Causal Modeling of Twitter Activity during COVID-19

Oguzhan Gencoglu [1,*] and **Mathias Gruber** [2]

[1] Faculty of Medicine and Health Technology, Tampere University, 33720 Tampere, Finland
[2] LEO Pharma, 2750 Ballerup, Denmark; nano.mathias@gmail.com
* Correspondence: oguzhan.gencoglu@tuni.fi

Abstract: Understanding the characteristics of public attention and sentiment is an essential prerequisite for appropriate crisis management during adverse health events. This is even more crucial during a pandemic such as COVID-19, as primary responsibility of risk management is not centralized to a single institution, but distributed across society. While numerous studies utilize Twitter data in descriptive or predictive context during COVID-19 pandemic, causal modeling of public attention has not been investigated. In this study, we propose a causal inference approach to discover and quantify causal relationships between pandemic characteristics (e.g., number of infections and deaths) and Twitter activity as well as public sentiment. Our results show that the proposed method can successfully capture the epidemiological domain knowledge and identify variables that affect public attention and sentiment. We believe our work contributes to the field of infodemiology by distinguishing events that correlate with public attention from events that cause public attention.

Keywords: Twitter; machine learning; causal inference; COVID-19; sentiment analysis; social media

1. Introduction

On 11 March 2020, Coronavirus disease 2019 (COVID-19) was declared a pandemic by the World Health Organization [1] and more than 30 million people have been infected by it as of 19 September 2020 [2]. During such crises, capturing the dissemination of information, monitoring public opinion, observing compliance to measures, preventing disinformation, and relaying timely information is crucial for risk communication and decision-making about public health [3]. Previous national and global adverse health events show that social media surveillance can be utilized successfully for systematic monitoring of public perception in real-time due to its instantaneous global coverage [4–9].

Due to its large number of users, Twitter has been the primary social media platform for acquiring, sharing, and spreading information during global adverse events, including the COVID-19 pandemic [10]. Especially during the early stages of the COVID-19 pandemic, millions of posts have been tweeted in a span of couple of weeks by users, that is, citizens, politicians, corporations, and governmental institutions [11–14]. Consequently, numerous studies proposed and utilized Twitter as a data source for extracting insights on public health as well as insights on public attention during the COVID-19 pandemic. Focus of these studies include content analysis [15], topic modeling [16], sentiment analysis [17], nowcasting or forecasting of the disease [18], early detection of the outbreak [19], quantifying and detecting misinformation, disinformation, or conspiracies [20], and measuring public attitude towards relevant health concepts (e.g., social distancing or working from home) [21].

Despite such abundance of studies on manual or automatic analysis of social media data during COVID-19, *causal* modeling of relationships between characteristics of the pandemic and social media activity has not been investigated at all, as of September 2020. While descriptive statistical analysis (e.g., correlation, cluster, or exploratory analysis) is beneficial for pattern and hypothesis

discovery, and standard machine learning methods are effective in predictive modeling of those patterns, causal inference of relevant phenomena will not be possible without causal computational modeling. Causal modeling in the context of social media and pandemic can enable the optimization of onset of risk communication interventions to increase dissemination of accurate information. Similarly, it can be utilized to prevent acute propagation of negative sentiment with timely interventions. Consequently, such causal modeling can help risk communication policies to shift from alerting people to reassuring them. Furthermore, causal modeling enables simulation of what-if scenarios to enhance disaster preparedness. Therefore, as public decision-making can benefit from adequate assessment of public attention and correct understanding of underlying causes affecting it, we hereby propose causal modeling of Twitter activity.

We hypothesize that daily Twitter activity and sentiment during the COVID-19 pandemic has a causal relationship with the characteristics of the pandemic as well as with certain country statistics. We propose a structural causal modeling approach for discovering causal relationships and quantifying likelihood of events under various conditions (i.e., causal queries). To validate our approach, we collect close to 1 million tweets with location information spanning 57 days and identify several attributes of COVID-19 pandemic that might affect Twitter activity. We first employ a structure learning method to automatically construct a graphical causal structure in a data-driven manner. Then, we utilize Bayesian Networks (BNs) to learn conditional probability distributions of daily Twitter activity (number of daily tweets) and average public sentiment with respect to several pandemic characteristics such as total number of deaths and number of new infections. Our results show that the proposed structure discovery method can successfully capture the epidemiological domain knowledge. Furthermore, causal inference of daily Twitter activity with cross-validation across 12 countries show that our approach provides accurate predictions of Twitter activity with interpretable and intuitive results. We have released the full source code of our study (https://github.com/ogencoglu/causal_twitter_modeling_covid19). We believe our study contributes to the field of infodemiology by proposing causal modeling of public attention during the crisis of COVID-19 pandemic.

