Inferring global-scale temporal latent topics from news reports to predict public health interventions for COVID-19

Zhi Wen

McGill University

Guido Powell

McGill University

Imane Chafi

McGill University

David L Buckeridge

McGill University

Yue Li (✉ yueli@cs.mcgill.ca)

McGill University <https://orcid.org/0000-0003-3844-4865>

Article

Keywords: COVID-19, public health, non-pharmacological interventions

Posted Date: May 24th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-555402/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at Patterns on February 1st, 2022. See the published version at <https://doi.org/10.1016/j.patter.2022.100435>.

Inferring global-scale temporal latent topics from news reports to predict public health interventions for COVID-19

Zhi Wen^{1,†}, Guido Powell^{2,†}, Imane Chafi¹, David Buckeridge^{2,*} and Yue Li^{1,*}

¹School of Computer Science, McGill University

²School of Population and Global Health, McGill University

*Correspondence to david.buckeridge@mcgill.ca, yueli@cs.mcgill.ca

[†]Equal contribution

Abstract

The COVID-19 global pandemic has highlighted the importance of non-pharmacological interventions (NPI) for controlling epidemics of emerging infectious diseases. Despite the importance of NPI, their implementation has been monitored in an *ad hoc* and uncoordinated manner, mainly through the manual efforts of volunteers. Given the absence of systematic NPI tracking, authorities and researchers are limited in their ability to quantify the effectiveness of NPI and guide decisions regarding their use during the progression of a global pandemic. To address this issue, we propose 3-stage machine learning framework called EpiTopics to facilitate the surveillance of NPI by mining the vast amount of unlabelled news reports about these interventions. Building on topic modeling, our method characterizes online government reports and media articles related to COVID-19 as a mixture of latent topics. Our key contribution is the use of transfer-learning to address the limited number of NPI-labelled documents and topic modelling to support interpretation of the results. At stage 1, we trained a modified version of the unsupervised dynamic embedded topic model (DETM) on 1.2 million international news reports related to COVID-19. At stage 2, we used the trained DETM to infer topic mixture from a small set of 2000 NPI-labelled WHO documents as the input features for predicting NPI labels on each document. At stage 3, we supply the inferred country-level temporal topics from the DETM to the pretrained document-level NPI classifier to predict country-level NPIs. We identified 25 interpretable topics, over 4 distinct and coherent COVID-related themes. These topics contributed to significant improvements in predicting the NPIs labelled in the WHO documents and in predicting country-level NPIs. Together, our work lay the machine learning methodological foundation for future research in global-scale surveillance of public health interventions. The EpiTopics code is available at GitHub: <https://github.com/li-lab-mcgill/covid-npi>.

1 Introduction

It has long been understood that organized community efforts are required to control the spread of infectious diseases [1]. These efforts, called public health interventions, include social or non-pharmaceutical measures to limit the mobility and contacts of citizens and pharmaceutical interventions to prevent and limit the severity of infections. Non-pharmaceutical interventions (NPI) are not easily evaluated through experimental studies [2], so evidence regarding the effectiveness of these public health interventions is usually obtained through observational studies. However, evaluation and even monitoring NPI is challenging because their use is not recorded consistently within or across countries.

In the early stages of an emerging infectious disease, such as SARS-CoV-2, NPI tend to be particularly important due to the lack of specific pharmaceutical interventions. At the outset of the COVID-19 pandemic, recognizing the absence of systems for recording NPI, multiple groups initiated projects to track the use of interventions for COVID-19 around the world. These projects have relied on the manual efforts of volunteers to review digital documents accessible online with minimal coordination across projects [3]. This approach to monitoring interventions is difficult to sustain and scale to other infectious diseases, and it has produced information about NPI that is of variable reliability and quality [4].

Given that changes in the status of NPI are usually described in digital online media (e.g., government announcements, news media), machine learning methods have the potential to support the monitoring of NPI. However, the application of machine learning methods to this task is complicated by the need to generate interpretable results, the limited amount and quality of labelled data, and the vast volume of heterogeneous online media.

In this paper, we present a machine learning framework that uses online media to monitor changes in the status of different NPI to control COVID-19 at a global scale. This type of surveillance can inform situational awareness, support policy evaluation, and guide evidence-based decisions about the use of public health interventions to control the COVID-19 pandemic. More generally, our overarching goal is to demonstrate the feasibility of a machine learning approach to surveillance of public health interventions.

To address the machine learning challenges inherent in this problem, we propose a three-step framework called EpiTopics (**Fig. 1**). At the center of our approach is the use of probabilistic topic modelling, which facilitates the interpretation of our results and is well-suited to the noisy nature of online media. Our approach also uses transfer learning, allowing us to exploit a small amount of labelled data to make inference using a larger amount of unlabelled data. By classifying online media news into NPI categories based on interpretable model-inferred distributions, we aim to automate many of the labour-intensive aspects of NPI tracking and enable systematic monitoring of interventions across larger geographical and conceptual scopes than are possible with currently used manual methods.

Our proposed approach is novel both in terms of the public health focus and the machine learning methods. Other researchers have applied supervised learning to online media [3] and Wikipedia articles [5] to identify COVID-19 NPI, but the ‘black box’ nature of the models has

made it difficult to interpret results and model parameters. Our approach also has less demand for labelled data as it exploits large-scale unlabeled data via unsupervised learning and transfer learning. Topic modelling, such as Latent Dirichlet Allocation (LDA) [6], has been used for surveillance of online media to detect epidemics [7–10] and to characterize news reports [11] and social media [12] in the context of COVID-19 or other outbreaks. Our work differs from these applications due to a focus on tracking NPI as opposed to detecting epidemics and by adopting a dynamic topic model that goes beyond the document-level to support inference at the country level.

2 Results

2.1 EpiTopics model overview

We developed, EpiTopics (**Fig. 1**), a machine learning framework for surveillance of non-pharmacological interventions (NPI) used to control the COVID-19 pandemic. The intended application of the framework is to predict each week, based on unlabelled news documents about COVID-19, any change in the status of multiple NPI at the country level. To this end, EpiTopics implements a three-stage machine learning strategy, which is described in general terms below and in greater detail in the **Methods** section.

2.2 Interpreting COVID-19 topics learned from the AYLIEN dataset

At the first stage, we sought to learn a set of unbiased latent topic distributions by mining a large corpus of news articles about COVID-19, but without NPI labels. We used an open-access dataset of articles from AYLIEN, which included 1.2 million news articles about COVID-19 from 42 countries recorded from November 1, 2019 to July 31, 2020 (**Section 4.1**). All articles were related to COVID-19, but not necessarily related to NPI. Although NPI labels were not available for these articles, in contrast to the smaller datasets with NPI labels (described next), this corpus was suitable for learning diverse semantic topics related to COVID-19 due to its size and coverage. To capture country-dependent topic dynamics, we developed an unsupervised model called MixMedia [13], which is adapted from the dynamic embedding topic model (DETM) [14] (**Fig. 1a**; **Methods**, **Discussion**). We experimented with different numbers of topics and chose 25 topics based on topic quality [14] (Section 4.6; Appendix S1; **Supplementary Fig. S1**).

We annotated the 25 topics learned from the AYLIEN data at Stage 1 based on the words with the highest probability under each topic (**Fig. 2**) and then grouped these topics into four themes: (1) central issues of COVID-19; (2) broader social impacts; (3) specific locations; and, (4) indirectly related issues.

Under theme 1, several topics addressed *central health-related issues* of pandemic. For instance, the most frequent words in topic M3 - COVID Burden included deaths and confirmed

cases; topic M12 - Healthcare/Hospitals included health, staff, doctors, beds; topic M22 - Testing/Tracing included tested, positive, symptoms; topic M25 - Vaccine/Research included vaccine, trial, drug; topic M17 - Masks included personal protective measures such as wearing masks, washing hands, and social distancing; and topic M9 - Global Pandemic included health, pandemic, and world.

