categories. Combining these metrics enables model performance evaluation without additional work by SMEs.

2.3. Results

We summarize the results of category discovery with topic modeling on D_{CAN} and D_{OXF} datasets in Table 3 using $K \in \{10, 25, 50, 100\}$. HDP leads to topics with significantly more coherence, while a clear tradeoff must be made between overall similarity and coverage. Similarity drops less precipitously as the number of topics increases when using LDA as opposed to HDP. D_{OXF} has a much smaller number of categories, so there is a clear early convergence in coverage. In practice, a higher number of topics will be harder for humans to check for interventions but should lead to coverage of existing categories and aid with discovery.

3. Conclusion and Future Work

In this paper, we briefly demonstrate LDA- and HDP-based methods for NPI discovery on D_{CAN} and D_{OXF} . As our results have verified, combining coherence, topic similarity, and coverage can provide intervention category suggestions for experts to explore and find new NPIs leveraging national and international data. For future work, we will explore prediction-focused supervised LDA and identification of NPIs in an online manner.

References

- COVID-19 situation report-19. <http://aiweb.techfak.uni-bielefeld.de/content/bworld-robot-control-software/>, 2020. [Online; accessed 12-May-2020].
- ACAPS. ACAPS COVID-19: Government measures dataset - humanitarian data exchange (version 1.0). 2020.
- Bird, S., Klein, E., and Loper, E. *Natural Language Processing with Python*. O'Reilly Media Inc, 2009. URL <http://nltk.org/book>.
- Biswas, M. H., Paiva, L. T., and De Pinho, M. d. R. A seir model for control of infectious diseases with constraints. *MBE*, 2014. doi: 10.3934/mbe.2014.11.761.
- Blei, D. M. Probabilistic topic models. *Communications of the ACM*, 55(4):77–84, 2012.
- Ferguson, N., Laydon, D., Nedjati Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Cucunuba Perez, Z., Cuomo-Dannenburg, G., et al. Report 9: Impact of non-pharmaceutical interventions (npis) to reduce covid19 mortality and healthcare demand. 2020.
- Hale, T., Petherick, A., Phillips, T., and Webster, S. Variation in government responses to covid-19 (version 4.0). 2020. doi: BSG-WP-2020/031.
- HIT COVID, T. Health interventions tracking for covid-19 (hit-covid) (version v0.3). 2020. doi: 10.5281/zenodo.3765628.
- Koo, J. R., Cook, A. R., Park, M., Sun, Y., Sun, H., and Lim, J. T. e. a. Interventions to mitigate early spread of sars-cov-2 in singapore: a modelling study. *The Lancet Infectious Diseases*, 2020. doi: [https://doi.org/10.1016/S1473-3099\(20\)30162-6](https://doi.org/10.1016/S1473-3099(20)30162-6).
- McCoy, L. G., Smith, J., Anchuri, K., Berry, I., Pineda, J., Harish, V., Lam, A. T., Yi, S. E., Hu, S., Canadian Open Data Working Group, N.-P. I., and Fine, B. CAN-NPI: A curated open dataset of canadian non-pharmaceutical interventions in response to the global covid-19 pandemic. 2020.
- Řehůřek, R. and Sojka, P. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pp. 45–50, Valletta, Malta, May 2010. ELRA. <http://is.muni.cz/publication/884893/en>.
- Röder, M., Both, A., and Hinneburg, A. Exploring the space of topic coherence measures. In *Proceedings of the eighth ACM international conference on Web search and data mining*, pp. 399–408, 2015.
- Syed, S. and Spruit, M. Full-text or abstract? examining topic coherence scores using latent dirichlet allocation. In *2017 IEEE International conference on data science and advanced analytics (DSAA)*, pp. 165–174. IEEE, 2017.
- Teh, Y. W., Jordan, M. I., Beal, M. J., and Blei, D. M. Hierarchical dirichlet processes. 2004.
- Tuite, A. R., Fisman, D. N., and Greer, A. L. Mathematical modelling of covid-19 transmission and mitigation strategies in the population of ontario, canada. *CMAJ*, 2020. doi: 10.1503/cmaj.200476.
- Wang, C. W., Li, L., Hao, X., Guo, H., Wang, Q., Huang, J., He, N., Yu, H., Lin, X., Pan, A., Wei, S., and Wu, T. Evolving epidemiology and impact of non-pharmaceutical interventions on the outbreak of coronavirus disease 2019 in wuhan, china. *JAMA*, 2020. doi: 10.1001/jama.2020.6130.