2. Going Beyond Correlations

Use of observational data from social media was proven to be beneficial in systematic monitoring of public opinion during adverse health events [4–9]. Such utilization of large, publicly available data becomes even more relevant during a global pandemic such as COVID-19, as neither enough time nor a practical way to run variety of randomized control trials for quantifying public opinion exist. Furthermore, as disease containment measures (e.g., lockdowns, quarantines, and curfews), associated financial issues (e.g., due to inability to work), and changes in social dynamics may impact mental health negatively [22–24], opinion surveillance methods that do not carry the risk of further stressing of the participants are pertinent.

Themes of previous studies that focus on exploration of, description of, correlation of, or predictive modeling with Twitter data during COVID-19 pandemic include sentiment analysis [17,25–28], public attitude/interest measurement [21,29–31], content analysis [15,32–36], topic modeling [16,26,27,37–40], analysis of misinformation, disinformation, or conspiracies [20,41–46], outbreak detection or disease nowcasting/forecasting [18,19], and more [47–52]. Similarly, data from other social media channels (e.g., Weibo, Reddit, Facebook) or search engine statistics are utilized for parallel analyses related to COVID-19 pandemic as well [53–69]. While these studies reveal important information and patterns, they do not attempt to uncover or model causal relationships between the attributes of COVID-19 pandemic and social media activity. As *correlation does not imply causation* (e.g., spurious correlations), the ability to identify truly causal relationships between pandemic characteristics and public behaviour (online or not) remains crucial for devising public policies that are more impactful. Without causal understanding, our efforts and decisions on risk communication, public health engagement, health intervention timing, and adjustment of resources for fighting disinformation, fearmongering, and alarmism will stay subpar.

The task of forging causal models comes with numerous challenges in various domains because, typically, domain knowledge and significant amount of time from the experts is required. For substantially complex phenomena such as a pandemic due to a novel virus, diagnosing causal attributions becomes even harder. Therefore, learning causal relationships automatically from observational data has been studied in machine learning. One of the primary challenges for this pursuit is that numerous latent variables that we can not observe exist in real world problems. In fact, numerous other latent variables that we are not even aware of may exist as well. As latent variables can induce statistical correlations between observed variables that do not have a causal relationship, *confounding factors* arise. While this phenomenon may not exhibit a considerable problem in standard probabilistic models, causal modeling suffers from it immensely.

Several machine learning methods are proposed for learning causal structures from observational data and some allow combination of statistical information (learned from the data) and domain expertise [70,71]. Bayesian networks are frequently utilized frameworks for learning models once the causal structure is fixed. As probabilistic graphical models, BNs flexibly unify graphical models, structural equations, and counterfactual logic [71–74]. A causal BN consists of a directed acyclic graph (DAG) in which nodes correspond to random variables and edges correspond to direct causal influence of one node on another [71]. This compact representation of high-dimensional probability spaces (e.g., joint probability distributions) provides intuitive and explainable models for us. In addition, BNs allow not only straightforward observational computations (e.g., calculation of marginal probabilities) but also interventional ones (e.g., *do-calculus*), enabling simulations of various what-if scenarios.

3. Methods

3.1. Data

We primarily utilized two data sources for our study, that is, daily number of officially reported COVID-19 infections and deaths from "COVID-19 Data Repository" by the Center for Systems Science and Engineering at Johns Hopkins University [2] and daily count of COVID-19 related tweets from Twitter [75]. A 57 day period between 22 January–18 March 2020 is chosen for this study to represent the early stages of the pandemic when disease characteristics are less known and public panic is elevated. We collected 954,902 tweets that have location information from Twitter by searching for *#covid19* and *#coronavirus* hashtags. Similar to other studies [18,20,46], geolocation of the tweets is inferred either by using user geo-tagging or geo-coding the information available in users' profiles. Timeline of daily log-distribution of collected tweet counts among 177 countries can be examined from Figure 1. The trend shows an increasing prevalence of high daily number of tweets as the pandemic spreads across the globe with time.