Under theme 2, topics were associated with *broader social impacts*. For example, topic M20 - Schools reflected changes in education including words such as school, online, classes, exams; topic M14 - International Travel focused on changes in international travel, with top words like flight and passenger; topic M10 - Government Restrictions included lockdowns and working from home; topic M23 - Time/Home was related to spending time at home and family; topic M8 - Food & Delivery discussed restaurant, delivery, and grocery stores; topic M21 - Sports Leagues was associated with sports teams and events; and, topic M6 - Entertainment was related to the impacts on entertainments such as film and music. There were also topics related to financial impacts such as topic M1 - Business/Work on businesses and companies, topic M5 - Economy/Markets on stock markets, and topic M7 - Finance/Economy on impacts on unemployment, income, government financial packages.

Under theme 3, some topics focus on impacts of the virus on *specific countries*. Topic M2 - COVID China focuses on the pandemic in China and particularly in Wuhan city, topic M4 - COVID USA/Local focused on the United States and its specific states and cities, and topic M13 - COVID India focused on the pandemic in India.

Under theme 4, a few topics were not directly related to COVID-19, although they were sometimes discussed in the context of COVID-19 as they reflected events occurring during the pandemic. For example, topic M18 - Outdoors discusses water, space, and parks, which are activities with reduced risk of COVID-19; topic M11 - Government/State was associated with states or governments; topic M15 - Internet/Social Media focused on social media and the internet, which are crucial for distributing COVID-19 information; topic M16 - Legal involves court, law, and order; topic M19 - Police Brutality/Protests was associated with police brutality which sparked large-scale protests during the pandemic; and topic M24 - US Election was associated with United States presidential election, for which the pandemic was a central issue.

Taken together, these observations indicate that the 25 topics inferred by our EpiTopics model from the AYLIEN news dataset represent diverse aspects of the COVID-19 pandemic. These inferred topics provide a rich foundation against which we can characterize online media reports about COVID-19.

2.3 Country topic dynamics

Continuing exploring the results generated by EpiTopics at stage 1, we examined temporal patterns at the country level in the posterior distributions of topic popularity (η) inferred by our EpiTopics model. These patterns reflect changes over time in the COVID-19-related media news as the pandemic unfolded in each country (**Fig. 3**). As expected, media news in January and February from all of the 42 countries are predominantly about topic M12 China COVID. In

March, the topic popularity began to diverge for different countries. Interestingly, the change in topic popularity for many countries meaningfully reflected the progression of the pandemic within those countries. For instance, topic M11 (India COVID) started to rise at the end of March and stayed high for India media news. Asian countries, South Korea, China, and Japan tended to focus more on Topic M12 China COVID even after February whereas the same topic dropped more sharply after February for countries in other global regions. Spain was the only country with high topic focus on sporting leagues, perhaps due to the strong impacts of COVID on these activities in Spain or possibly reflecting sampling bias from the data, which was focused on English-language media. In contrast to other countries, we observed more dramatic changes in topic dynamics for the United States with the focus shifting from topic M12 China COVID in February to COVID in local regions of the US, then to multiple topics around June including topic M3 for the global pandemic, topic M17 US election, and the more general topic M15 for staying home. We caution that part of the reasons could be due to the greater amount of articles collected for the United States compared to other countries.

2.4 Predicting document-level interventions

At the second stage, we trained a linear classifier to predict the 15 NPI labels mentioned in news articles in the ~ 2000 WHO dataset (**Section 4.1**) using their topic mixture (**Fig. 1b**). Each article in this dataset reported a change in the status of an NPI within a country and was labelled with one or more NPIs. To build the classifier, we first inferred the topic mixture of the WHO labelled articles using the MixMedia trained on the AYLIEN data in the first stage. We then projected the topic mixture (θ) by the topic embedding (α – also learned from Stage 1) to yield a document embedding input matrix. This α -embedding allowed for higher dimensional representation of concepts relevant to NPI for better prediction performance. To predict the NPI mentioned in an article, we fit an embedding-by-NPI linear regression coefficient matrix (ω). Notably, multiplying the topic-embedding projection matrix (α) by the embedding-by-NPI coefficient (ω) gives the topic-by-NPI matrix, which we explored in a subsequent analysis.

We compared the prediction performance of EpiTopics and two baseline methods using the macro AUPRC and the weighted AUPRC (**Table 1**). The baseline methods predicted document NPIs from bag-of-words (BOW) representations of documents using a linear classifier and a two-layered feed-forward network, respectively. Overall, we observed that our method outperformed the two baseline methods by a large margin, highlighting the benefit of incorporating topic mixtures, although we observed broad range of prediction accuracy values across individual NPIs (**Fig. 4a**). We achieved the best prediction performance for “Financial Packages”, “School Measures”, and “International Travel Measures” (AUPRC > 0.6). In contrast to the baseline prediction models that operate directly on the raw word count data, our method is also interpretable by examining the learned classifier weights which quantify the associations between the 25 topics and the 15 NPIs (**Fig. 4b**). For example, we observe that topic M20 “Schools” is strongly associated with the NPI “School measures”, topic M5 “Economy/Markets” is associated with “Financial packages”, and topic M21 “Sports League” is associated with

"Gatherings, businesses and services".

2.5 Predicting country-level interventions

At the third stage, we performed transfer learning to predict the 15 NPI labels at the country-level for 42 countries by transferring weights learned from Stages 1 and 2 to Stage 3. We aimed to predict whether there is *any change* in the status of a specific NPI at the country level from online media reports (**Fig. 1c**). Here a change in NPI status includes implementation, modification, and phase-out. We do not distinguish the types of change in our current model due to the limitations of the labelled data.

For input features, we used the α -projected country-dependent topic posterior ($\eta\alpha$) inferred from the AYLIEN dataset at Stage 1. The weights (ω) learned at Stage 2 for the document-level NPI prediction were reused as the NPI-topic classifier weights. These weights were first used directly to predict NPIs at the country level without further training, an approach which we call *zero-shot* transfer. To improve upon this model, we also fine-tuned the topic-NPI coefficients ω to tailor the model towards country-level NPI prediction. We call this approach the *Fine-tuned* model. Lastly, as a baseline model, we also trained a classifier *from scratch* by randomly initializing and fitting the topic-NPI coefficients ω while fixing the same α -projected country-dependent topic posterior inferred from Stage 1. We call this approach as the *From-scratch* model.

Overall, zero-shot transfer performed significantly better than random prediction and fine-tuning the transfer learning performed significantly better than training from scratch (**Fig. 5c**; **Table 2**). The results are mostly consistent for individual NPI prediction (**Fig. 5a**; **Fig. S2**). Furthermore, the predicted probabilities of NPIs are explained by the dominant country-level topic priors and their associations with each NPI (**Fig. 6**). For instance, the high probabilities for "Financial packages" in Philippines starting from March (**Fig. 6**), match well with the true labels, and the high probabilities of this NPI are explained by the increasing probabilities of topics M5 - Economy/Markets and M7 - Finance/Economy, and their strong association with NPI "Financial packages" learned in the classifier (**Fig. 4b**, **Fig. 6**).

We observe that the prevalence of NPIs is highly correlated with individual NPIs' AUPRC scores, especially for the random baseline (**Fig. 5a** and **Fig. S4b**). This is expected as the AUPRC of purely random predictions should be equal to the proportion of the positive labels [15]. For non-random predictions, on a small dataset such as ours, it is also reasonable that the prevalence plays an important role. In addition to the prevalence of a NPI, the nature a NPI and the way it is defined and grouped can impact performance. In some groups the NPIs can be diverse (e.g. "Other measures"), making these groups less conceptually coherent than other groups and harder to predict. Also, it is possible that some NPI status changes are not recorded by the trackers and therefore missing in our dataset.

We also explored the feasibility of predicting country-level NPIs from *document-level topic mixtures* θ_d inferred from the AYLIEN dataset (**Section 4.5**). Compared with predicting from country-level topic priors (**Table 3**), we observed minor difference in performance of zero-shot

transfer and fine-tuning. With orders of magnitude more documents, training from scratch performs similarly to fine-tuning. However, directly training on country-level topic mixture as input is much more efficient than training on the document-level topic mixture as there are 1 million documents and only a few hundred country-week pairs. Therefore, compared to the document-level alternative, the results highlight the advantage of directly operating on the inferred country-topic mixture in terms of both the interpretability and the computational efficiency.

3 Discussion

In this paper, we present EpiTopics, a framework for surveillance of public health interventions using online news reports. This framework makes use of transfer-learning to address the limited amount of labelled data and topic modelling to support interpretation of the results. Transfer learning is used in two ways. First, we exploited the flexibility of the topic mixture encoder network at the document-level trained on 1.2 million unlabelled news reports related to COVID-19 at Stage 1 and applied to the 2000 NPI-labelled WHO documents at Stage 2 (**Fig. 1b**). Second, we used country-specific topics generated by the recurrent neural network encoder from Stage 1 as input features and the already trained document-level NPI linear classifier to predict the country-level NPI at Stage 3 (**Fig. 1c**). In terms of topic-modelling, our approach learned 25 interpretable topics, over 4 distinct and coherent COVID-related themes, whose country-specific dynamics appear to reflect evolving events in different countries. These topics contributed to significant improvements in predicting the NPI from documents in the WHO dataset. While country-level predictions of some interventions from the AYLIEN dataset were less accurate, the topic-NPI associations drawn from the previous step improved performance.

From a machine learning perspective, EpiTopics makes several contributions. In contrast to existing supervised learning approaches that require labelled data, our approach exploits the rich information from the vast amount of unlabelled news reports. We accomplish this by adapting the Dynamic Embedded Topic Model (DETM) framework [14]. We made two important modifications to the original DETM so that it would be more suitable for surveillance of NPI. First, we infer temporal *country-dependent* topic probabilities to capture the evolving pandemic situation within each country. Second, compared to the time-dependent topic embedding in the original DETM, we infer a set of *static* topic embedding to improve interpretation. This allows us to analyze each topic separately without keeping track of their own evolution over time within small time intervals (i.e., CDC week). Taken together, these modifications to the DETM allowed us to analyze the country-dependent topic trajectories over time by leveraging a large set of highly interpretable global topic distributions.

From a public health perspective, EpiTopics framework demonstrates the feasibility of using machine learning methods to meaningfully improve the surveillance of interventions. In the context of the COVID-19 pandemic, our methods can be used to enhance and facilitate manual efforts by research teams tracking interventions. While automated methods for screening have been used by existing COVID-19 trackers [3], the interpretable country-level topic dynamics

provided by our model offer high-level insights along with classification guideposts for analysts attempting to monitor tens of thousands of NPI events around the world. More generally, expanding our approach to new corpora could further the development of rigorous methods for surveillance of public health interventions, which are not usually subject to routine surveillance. Therefore, EpiTopics complements the current, mainly manual approaches to monitoring NPI and has the potential to be expanded to the surveillance of public health interventions in general.

In the future, this work can be extended in several directions. First, our study is focused on the time frame from the beginning of COVID-19 to July 2020. Application of the EpiTopics framework to data covering a wider time frame is possible, for example through the Fall of 2020 when many countries experienced a second wave. Since the beginning of the second wave, many interventions were implemented, and vaccines started to be approved and administered across the globe, which could affect the media coverage and NPI prediction.

Second, this study is conducted using two data sets consisting of COVID-19-related news reports. We intend to study the generalizability of EpiTopics by exploring other sources of data which may cover a broader range of countries, NPIs, tasks and documents [5, 16]. For instance, the proposed method could be integrated into existing automated surveillance systems such as Global Public Health Intelligence Network (GPHIN) [17], which monitors general media news potentially related to public health, instead of being applied exclusively to COVID-19-related news.

Third, the documents are represented as bag-of-words (BOW) for topic modelling. This representation omits rare words and stop words that are infrequent and too frequent, respectively. The order of words in sentences are also ignored. Therefore, we may lose some semantic meaning of the documents, especially when a sentence contains negation (e.g., deaths are not high). To this end, we will explore neural language models such as ELECTRA [18] and BERT [19] at the expense of computations and more sophisticated techniques such as the attention mechanisms for model interpretability [20]. Also, these neural language models operate at the sentence level, whereas our goal in this study is to model at both document and country levels. Future work can explore ways of addressing these issues and investigate their performance on the task of NPI prediction.

Fourth, we can predict future country-level NPI at the next time point based on the previous time points using the country-level topic trend and previous NPI states. This can be done in an auto-regressive framework. The challenge here is the scarce observations of positive NPI changes at the country-level over time. Auxiliary labels such as number of case counts may help improve signals.

Lastly, geolocation extraction is a challenging task [21] and we expect that some proportion of countries were misclassified in the construction of the AYLIEN media dataset. For example, the country of publication could be classified as the location of an events and some documents may discuss multiple countries. While the large size of the dataset prohibits exhaustive validation, a non-exhaustive manual inspection of a subset of documents confirm that a significant majority of them have accurate country assignments. To reduce the impact of this misclassifi-

cation on our results, we limited our analyses to the country level rather than to the sub-region level, where misclassification is more likely. We will explore regional prediction in a more focus study.

In conclusion, our current work lays the methodological foundation for global-scale surveillance of the implementation of public health interventions and we have identified directions for future research to build on this foundation. We envision that our work will inspire further research to transform the way that public health interventions are monitored with advanced machine learning approaches.

4 Methods

4.1 Datasets

WHO Supported by the World Health Organization and led by a team at the London School of Health and Tropical Medicine, the WHO Public Health and Social Measures (WHO-PHSM) dataset (<https://www.who.int/emergencies/diseases/novel-coronavirus-2019/phsm>) represents merged and harmonized data on COVID-19 public health and social measures combining databases from seven international trackers. These individual trackers used mostly similar methods of manually collecting media articles and official government reports on COVID-19 interventions. The database has de-duplicated, verified and standardized data into a unified taxonomy, geographic system and data format. We selected English-language records labeled as having URL links to static web pages that are representative of the coded information. These pages were scraped using the Beautiful Soup and Newspaper3k Python libraries. Of over 30,000 records available at the time of writing, 8,432 records had valid URLs for scraping, of which 2,049 represented unique documents (articles or reports) that aligned with the period and countries in the larger AYLIEN dataset.

To balance the need for fine-grained NPI classes and the need for sufficient data points for model development, the WHO taxonomy covering 44 different categories of NPIs has been regrouped into 15 interventions (**Table S1**). The selection of categories to group together was determined by their conceptual similarity as well as their frequency. The prevalence of the 15 NPIs at the document and country levels are shown in **Fig. S4** in Appendix S3. We split the data into a training set and a testing set with a 80%-20% ratio and tuned the hyperparameters on the training set.

AYLIEN While the WHO dataset is expert-curated and its documents have corresponding NPI labels, the documents available were too few and not representative of general media content on which a surveillance systems may operate. Therefore, in addition to the WHO dataset, we also use another dataset created by a News Intelligence Platform called AYLIEN (<https://aylien.com/resources/datasets/coronavirus-dataset>). Although this dataset has no NPI label, it has a lot more COVID-19 related news reports. In particular, the AYLIEN dataset has

1.2 million news articles related to COVID-19, spanning from November 2019 to July 2020. Half of the data points (country-time pairs) are used for training and the remaining half of testing.