We select the following 12 countries for our causal modeling analysis: Italy, Spain, Germany, France, Switzerland, United Kingdom, Netherlands, Norway, Austria, Belgium, Sweden, and Denmark. These are the countries with substantial number of reported COVID-19 cases (listed in descending order) in Europe as of 18 March 2020, yet still exhibiting a high diversity in terms of the timeline of the pandemic. For instance, while Italy located further in the pandemic timeline due to being hit first in Europe, United Kingdom could be considered in the very initial stages of it for the analysis period of our study. Figure 2 depicts the cumulative number of tweet counts alongside with that of reported infections and deaths for the selected countries. Evident correlations between these variables can be noticed. A sharp increase in Twitter activity is observed after 28–29 February, which corresponds to the period of each country having at least one confirmed COVID-19 case.
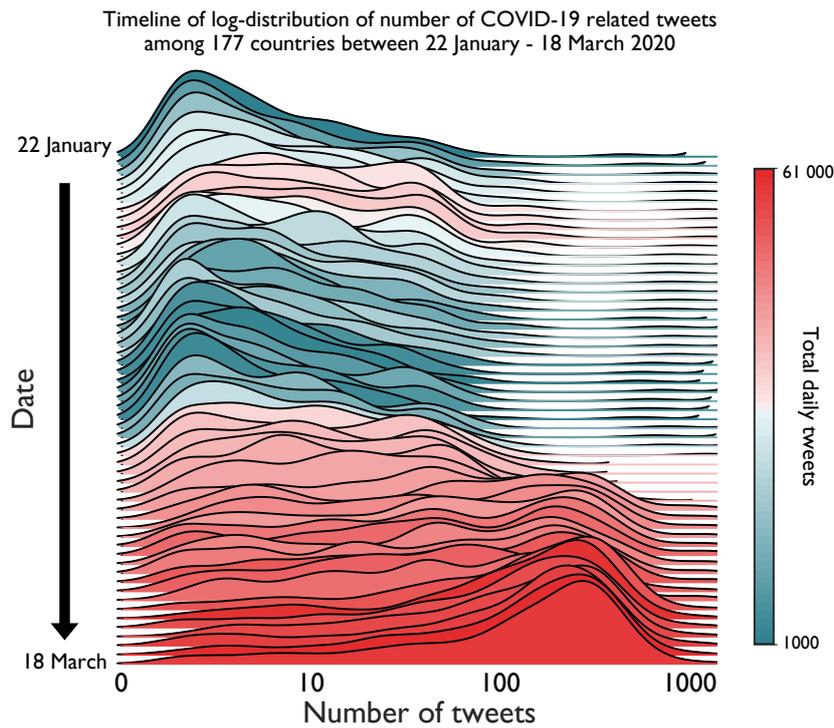
Timeline of log-distribution of number of COVID-19 related tweets
among 177 countries between 22 January - 18 March 2020

Figure 1. Evolution of COVID-19 related Twitter activity between 22 January–18 March 2020.

Cumulative number of tweets, infections, and deaths for the selected 12 countries
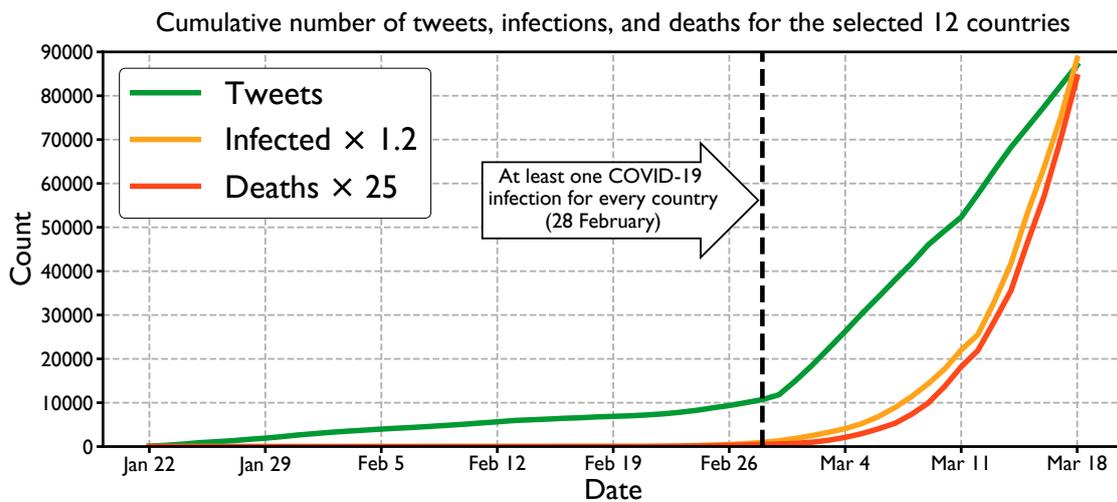
Figure 2. Cumulative counts of Twitter activity and COVID-19 statistics for the selected countries during the study period.