4.2 Data Processing

We adopted the same data processing pipeline for both datasets. Specifically, we removed white spaces, special characters, non-English words, single-letter words, infrequent words (i.e., words appearing in fewer than 10 documents). We also removed the same stop words as in [22] as well as common non-English stop words "la", "se", "el", "na", "en", "de". For WHO documents, we obtained country information from the "SOURCE" field associated with each WHO document. For AYLIEN documents, we took the "country" sub-field under the "source" field, indicating the country the document is relevant to. After removing documents without empty source fields, we have 1.2 million documents in the AYLIEN dataset. To minimize the effect of noise in time stamps and also to reduce sparsity of the data, we reduced the temporal resolution to weeks by grouping the documents and the NPIs within the same CDC weeks (<https://wwwn.cdc.gov/nndss/downloads.html>).

In total, there are 1,176,916 documents in the AYLIEN dataset, over 42 countries and spanning 38 CDC weeks from November 2019 to July 2020. After applying the same pre-processing procedure to WHO and discarding the documents whose country or source is not seen in AYLIEN, we have 2,049 unique documents in the WHO dataset with NPI labels. There are 109377 and 8012 unique words in AYLIEN and WHO respectively. Among them, 366 unique words in WHO are not present in AYLIEN. Manual inspection confirms these 366 words are non-lexical or non-English and were discarded.

4.3 EpiTopics-Stage 1: Unsupervised dynamic embedded topic

The first stage of EpiTopics (**Fig. 1a**) is built upon our previous model called MixMedia [13], which was adapted from The Embedded Topic Model (ETM) [22] and The Dynamic Embedded Topic Model (DETM) [14]. The data generative process of EpiTopics is as follows:

1. Draw a topic proportion θ_d for a document d from logistic normal $\theta_d \sim \mathcal{LN}(\eta_{s_d}^{(t_d)}, \delta^2 \mathbf{I})$:

$$\delta_d \sim \mathcal{N}(0, I); \quad \theta_d = \text{softmax}(\delta_d) = \frac{\exp(\delta_d)}{\sum_k \exp(\delta_{kd})}$$

where s_d and t_d are indices for the country of the document d and the time at which document d was published, respectively;

2. For each token n in the document,

(a) Draw topic assignment $z_{dn} \sim \text{Cat}(\theta_d)$

(b) Draw word $w_{dn} \sim \text{softmax}(\rho^T \alpha_{z_{dn}})$

Here the time-varying topic prior $\boldsymbol{\eta}_s^{(t)}$ is a dynamic Gaussian variable, which depends on the topic prior at the previous time point of the same source:

$$p(\boldsymbol{\eta}_s^{(0)}) = \mathcal{N}(0, I) \quad \text{if } t = 0 \quad (1)$$

$$p(\boldsymbol{\eta}_s^{(t)} | \boldsymbol{\eta}_s^{(t-1)}) = \mathcal{N}(\boldsymbol{\eta}_s^{(t-1)}, \zeta^2 I) \quad \text{if } t > 0 \quad (2)$$

For the ease of interpretation, we assume that the topic embeddings $\boldsymbol{\alpha}$ is time-invariant [13]:

$$p(\boldsymbol{\alpha}_k) = \mathcal{N}(0, \gamma^2) \quad (3)$$

Inference Our model has several latent variables including the dynamic topic prior $\boldsymbol{\eta}_s$ per source s , the topic mixture per document $\boldsymbol{\theta}_d$, the topic assignment per word per document token z_{dn} , and the topic embedding $\boldsymbol{\alpha}_k$ per topic k . Word embedding ρ is treated as fixed point estimates and optimized via empirical Bayes. For the ease of inference, we marginalize the topic assignments z_{dn} in the conditional data multinomial likelihood:

$$p(\mathcal{D} | \boldsymbol{\theta}, \boldsymbol{\beta}) = \prod_{d,n} \sum_k p(w_{dn} | z_{dn} = k, \boldsymbol{\beta}_k) p(z_{dn} = k | \boldsymbol{\theta}_d) = \prod_{d,n} \sum_k \boldsymbol{\theta}_{dk}^\top \boldsymbol{\beta}_{w_{dn}k} \quad (4)$$

where $\mathcal{D} = \{\mathbf{w}_1, \mathbf{w}_1, \dots, \mathbf{w}_N\}$ denote the corpus, where $\mathbf{w}_d$ is the bag-of-word representation of document d and N is the number of documents in the corpus, and $\boldsymbol{\beta}_{vk} = \text{softmax}(\rho_v^\top \boldsymbol{\alpha}_k)$ for word $w_{dn} = v$ and topic k .

The posterior distribution of the other latent variables $p(\boldsymbol{\eta}, \boldsymbol{\alpha}, \boldsymbol{\theta} | \mathcal{D})$ are intractable. To approximate them, we took an amortized variational inference approach using a family of proposed distributions $q(\boldsymbol{\eta}, \boldsymbol{\alpha}, \boldsymbol{\theta})$ to approximate the true posterior [14]:

$$q(\boldsymbol{\eta}, \boldsymbol{\alpha}, \boldsymbol{\theta}) = \prod_d q(\boldsymbol{\theta}_d | \boldsymbol{\eta}_{s_d}^{(t_d)}, \mathbf{w}_d) \prod_s \prod_t q(\boldsymbol{\eta}_s^{(t)} | \boldsymbol{\eta}_s^{(t-1)} \tilde{\mathbf{w}}_s^{(t)}) \prod_k \prod_t q(\boldsymbol{\alpha}_k)$$

where

$$q(\boldsymbol{\theta}_d | \boldsymbol{\eta}_{s_d}^{(t_d)}, \mathbf{w}_d) = \text{softmax}(\boldsymbol{\delta}_d), \quad \boldsymbol{\delta}_d \sim \mathcal{N}(\boldsymbol{\mu}_d, \text{diag}(\sigma_d^2)) = \boldsymbol{\mu}_d + \text{diag}(\sigma_d) \mathcal{N}(0, I) \\ [\boldsymbol{\mu}_d, \log \sigma_d^2] = \text{NNET}([\boldsymbol{\eta}_{s_d}^{(t_d)}, \tilde{\mathbf{w}}_d]; \mathbf{W}_\theta) \quad (5)$$

$$q(\boldsymbol{\eta}_s^{(t)} | \boldsymbol{\eta}_s^{(t-1)}, \tilde{\mathbf{w}}_s^{(t)}) = \mathcal{N}(\boldsymbol{\mu}_s^{(t)}, \text{diag}(\mathbf{v}_s^{(t)})) = \boldsymbol{\mu}_s^{(t)} + \text{diag}(\mathbf{v}_s^{(t)}) \mathcal{N}(0, I) \\ [\boldsymbol{\mu}_s^{(t)}, \mathbf{v}_s^{(t)}] = \text{LSTM}([\boldsymbol{\eta}_s^{(t-1)}, \tilde{\mathbf{w}}_s^{(t)}]; \mathbf{W}_\eta) \quad (6)$$

$$q(\boldsymbol{\alpha}_k) = \mathcal{N}(\mathbf{m}_k, \mathbf{v}_k^2) \quad (7)$$

Here using the bag-of-word representation, $\mathbf{w}_d$ denotes a $V \times 1$ vector of the word frequency of document d over the vocabulary of size V ; $\tilde{\mathbf{w}}_d$ is the normalized word frequencies; $\tilde{\mathbf{w}}_s^{(t)}$ denotes average word frequency at time t for source s .

Using the variational autoencoder framework [23], the function $\text{NNET}(\mathbf{x}; \mathbf{W}_\theta)$ is a feed-forward neural network parametrized by $\mathbf{W}_\theta$; $\text{LSTM}(\mathbf{x}; \mathbf{W}_\eta)$ is a Long Short Term Memory

(LSTM) network [24] parametrized by $\mathbf{W}_\eta$. Because of the Gaussian properties, we use the re-parameterization trick [23] to stochastically sample the latent variable θ_d , $\eta_s^{(t)}$, and α_k from their means with added Gaussian noise weighted by their variances as shown in Equations (5), (6), and (7), respectively.

To learn the above variational parameters $\Theta_q = \{\mathbf{W}_\theta, \mathbf{W}_\eta, \mathbf{m}_\alpha, \mathbf{v}_\alpha\}$, we optimize the evidence lower bound (ELBO), which is equivalent to minimizing the Kullback-Leibler (KL) divergence between the true posterior and the proposed distribution $KL(q(\Theta)|p(\Theta|\mathcal{D}))$:

$$ELBO = \mathbb{E}_q[\log p(\mathcal{D}|\boldsymbol{\eta}, \boldsymbol{\alpha}, \boldsymbol{\theta})] + KL[p(\boldsymbol{\eta}, \boldsymbol{\alpha}, \boldsymbol{\theta})|q(\boldsymbol{\eta}, \boldsymbol{\alpha}, \boldsymbol{\theta})] \quad (8)$$

We optimize ELBO with respect to the variational parameters amortized variational inference [14, 23, 25, 26]. Specifically, we sample the latent variables from $q(\boldsymbol{\eta}|\tilde{w})$ (6), $q(\boldsymbol{\alpha})$ (7), $q(\boldsymbol{\theta}|\tilde{w})$ (5) based on a minibatch of data. We then use those samples as the noisy estimates of the variational expectation for the ELBO (8). The ELBO is optimized with gradient updates using Adam [27].

Selecting the topic number based on topic quality We use topic quality to select the best number of topics. Topic quality is calculated as the product of topic diversity and topic coherence. Topic diversity is defined as the percentage of unique words in the top 25 words across all topics, and topic coherence is defined as the average point-wise mutual information of the top-10 most likely words under each topic:

$$TC = \frac{1}{K} \sum_{k=1}^K \frac{1}{45} \sum_{i=1}^{10} \sum_{j=i+1}^{10} f(w_i, w_j) \quad (9)$$

Implementation EpiTopics-Stage1 or unsupervised MixMedia [13] is implemented and trained with PyTorch 1.5.0. We used 300-dimensional word and topic embeddings. We implemented the neural network in Eq (5) with a two-layer feed-forward network, with a hidden size of 800 and ReLU activation. We implemented the LSTM in Eq (6) with a 3-layer LSTM with a hidden size of 200 and a dropout rate of 0.1. We used a learning rate of 10^{-3} , reducing to 1/4 of the previous value if validation perplexity is not reduced within the last 10 epochs, to a minimal learning rate of 10^{-7} . The model was trained for 400 epochs with a batch size of 128. The number of topics 25 was chosen based on the best topic quality, as shown in **Fig. S1**.

4.4 EpiTopics-Stage 2: Predicting document-level interventions

At this stage, we use the unsupervised MixMedia trained at stage 1 to generate the topic mixtures θ_d for WHO document d from country s_d observed at the t_d^{th} CDC week:

$$\boldsymbol{\mu}_d, \log \boldsymbol{\sigma}_d^2 = NNET([\hat{\boldsymbol{\eta}}_{s_d}^{(t_d)}, \tilde{\mathbf{w}}_d]) \quad (10)$$

$$\boldsymbol{\theta}_d = softmax(\boldsymbol{\mu}_d) \quad (11)$$

where $NNET$ is the trained encoder of MixMedia in Eq (5), $\tilde{\mathbf{w}}_d$ is the normalized word frequencies of document d , and $\hat{\boldsymbol{\eta}}_{s_d}^{(t_d)} \equiv \boldsymbol{\mu}_{s_d}^{(t_d)}$ is the expectation of the country-level topic mixture for country s_d at time t_d , which is produced by the trained LSTM of MixMedia in Eq (6).

We then fit a linear classifier to predict the 15 NPIs for document d using its topic mixture $\boldsymbol{\theta}_d$:

$$\hat{\mathbf{y}}_d = \sigma(\boldsymbol{\theta}_d \boldsymbol{\alpha} \boldsymbol{\omega} + \mathbf{b}) \quad (12)$$

where $\sigma(x) = 1/(1 + e^{-x})$ is the sigmoid function, $\hat{\mathbf{y}}_d$ is the vector of predicted probabilities of the 15 NPIs, $\boldsymbol{\omega}$ is the linear weights and $\mathbf{b}$ is the bias vector.

The loss function here is cross-entropy plus the L2-regularization:

$$L = \sum_d -\mathbf{y}_d \log \hat{\mathbf{y}}_d - (1 - \mathbf{y}_d) \log(1 - \hat{\mathbf{y}}_d) + \lambda \sum_{c=1}^{C=15} \|\boldsymbol{\omega}_l\|^2 \quad (13)$$

We chose a linear classifier here for the ease of interpretability and also due to small training sample size. The associations between topics and NPIs can be obtained by inspecting the inner product of the topic embedding and regression coefficients: $\boldsymbol{\alpha} \boldsymbol{\omega}$, which gives a $K \times 15$ matrix indicating the associations between the K topics and 15 NPIs.

Implementation The linear classifier was implemented with PyTorch 1.5.0. We used Adam optimizer with a learning rate of 6×10^{-3} and L2-penalty λ of 10^{-5} . We trained the classifier for 400 epochs with a batch size of 512. To mitigate class imbalance over the 15 NPIs, we adjusted the weights for the per-NPI loss by up-weighting the loss of minority NPI class labels and down-weight the loss of majority NPI class labels:

$$\text{weight of class } c = \frac{\# \text{ of positive labels in total} / \# \text{ of positive labels in class } c}{\sum_c (\# \text{ of positive labels in total} / \# \text{ of positive labels in class } c)} \quad (14)$$

4.5 EpiTopics-Stage 3: predicting country-level interventions

Prediction from country-level topic priors After the linear classifier for document-level intervention prediction is trained, its weights are transferred to the prediction task for country-level NPI labels:

$$\hat{\mathbf{y}}_s^{(t)} = \sigma(\hat{\boldsymbol{\eta}}_s^{(t)} \hat{\boldsymbol{\alpha}} \hat{\boldsymbol{\omega}} + \hat{\mathbf{b}}) \quad (15)$$

For zero-shot transfer, we directly use the *country-level topic mixture* $\boldsymbol{\eta}$ and topic embedding $\boldsymbol{\alpha}$ from learned Stage 1 and the trained linear classifier coefficients $\hat{\boldsymbol{\omega}}$ and biases $\hat{\mathbf{b}}$ from Stage 2.

For the fine-tune, the classifier weights are further updated for predicting country-level interventions, i.e. initialize classifier weights with the learned parameters $\hat{\boldsymbol{\omega}}$ and further adjust it to predict country-level NPIs. For the from-scratch model, we initialize the classifier weights by sampling from standard Normal.

Prediction from Document-level Topic Mixtures We also experimented with predicting country-level NPIs using *documents' topic mixtures*. Specifically, the country-level NPI probabilities are calculated by taking the average of the predictions of the document-level NPI that for corresponding country and time:

$$\hat{\mathbf{y}}_s^{(t)} = \frac{1}{|\mathcal{D}_s^{(t)}|} \sum_{d \in \mathcal{D}_s^{(t)}} \sigma(\boldsymbol{\theta}_d \boldsymbol{\alpha} \boldsymbol{\omega} + \mathbf{b}) \quad (16)$$

where $\mathcal{D}_s^t \subset \mathcal{D}$ is the set of documents associated with country s and time point t . Therefore, this method can only be trained on predicting NPIs for the country and time pairs, where there is at least one document observed. Similarly to the method above, predictions can be made in a zero-shot transfer way, or the classifier weights are updated from the learned weights or random weights for the fine-tuned and from-scratch models, respectively.