3.2. Feature Selection

In order to characterize the pandemic straightforwardly, we calculate the following six features (attributes) from the official COVID-19 incident statistics for each day for 12 selected countries: (1) *total number of infections up to that day* (normalized by the country's population), (2) *number of new infections* (normalized by the country's population), (3) *percentage increase in infections* (with respect to previous day), and the same three statistics for *deaths* (4-5-6).

Recent epidemiological studies on COVID-19 reveal the following: people over the age 65 are the primary risk group both for infection and mortality [76–79] and human-to-human transmission of the virus is largely occurring among family members or among people who co-reside [77,80,81]. In order to be able to test whether our approach can capture this scientific domain knowledge or not, we collect the following two features for each country: (7) *percentage of population over the age of 65* [82] and (8)

percentage of single-person households [83]. Finally, as we know that popularity of Twitter in a country and announcement of national lockdown (e.g., closing of schools, banning of gatherings) unequivocally affect the Twitter activity in that country, we add (9) *percentage of population using Twitter* [84] and (10) *is_lockdown_announced?* (3 day period is encoded as `Yes` if government restriction is announced [85], `No` otherwise) features as well. We represent Twitter activity by simply counting the (11) *number of daily tweets* (normalized by the country's population). We also calculate the (12) *average daily sentiment* (in range $[-1, 1]$) of English tweets (corresponding to over 80% of all tweets) by utilizing a pre-trained sentiment classifier (DistilBERT [86]). We treat each day as an observation and represent each day with these 12 attributes ($n = 12$) for structure learning, resulting in a feature matrix of dimensions 684×12. 684 observations come from 12 countries times 57 days.

For the purpose of increasing interpretability, we discretize the daily numerical features by mapping them to 2 categorical levels, namely `High` or `Low`. Features related to the pandemic (infections and deaths) and Twitter activity employ a cut-off value of 75th percentile and remaining numerical features employ a cut-off value of 50th percentile (corresponding to median). Such categorization, for instance, turns the numerical value of "population-normalized increase in deaths of 1.7325×10^{-7}" into a relatively calculated category of `High` for a given day. Sentiment scores are mapped to `Positive` (≥ 0) or `Negative` (< 0) as well.

3.3. Structure Learning and Causal Inference

In structure learning we would like to learn a directed acyclic graph, G, that describes the conditional dependencies between variables in a given data matrix. A typical formulation of this problem is a *structural equation model* (more generally a *generalized linear model*) in which a weighted adjacency matrix, $W \in \mathbb{R}^{n \times n}$, defines the graph. This is essentially a parametric model that enables operations on the continuous space of $n \times n$ matrices instead of discrete space of DAGs. Such formulation enables a score-based learning of DAGs, that is,

$$\min_{W \in \mathbb{R}^{n \times n}} L(W)$$
$$\text{subject to } G(W) \in \text{DAGs,}$$

(1)

where $G(W)$ is the n-node graph induced by the weighted adjacency matrix, W, and L is the score/loss function to be minimized. Even though the loss function is continuous, solving Equation (1) is still a non-convex, combinatorial optimization problem as the acyclicity constraint is discrete and difficult to enforce. Note that acyclicity is a strict requirement for causal graphs. In order to tackle this problem efficiently, we utilize the recently proposed NOTEARS (corresponding to *Non-combinatorial Optimization via Trace Exponential and Augmented lagRangian for Structure learning*) algorithm for structure learning [87].

NOTEARS algorithm discovers a directed acyclic graph from the observational data by re-formulating the structure learning problem as a purely continuous optimization. This approach differs significantly from existing work in the field which predominantly operates on discrete space of graphs. Re-formulation is achieved by introducing a continuous measure of "DAG-ness", $h(W)$, which quantifies the severity of violations from acyclicity as W changes. Consequently, the problem formulation becomes

$$\min_{W \in \mathbb{R}^{n \times n}} L(W)$$
$$\text{subject to } h(W) = 0,$$

(2)

which enables utilization of standard numerical solving methods and scales cubically, $\mathcal{O}(n^3)$, with the number of variables instead of exponentially as in other structure learning methods. We have chosen the score to be the least squared loss (can be any smooth loss function) with l_1-regularization term to discover a sparse DAG and use a gradient-based minimizer to solve Equation (2). In our context, we discover such an adjacency matrix that the graph it defines encodes the dependencies between our

features in a close-to-optimal manner (finding the global optimum is NP-hard [88,89]) and is a DAG. Efficiency of this approach enables structure learning in a scalable manner.

As NOTEARS algorithm allows incorporation of expert knowledge, we also put certain constraints on the structure in our experiment. These constraints correspond to prohibited causal attributions based on simple logical assumptions, for example, Twitter activity on a given day can not have a causal effect on number of deaths from COVID-19 on that day. Full list of these constraints can be found in Table A1 in the Appendix A. Once the structure is learned (both by data and logical constraints), we treat it as a causal model and learn the parameters of a Bayesian network on it with the training data in order to capture the conditional dependencies between variables. During inference on test data, probabilities of each possible state of a node with respect to the given input data is computed from the conditional probability distributions.

Our approach allows straightforward querying of the model with varying observations. For instance for a given day, the probability of Twitter activity being High, when total number of infections are Low and new deaths are High, that is,

$$\Pr(\text{Twitter Activity} = \text{H} \mid \text{Total Infections} = \text{H}, \text{New Deaths} = \text{L}), \tag{3}$$

can be computed by propagating the impact of these queries through the nodes of interest. By utilizing this property of our approach, we compute marginal probabilities for gaining further insights on likelihoods of various events.

Essentially, we expect two observations from our experiment. First, we expect the structure learning algorithm to discover the causal relations verified by domain/expert knowledge (e.g., % of single-person households and % of 65+ people affecting infections) and common sense/elementary algebra (e.g., new deaths affecting percentage change in deaths). Second, we expect the calculated likelihoods from the Bayesian network are in parallel with domain knowledge as well, for example, high % of people over 65 increasing the marginal likelihood of deaths instead of decreasing it or high % of single households (better social isolation) decreasing the marginal likelihood of infections instead of increasing it. Realization of these expectations will show that the proposed method can indeed capture causal relationships and will increase our confidence in discovered relationships between the pandemic attributes and Twitter activity as well as confidence in corresponding likelihoods.

3.4. Evaluation

We validate our approach first by inspecting whether the expected causal relationships (e.g., domain knowledge on COVID-19) are captured or not. Then, we infer the Twitter activity of each day from the learned Bayesian Network. Essentially, this corresponds to a binary classification task, that is, predicting the Twitter activity as High or Low from the rest of the variables. We utilize a Leave-One-Country-Out (LOCO) cross-validation scheme in which each fold consists of training set from 11 countries (627 samples) and test set (57 samples) from the remaining country. We do not perform standard k-fold cross-validation as we would like to measure the generalization performance across countries and prevent overly optimistic results. Therefore, we ensure that the observations from the same country fall in the same set (either training or test) for every fold. We evaluate the performance of our approach by calculating the average Area Under the Receiver Operating Characteristic curve (AUROC) of the cross-validation runs. For quantifying the causal effect of characteristics of pandemic and relevant country statistics on Twitter activity, we report likelihoods from the model by querying various conditions.

4. Results

The jointly (with statistical learning from data and user-defined logical constraints) discovered causal model by the structure learning algorithm can be examined from Figure 3. Different families of attributes are colored differently for ease of inspection—blue for COVID-19 pandemic related variables,

yellow for country-specific statistics, green for government interventions, and red for representing variables related to public attention and sentiment in Twitter. Daily Twitter activity is affected by 4 variables, namely Twitter usage statistics of that country, new infections on that day, new deaths on that day, and whether national lockdown is announced or not. Similarly, 4 variables affecting the average daily sentiment in Twitter are new infections on that day, new deaths on that day, total deaths up to that day, and again lockdown announcements. Total number of infections did not show any causal effect on Twitter activity or on average public sentiment.
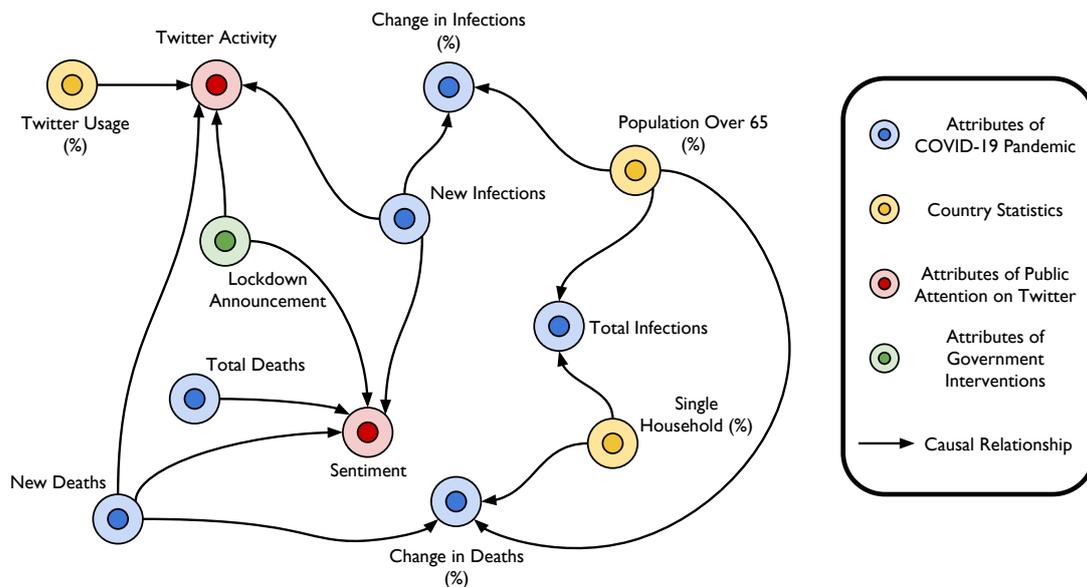