4.6 Experiments

Baselines We omitted the comparison with other models at EpiTopics-Stage1 and refer the readers to our MixMedia work for comparing the topic quality with baseline topic models [13]. For document NPI predictions (EpiTopics-Stage2), we compared against using linear and feed-forward network on bag-of-word (BOW) input features as the baselines.

The BOW + linear baseline is trained with learning rate 8×10^{-5} and weight decay 0.002 for 100 epochs. The BOW + feed-forward baseline is trained with learning rate 10^{-4} and weight decay 0.002 for 100 epochs. The feed-forward network is implemented with two fully connected layers with a hidden size of 256 and ReLU activation. Fine-tuning used learning rate 6×10^{-2} and weight decay 1×10^{-3} with a batch size of 1024, and training from scratch used learning rate 2×10^{-3} and weight decay of 2×10^{-1} with a batch size of 1024.

For country-level NPI prediction (EpiTopics-Stage3), we compared against training from scratch in addition to random model generated weights as baselines. For the random model, we randomly initialize the country-specific topic mixture $\boldsymbol{\eta}_s^t$, topic embedding $\boldsymbol{\alpha}$ and/or classifier parameters $\{\boldsymbol{\omega}, \mathbf{b}\}$ and do not update their weights. To measure the effect of different levels of randomness, we experimented with three random baselines as follows:

- Random eta, alpha and omega: setting $\boldsymbol{\eta}_s^t$, $\boldsymbol{\alpha}$ and $\{\boldsymbol{\omega}, \mathbf{b}\}$ to random numbers
- Random eta and omega: setting $\boldsymbol{\eta}_s^t$, and $\{\boldsymbol{\omega}, \mathbf{b}\}$ to random numbers
- Random omega: setting $\{\boldsymbol{\omega}, \mathbf{b}\}$ to random numbers

Intuitively, the first random baseline preserves no information gained from document-level prediction, and thus represents complete randomness, while the other two preserve some level of information and thus can potentially perform better than the first baseline.

Evaluation Since NPI prediction is a multi-label classification problem, and that different NPIs can have different optimal thresholds, we used weighted AUPRC and macro AUPRC as evaluation metrics. We computed the Area Under the Precision-Recall Curve (AUPRC) for each NPI individually, and either took the unweighted average of them (macro AUPRC) or the average weighted by each classes' prevalence (weighted AUPRC).

5 Figures

Stage 1. Learning COVID-19 topic and word embeddings from 1.2 million AYLIEN dataset over 42 countries

Stage 2. Supervised transfer-learning for document-level NPI prediction

Stage 3. Transfer learning from document-level to country-level NPI retrospective prediction

Figure 1: EpiTopics model overview. a. Unsupervised learning of COVID-19 topic and word embeddings from 1.2 million AYLIEN dataset over 42 countries. For each country, we extracted a set of news articles related to COVID-19 observed from November 1, 2019 to July 31, 2020 from AYLIEN. For each document (e.g., $w_{CAN,t,d}$ for document d published at time t from Canada), we model them as a bag of words by inferring their topic mixture (θ_d) and global topic embedding (α) and word embedding (ρ). b. Supervised transfer-learning for document-level NPI prediction. We used the trained topic mixture encoder from Stage 1 to infer the topic mixture of the 2000 WHO news reports. The resulting the topic mixture are then used as the input features to predict NPI labels in a logistic regression model. c. Supervised transfer-learning from document-level to country-level NPI prediction. Here we used country-specific topic trajectories η inferred at stage 1 as the input features to predict country-level NPI at each time point in a pretrained logistic regression model, where the linear coefficients were already fit at Stage 2 for the document-level NPI prediction task.

Figure 2: Learned topics and the top words under each topic. The size of the words reflects their probability values. The background colors indicate the themes we gave to the topics.

Figure 3: Temporal topics progression in example countries from January 2020 to July 2020. Different topics are represented with different colors, and the vertical span of each color block reflects the probability value of that topic, in that country at that specific time. The topics are ranked in descending order vertically. To avoid cluttering the plot, only the top 5 topics are displayed for each country.

Figure 4: Document-level NPI prediction results. (a) AUPRC scores on individual NPI predictions. We compared 3 methods: *BOW + linear* uses bag of word (BOW) features to predict NPI with a linear classifier; *BOW + feed-forward* uses a feed-forward neural network to predict NPI using the BOW features as inputs; *Document topic mixture + linear*: our EpiTopics that uses the inferred topic mixture from the trained encoder at Stage 1 to predict NPI by fitting a linear classification model. (b) Topic-NPI associations learned by the linear classifier displayed in heatmap with red, white, and blue indicating positive association, no association, and negative association, respectively.

a. AUPRCs of individual NPIs. Values in brackets are proportion of positive labels.

b. Precision-recall curve of “Financial packages”

c. AUPRC distributions across random seeds

Figure 5: Country-level NPI predictions. (a) AUPRC scores on individual NPI predictions at the country-level. Random, zero-shot transfer, from-scratch, and fine-tuned are methods that predict NPI using random features, pre-trained linear coefficients from Stage 2 on document-level NPI predictions, training the linear coefficients from random initialization, and fine-tuning the pre-trained linear coefficients from Stage 2, respectively. The numbers in the brackets indicate the proportion of the positive labels for that NPI, which are positively correlated with the corresponding AUPRC. (b) Precision-recall curve of different methods on predicting “Financial packages”. (c) Distributions of the weighted AUPRC scores and macro AUPRC scores over the 15 NPI predictions for the four methods across 100 repeated experiments with random initializations.

Figure 6: Example of topic dynamics and NPI predictions (using the *fine-tuned* method) for select countries. In the top panel, the predicted probabilities over time of "Financial packages" in the Philippines, "School measures" in the UK and "Stay-at-home orders" in Italy are represented by the dashed lines, The dots represents the date of the NPI. These examples were selected to highlight a range of good to poor prediction accuracy (left to right). Predictions are informed by temporal topic dynamics within the AYLIEN dataset, represented in the middle panel. The bottom panel shows weights of each topic for a given NPI prediction, based on topic-NPI associations (ω) learned via predicting NPIs from media articles in the WHO dataset. Topics with ω associations below 0 are not shown.

6 Tables

	Weighted AUPRC	Macro AUPRC
BOW + linear	0.165 (0.001)	0.108 (0.001)
BOW + feed-forward	0.149 (0.022)	0.099 (0.007)
Document topic mixtures + linear	0.408 (0.001)	0.316 (0.001)

Table 1: Area under the precision-recall curve (AUPRC) scores for document-level NPI prediction. The AUPRC scores are computed on individual NPIs, and then averaged without weighting (macro AUPRC) or weighted by NPIs’ prevalence (weighted AUPRC). Both BOW+linear and BOW+feed-forward use the normalized word vector (i.e., bag of words or BOW) for each document to predict NPI label. All methods are each repeated 100 times with different random seeds. Values in the brackets are standard deviations over the 100 experiments.

	Weighted AUPRC	Macro AUPRC
Random eta, alpha and omega	0.394 (0.007)	0.307 (0.006)
Random eta and omega	0.394 (0.007)	0.307 (0.006)
Random omega	0.398 (0.014)	0.312 (0.011)
Zero-shot transfer	0.427 (0.001)	0.336 (0.001)
Trained from scratch	0.446 (0.001)	0.358 (0.001)
Fine-tuned	0.488 (0.000)	0.390 (0.000)

Table 2: Area under the precision-recall curve (AUPRC) scores for country-level NPI prediction. Random baselines are each repeated 1000 times with different random seeds, and the rest are each repeated 100 times with different random seeds. Values in the brackets are standard deviations over the repeated experiments.