Figure 3. Discovered graph depicting causal relationships between various attributes.

Leave-One-Country-Out cross-validation results in terms of AUROCs can be seen in Table 1. Each row in the table corresponds to a cross-validation fold in which the Twitter activity in that particular country was tried to be predicted. The Bayesian network model achieves an average AUROC score of 0.833 across countries when trying to infer the Twitter activity from the rest of the variables for a given day. Daily Twitter patterns of Germany, Italy, and Sweden show very high predictability with AUROC scores above 0.97. United Kingdom shows the worst predictability with an AUROC of 0.68.

Calculation of marginal probabilities for several queries are presented in Table 2. Public attention and sentiment-related target variables and states are set to `High Twitter Activity` and `Negative Sentiment`.

Table 1. Area Under the Receiver Operating Characteristic curve (AUROC) result for each fold of Leave-One-Country-Out cross-validation.

Cross Validation Test Country	AUROC
Austria	0.798
Belgium	0.728
Denmark	0.831
France	0.776
Germany	0.992
Italy	0.976
Netherlands	0.746
Norway	0.907
Spain	0.766
Sweden	0.998
Switzerland	0.789
United Kingdom	0.684
Average	*0.833*

Table 2. Examples of queries and computed marginal probabilities for Twitter activity and average sentiment.

Query	Variable and State	Pr()
Single-person household (%) = H 65+ (%) = L	Total Infections = H	0.178
Single-person household (%) = L 65+ (%) = H	Total Infections = H	0.241
New Infections = H New Deaths = H	Twitter Activity = H	0.496
New Infections = L New Deaths = L	Twitter Activity = H	0.184
New Infections = H New Deaths = H Twitter Usage = H Lockdown Announcement = Yes	Twitter Activity = H	0.800
New Infections = L New Deaths = L Twitter Usage = L Lockdown Announcement = No	Twitter Activity = H	0.120
New Deaths = H	Sentiment = Neg	0.624
New Deaths = L	Sentiment = Neg	0.277
Total Deaths = H	Sentiment = Neg	0.344
Total Deaths = L	Sentiment = Neg	0.290
Lockdown Announcement = Yes	Sentiment = Neg	0.501
Lockdown Announcement = No	Sentiment = Neg	0.286

5. Discussion

By analyzing observational data, we attempt to discover causal associations between national COVID-19 patterns and Twitter activity as well as public sentiment during the early stages of the pandemic. Some of our findings are expected associations such as popularity of Twitter in a country (Twitter usage) affecting Twitter activity. Other expected causal relationships were new deaths affecting change in deaths and new infections affecting change in infections, due to trivial mathematical definitions. These were captured successfully as well. It is important to note that no causal relationship between infection statistics and death statistics was discovered which might seem against intuition. This is because in this study we treat each day as an observation in our modeling and do not create time-lagged version of variables. While some of our results imply expected associations, we also observe more interesting implications that are in alignment with recent scientific literature on COVID-19. For instance, percentage of single-person households affects the total number of COVID-19 infections. Similarly, the percentage of 65+ population affects the percentage change in deaths (essentially corresponding to rate of deaths). When the queries regarding domain knowledge are examined, we see that low percentage of single-person households (less social isolation) and high percentage of 65+ population increases the probability of total infections being high when compared to the opposite settings. This is in line with recent scientific literature on COVID-19 transmission characteristics [76–81].