	Topic level	Weighted AUPRC	Macro AUPRC
Random eta, alpha and omega	Country	0.411 (0.007)	0.323 (0.006)
Random eta and omega	Country	0.411 (0.007)	0.323 (0.006)
Random omega	Country	0.413 (0.017)	0.326 (0.013)
Zero-shot	Country	0.451 (0.001)	0.352 (0.001)
	Document	0.448 (0.001)	0.350 (0.001)
From scratch	Country	0.455 (0.001)	0.367 (0.001)
	Document	0.506 (0.000)	0.398 (0.000)
Fine-tuned	Country	0.503 (0.000)	0.404 (0.000)
	Document	0.506 (0.000)	0.397 (0.000)

Table 3: AUPRC scores for country-level NPI prediction from topics at document and country level. Values in the brackets are standard deviations. Random baselines are each repeated 1000 times with different random seeds, and the rest are each repeated 100 times with random seeds.

References

1. C-EA Winslow. The untilled fields of public health. *Science*, pages 23–33, 1920.
2. Robert William Sanson-Fisher, Billie Bonevski, Lawrence W Green, and Cate D’Este. Limitations of the randomized controlled trial in evaluating population-based health interventions. *American journal of preventive medicine*, 33(2):155–161, August 2007.
3. Cindy Cheng, Joan Barceló, Allison Spencer Hartnett, Robert Kubinec, and Luca Messerschmidt. Covid-19 government response event dataset (coronanet v. 1.0). *Nature human behaviour*, 4(7):756–768, 2020.
4. Jan M. Brauner, Sören Mindermann, Mrinank Sharma, David Johnston, John Salvatier, Tomáš Gavenčiak, Anna B. Stephenson, Gavin Leech, George Altman, Vladimir Mikulík, Alexander John Norman, Joshua Teperowski Monrad, Tamay Besiroglu, Hong Ge, Meghan A. Hartwick, Yee Whye Teh, Leonid Chindelevitch, Yarin Gal, and Jan Kulveit. Inferring the effectiveness of government interventions against covid-19. *Science*, 371(6531), 2021.
5. Parthasarathy Suryanarayanan, Ching-Huei Tsou, Ananya Poddar, Diwakar Mahajan, Bharath Dandala, Piyush Madan, Anshul Agrawal, Charles Wachira, Osebe Mogaka Samuel, Osnat Bar-Shira, Clifton Kipchirchir, Sharon Okwako, William Ogallo, Fred Otieno, Timothy Nyota, Fiona Matu, Vesna Resende Barros, Daniel Shats, Oren Kagan, Sekou Remy, Oliver Bent, Pooja Guhan, Shilpa Mahatma, Aisha Walcott-Bryant, Divya Pathak, and Michal Rosen-Zvi. AI-assisted tracking of worldwide non-pharmaceutical interventions for COVID-19. *Scientific Data*, 8(1):94, 2021.
6. David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003.
7. Dror Walter and Yotam Ophir. News frame analysis: An inductive mixed-method computational approach. *Communication Methods and Measures*, 13(4):248–266, 2019.
8. Yotam Ophir. Coverage of epidemics in american newspapers through the lens of the crisis and emergency risk communication framework. *Health security*, 16(3):147–157, 2018.
9. Yunli Wang and Cyril Goutte. Real-time Change Point Detection using On-line Topic Models. *COLING*, 2018.
10. Saurav Ghosh, Prithwish Chakraborty, Elaine O Nsoesie, Emily Cohn, Sumiko R Mekaru, John S Brownstein, and Naren Ramakrishnan. Temporal Topic Modeling to Assess Associations between News Trends and Infectious Disease Outbreaks. *Scientific Reports*, pages 1 – 12, 01 2017.

11. William Poirier, Catherine Ouellet, Marc-Antoine Rancourt, Justine Béchar, and Yannick Dufresne. (un) covering the covid-19 pandemic: Framing analysis of the crisis in canada. *Canadian Journal of Political Science/Revue canadienne de science politique*, 53(2):365–371, 2020.
12. Xuehua Han, Juanle Wang, Min Zhang, and Xiaojie Wang. Using social media to mine and analyze public opinion related to covid-19 in china. *International Journal of Environmental Research and Public Health*, 17(8):2788, 2020.
13. Yue Li, Pratheeksha Nair, Zhi Wen, Imane Chafi, Anya Okhmatovskaia, Guido Powell, Yannan Shen, and David Buckeridge. Global surveillance of covid-19 by mining news media using a multi-source dynamic embedded topic model. In *Proceedings of the 11th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics*, pages 1–14, 2020.
14. Adji B Dieng, Francisco JR Ruiz, and David M Blei. The dynamic embedded topic model. *arXiv preprint arXiv:1907.05545*, 2019.
15. Takaya Saito and Marc Rehmsmeier. The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets. *PloS one*, 10(3):e0118432, 2015.
16. Jan M Brauner, Sören Mindermann, Mrinank Sharma, David Johnston, John Salvatier, Tomáš Gavenčiak, Anna B Stephenson, Gavin Leech, George Altman, Vladimir Mikulík, Alexander John Norman, Joshua Teperowski Monrad, Tamay Besiroglu, Hong Ge, Meghan A Hartwick, Yee Whye Teh, Leonid Chindelevitch, Yarin Gal, and Jan Kulveit. Inferring the effectiveness of government interventions against COVID-19. *Science (New York, NY)*, 371(6531), February 2021.
17. Marie Dion, Philip AbdelMalik, and Abba Mawudeku. Big data: big data and the global public health intelligence network (gphin). *Canada Communicable Disease Report*, 41(9):209, 2015.
18. Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. ELECTRA: Pre-training text encoders as discriminators rather than generators. In *ICLR*, 2020.
19. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. *Proceedings of the 2019 Conference of the North*, 2019.
20. Sarthak Jain and Byron C Wallace. Attention is not explanation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3543–3556, 2019.

21. Muhammad Nur Yasir Utomo, Teguh Bharata Adji, and Igi Ardiyanto. Geolocation prediction in social media data using text analysis: A review. In *2018 International Conference on Information and Communications Technology (ICOIACT)*, pages 84–89. IEEE, 2018.
22. Adji B Dieng, Francisco JR Ruiz, and David M Blei. Topic modeling in embedding spaces. *Transactions of the Association for Computational Linguistics*, 8:439–453, 2020.
23. Diederik P Kingma and Max Welling. Auto-encoding variational bayes. 2014.
24. Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
25. Rajesh Ranganath, Sean Gerrish, and David Blei. Black box variational inference. In *Artificial intelligence and statistics*, pages 814–822. PMLR, 2014.
26. Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. Stochastic variational inference. *Journal of Machine Learning Research*, 14(5), 2013.
27. D Kingma and J Ba. Adam: A method for stochastic optimization in: Proceedings of the 3rd international conference for learning representations (iclr’15). *San Diego*, 2015.

Supplementary Information

Inferring global-scale temporal latent topics from news reports to predict public health interventions for COVID-19

Zhi Wen^{1,†}, Guido Powell^{2,†}, Imane Chafi¹, David Buckeridge^{2,*}, Yue Li^{1,*}

¹School of Computer Science, McGill University, ²School of Population and Global Health, McGill University

[†]Equal contribution ^{*}Correspondence: jian.tang@hec.ca, yueli@cs.mcgill.ca

S1 Experiment Details

The number of topics 25 is chosen based on the best topic quality, as shown in **Fig. S1**.

Figure S1: Topic quality versus the number of topics. Using 25 topics leads to the highest topic quality, and that number is used throughout this work. The x-axis is log-scaled.

S2 Precision-Recall Curves of All NPIs

The precision-recall curves for the all NPIs are in **Fig. S2**.

Figure S2: Country-level precision-recall curves of all 15 NPIs over 42 countries. Similar to **Fig. 5b**, each subfigure corresponds to one NPI, and together they correspond to all the NPIs in **Fig. 5a**. The predictions are made by a linear classifier on the country-specific topic mixture, and different methods in comparison learn the classifier weight differently as described in **Section 4.5**. In particular, Random, zero-shot transfer, from-scratch, and fine-tuned are methods that predict NPI using random features, pre-trained linear coefficients from Stage 2 on document-level NPI predictions, training the linear coefficients from random initialization, and fine-tuning the pre-trained linear coefficients from Stage 2, respectively. The numbers in the brackets indicate the proportion of the positive labels for that NPI, which are positively correlated with the corresponding AUPRC.