By inferring Twitter activity, we show the generalization ability of causal inference across 12 countries with reasonable accuracy. Factors affecting Twitter activity and sentiment are discussion-worthy as well. By observing correlations, Wong et al. hints that there may be a link between announcement of new infections and Twitter activity [17]. Our results in Figure 3 and Table 2 suggest the same with a causal point of view. Similarly, our finding of negative impact of declaration of

government measures on public sentiment is also in parallel with recent research. By analyzing Chinese social media, Li et al. show that official declaration of COVID-19 (epidemic at that time) correlates with increased negative emotions such as anxiety, depression, and indignation [56]. When new infections, new deaths, total deaths are high and an announcement of lockdown is made, Twitter activity on that day becomes more than 6 times more likely than when the situation is opposite (probabilities of 0.8 vs. 0.12). High number of new deaths for a given day causes the sentiment to be much more negative than low number of new deaths (probabilities of 0.624 vs. 0.277). Similarly, an announcement of lockdown is causally associated with an increase in negative sentiment in Twitter (probabilities of 0.501 vs. 0.286).

As it is important to observe the countries that are ahead in terms of pandemic timeline and learn the behaviour of the pandemic, it is equally important to understand also the public attention and sentiment characteristics from those countries. Wise et al. show that risk perception of people and their frequency of engagement in protective behaviour change during the early stages of the pandemic [90]. Inference of such patterns in a causal manner from social media can aid us in the pursuit of timely decisions and suitable policy-making, and consequently, high public engagement. After all, primary responsibility of risk management during a global pandemic is not centralized to a single institution, but distributed across society. For example, Zhong et al. shows that people's adherence to COVID-19 control measures is affected by their knowledge and attitudes towards it [91]. In that regard, computational methods such as causal inference and causal reasoning can help us disentangle correlations and causation between the observed variables of the adverse phenomenon.

In real-world scenarios, it is virtually impossible to correctly identify all the causal associations due to presence of numerous confounding factors. As in with all methods in machine learning, a trade-off between false positive associations and false negative ones exists in our approach as well. While we rely on official COVID-19 statistics, testing and reporting methodologies as well as policies can change during the course of the pandemic. Furthermore, in the context of this study, ground truth causal associations do not exist even for a few variables, preventing the direct measurement of performance of causal discovery methods. We would like to emphasize that we acknowledge these and other relevant limitations of our study. Our study has further limitations regarding the simplifications on our problem formulation and data. For instance, we do not attempt to model temporal causal relationships in this study, for example, high deaths numbers having an impact on the public sentiment possibly for several following days. We have not taken into account remarks by famous politicians, public figures, or celebrities which may indeed impact social media discussions. We have not incorporated "retweets" or "likes" into our models either. We would also like emphasize that with this study we wanted to introduce an uncomplicated example of causal modeling perspective to social media analysis during COVID-19.

Future work includes investigating the effect of dynamics of the pandemic on the spreading mechanisms of information, including relevant health topics in Twitter and other social media. As social media can be exploited for deliberately creating panic and confusion [92], causal inference on patterns of misinformation and disinformation propagation in Twitter will be studied as well. Finally, country-specific models with more granular statistics of the country and time-delayed variables will be investigated for a longer analysis period.

6. Conclusions

Distinguishing epidemiological events that correlate with public attention from epidemiological events that cause public attention is crucial for constructing impactful public health policies. Similarly, monitoring fluctuations of public opinion becomes actionable only if causal relationships are identified. We hope our study serves as a first example of causal inference on social media data for increasing our understanding of factors affecting public attention and sentiment during COVID-19 pandemic.

Abbreviations

The following abbreviations are used in this manuscript:

AUROC	Area Under the Receiver Operating Characteristic curve
COVID-19	Coronavirus Disease 2019
BN	Bayesian Network
DAG	Directed Acyclic Graph
LOCO	Leave-One-Country-Out
NOTEARS	Non-combinatorial Optimization via Trace Exponential and Augmented lagRangian for Structure learning

Appendix A

Prohibited causal associations are listed in Table A1 below. For example, *Twitter activity* can not cause any other variable for a given day. Similarly, *Twitter usage percentage* or *lockdown announcement* can not have a causal relationship with *new deaths* for a given day.

Table A1. Prohibited causal associations (constraints) for structure learning.

From	To
Any node	Population Over 65 (%) Twitter Usage (%) Single Household (%)
Twitter Activity Sentiment	Any node
Twitter Usage (%) Lockdown Announcement	Total Infections New Infections Change in Infections (%) Total Deaths New Deaths Change in Deaths (%)
Population Over 65 (%) Single Household (%)	Twitter Activity Sentiment
Twitter Usage (%)	Sentiment