S3 Dataset Details

The number of positive labels at the document-level in the WHO dataset and at the country-level are shown in **Fig. S4**.

Figure S3: Number of documents of each country in the AYLIEN dataset. The y-axis is log-scaled.

(a) WHO documents

(b) Country-time pairs

Figure S4: The number of positive labels at the document and country levels. The x-axis is log-scaled.

NPI label	WHO code	WHO Category	WHO Subcategory	WHO Measure	WHO PHSM definition (where available)
Communications and engagement	8.4	Other measures	Communications and engagement		
	8.4.1			General public awareness campaigns	
	8.4.2			Other communications	
Detecting and isolating cases	3.1.1	Surveillance and response measures	Detecting and isolating cases	Passive case detection	Detection of suspect COVID-19 cases triggered by patients seeking care for their illness.
	3.1.2			Active case detection	Detection of suspect COVID-19 cases by health workers or other persons who move out into the community to find persons with symptoms and refer them for testing. Health checkpoints are a type of active case detection.
	3.1.3			Isolation	Separation of ill or infected persons from others.
Domestic Travel	4.5.1	Social and physical distancing measures	Domestic travel	Suspending or restricting movement	Limitation of the free movement of the population within the country.
	4.5.3			Restricting entry	Denial or restriction of entry for travelers proceeding from other subnational areas (e.g., states or provinces).
	4.5.4			Closing internal land borders	Partial or complete closure of ground crossings through land borders or other means of travel (e.g., suspension of flights or public transport) across subnational administrative areas (e.g., states or provinces).
Environmental measures	2.1	Environmental measures		Cleaning and disinfecting surfaces and objects	Removal of dirt and impurities from surfaces and objects and using chemicals or ultraviolet light to kill virus particles.
	2.2			Improving air ventilation	Natural or mechanical ventilation of air in rooms with the intent to provide a higher rate of air exchange and reduce the risk of airborne spread.
Financial packages	8.3	Other measures		Financial packages	
Gatherings, businesses and services	4.3.1	Social and physical distancing measures	Gatherings, businesses and services	Restricting private gatherings at home	Restriction of the number of non-residents who can participate in private gatherings or events in the home.
	4.3.2			Cancelling, restricting or adapting private gatherings outside the home	Cancellation or restriction of the size or frequency of private gatherings or events outside the home (e.g., weddings, funerals, parties). If gatherings or events are allowed to continue, adaptation measures may be put in place to prevent spread of COVID-19.
	4.3.3			Cancelling, closing, restricting or adapting public gatherings outside the home	Cancellation, closure or restriction of the size or frequency of public gatherings outside the home in public places. If gatherings or events are allowed to continue, adaptation measures may be put in place to prevent spread of COVID-19.
	4.3.4			Cancelling, restricting or adapting mass gatherings	Cancellation or restriction of the size of planned or spontaneous events where the number of people attending could strain the planning and response resources of the community or country hosting the event (e.g., sports events, concerts, the Hajj, conferences). If mass gatherings or events are allowed to continue, adaptation measures may put in place to prevent spread of COVID-19.
Individual measures	1.1	Individual measures		Performing hand hygiene	Washing hands with soap and water or cleaning hands by rubbing them with an alcohol-based formulation.
	1.3			Performing respiratory etiquette	Covering mouth and nose with a bent elbow or tissue when coughing or sneezing followed by disposing of any tissues into a closed trash receptacle and performing hand hygiene.
	1.5			Using other personal protective equipment	Using goggles, visors or gloves.
	1.6			Physical distancing	Remaining at a distance of at least 1 metre from other individuals without touching.
	5.1			Providing travel advice or warning	Health advice or warnings provided by government authorities on country-level transmission of COVID-19 to guide individual decisions on travel.
	5.2			Restricting visas	Suspension of visa on arrival or restriction of issuing visa for travelers originating from COVID-19 affected countries.
	5.3			Restricting entry	Denial of entry for travelers proceeding from specific countries or any country without there being a general international flight ban or land border closure in place.

International travel measures	5.4	International travel measures		Restricting exit	Denial of exit for travelers proceeding to specific countries.
	5.5			Entry screening and isolation or quarantine	Evaluation of the health or exposure status of travelers entering the country where exposure may be defined as arriving from an affected country. Screening for symptoms or exposure may be followed by testing and isolation or quarantine.
	5.6			Exit screening and isolation or quarantine	Evaluation of the health or exposure status of travelers leaving the country. Screening for symptoms or exposure to confirmed cases is usually followed by testing, if the person meets the national testing criteria, and isolation or quarantine.
	5.7			Suspending or restricting international flights	Stopping arrival of international flights, restricting the origin or number of flights or rescheduling of flights.
	5.8			Suspending or restricting international ferries or ships	Stopping arrival of ferries or ships, restricting the origin or number of ferries or ships or rescheduling of ferries or ships.
	5.9			Closing international land borders	Partial or complete closure of ground crossings through land borders.
Masks	1.4	Individual measures		Wearing a mask	Wearing a medical or cloth mask on the face, covering at least the nose and mouth.
Offices, businesses, institutions and operations	4.2.1	Social and physical distancing measures	Offices, businesses, institutions and operations	Adapting	Actions taken to reduce spread of COVID-19 while keeping offices, businesses, institutions and operations open.
	4.2.2			Closing	Closure of offices, businesses, institutions and services (e.g., offices, restaurants, bars, sports clubs, hotels, fitness centres, gyms, places of worship, cultural centres, conference centers).
Other measures	6.1	Drug-based measures		Using medications for prevention	Use of medications to prevent infection with the COVID 19 virus.
	6.2			Using medications for treatment	Use of medication to reduce infectiousness or shorten infectious period.
	8.1	Other measures		Legal and policy regulations	
	8.2			Scaling up	
	8.5			Other	
School measures	4.1.1	Social and physical distancing measures	School measures	Adapting	Actions taken to reduce spread of COVID-19 while keeping schools open.
	4.1.2			Closing	Closure or partial closure of the school building to students, teachers and staff. In partial closures, teachers and staff may remain to provide meals to low-income students or to provide lessons to children of essential workers.
Special populations	4.4.1	Social and physical distancing measures	Special populations	Shielding vulnerable groups	Measures to protect vulnerable persons at increased risk of severe disease from COVID-19 (e.g., older individuals and those with underlying conditions) or increased risk of infection (e.g., healthcare workers, homeless persons) including stay-at-home orders targeted to specific groups.
	4.4.2			Protecting populations in closed settings	Measures taken to reduce spread of COVID-19 in settings where populations reside in groups or are restrained or limited in movement or autonomy (e.g., some longer-term health care settings, seniors' residences, shelters, prisons). May include limiting visitors or outside excursions, cohorting of infected persons or green zones
	4.4.3			Protecting displaced populations	Measures taken to reduce spread in migrant, refugee or displaced persons settlements, including shielding strategies such as green zones
Stay-at-home order	4.5.2	Social and physical distancing measures	Domestic travel	Stay-at-home order	All residents ordered to stay at home except for essential activities (e.g., food shopping, traveling to work at essential businesses, medical visits).
Tracing and quarantining contacts	3.2.1	Surveillance and response measures	Tracing and quarantining contacts	Contact tracing	Interviews of cases to identify potential contacts.
	3.2.2			Quarantine of contacts	Restriction of activities or the separation of persons who are not ill but who may have been exposed to a case to monitor their health for the development of illness

Table S1: Grouping of WHO-NPI categories into 14 groups, which are treated as labels in our application of NPI predictions.