References

1. Cucinotta, D.; Vanelli, M. WHO Declares COVID-19 a Pandemic. *Acta Bio-Medica Atenei Parm.* **2020**, *91*, 157–160. [CrossRef]
2. Dong, E.; Du, H.; Gardner, L. An Interactive Web-based Dashboard to Track COVID-19 in Real Time. *Lancet Infect. Dis.* **2020**. [CrossRef]
3. Van Bavel, J.J.; Baicker, K.; Boggio, P.S.; Capraro, V.; Cichocka, A.; Cikara, M.; Crockett, M.J.; Crum, A.J.; Douglas, K.M.; Druckman, J.N.; et al. Using Social and Behavioural Science to Support COVID-19 Pandemic Response. *Nat. Hum. Behav.* **2020**, 1–12. [CrossRef]
4. Signorini, A.; Segre, A.M.; Polgreen, P.M. The Use of Twitter to Track Levels of Disease Activity and Public Concern in the US During the Influenza A H1N1 Pandemic. *PLoS ONE* **2011**, *6*. [CrossRef]

5. Ji, X.; Chun, S.A.; Geller, J. Monitoring Public Health Concerns Using Twitter Sentiment Classifications. In Proceedings of the IEEE International Conference on Healthcare Informatics, Philadelphia, PA, USA, 9–11 September 2013; pp. 335–344. [CrossRef]

6. Ji, X.; Chun, S.A.; Wei, Z.; Geller, J. Twitter Sentiment Classification for Measuring Public Health Concerns. *Soc. Netw. Anal. Min.* **2015**, *5*, 13. [CrossRef]

7. Weeg, C.; Schwartz, H.A.; Hill, S.; Merchant, R.M.; Arango, C.; Ungar, L. Using Twitter to Measure Public Discussion of Diseases: A Case Study. *JMIR Public Health Surveill.* **2015**, *1*, e6. [CrossRef]

8. Mollema, L.; Harmsen, I.A.; Broekhuizen, E.; Clijnk, R.; De Melker, H.; Paulussen, T.; Kok, G.; Ruiter, R.; Das, E. Disease Detection or Public Opinion Reflection? Content Analysis of Tweets, Other Social Media, and Online Newspapers During the Measles Outbreak in the Netherlands in 2013. *J. Med. Internet Res. (JMIR)* **2015**, *17*, e128. [CrossRef]

9. Jordan, S.E.; Hovet, S.E.; Fung, I.C.H.; Liang, H.; Fu, K.W.; Tse, Z.T.H. Using Twitter for Public Health Surveillance from Monitoring and Prediction to Public Response. *Data* **2019**, *4*, 6. [CrossRef]

10. Rosenberg, H.; Syed, S.; Rezaie, S. The Twitter Pandemic: the Critical Role of Twitter in the Dissemination of Medical Information and Misinformation During the COVID-19 Pandemic. *Can. J. Emerg. Med.* **2020**, 1–7. [CrossRef]

11. Chen, E.; Lerman, K.; Ferrara, E. Covid-19: The First Public Coronavirus Twitter Dataset. *arXiv* **2020**, arXiv:2003.07372.

12. Gao, Z.; Yada, S.; Wakamiya, S.; Aramaki, E. NAIST COVID: Multilingual COVID-19 Twitter and Weibo Dataset. *arXiv* **2020**, arXiv:2004.08145.

13. Lamsal, R. Corona Virus (COVID-19) Tweets Dataset. *IEEEDataPort* **2020**. [CrossRef]

14. Aguilar-Gallegos, N.; Romero-García, L.E.; Martínez-González, E.G.; García-Sánchez, E.I.; Aguilar-Ávila, J. Dataset on Dynamics of Coronavirus on Twitter. *Data Brief* **2020**, *30*, 105684. [CrossRef] [PubMed]

15. Thelwall, M.; Thelwall, S. Retweeting for COVID-19: Consensus Building, Information Sharing, Dissent, and Lockdown Life. *arXiv* **2020**, arXiv:2004.02793.

16. Sha, H.; Hasan, M.A.; Mohler, G.; Brantingham, P.J. Dynamic Topic Modeling of the COVID-19 Twitter Narrative Among US Governors and Cabinet Executives. *arXiv* **2020**, arXiv:2004.11692.

17. Wong, C.M.L.; Jensen, O. The Paradox of Trust: Perceived Risk and Public Compliance During the COVID-19 Pandemic in Singapore. *J. Risk Res.* **2020**, 1–10. [CrossRef]

18. Turiel, J.; Aste, T. Wisdom of the Crowds in Forecasting COVID-19 Spreading Severity. *arXiv* **2020**, arXiv:2004.04125.

19. Gharavi, E.; Nazemi, N.; Dadgostari, F. Early Outbreak Detection for Proactive Crisis Management Using Twitter Data: COVID-19 a Case Study in the US. *arXiv* **2020**, arXiv:2005.00475.

20. Chary, M.; Overbeek, D.; Papadimoulis, A.; Sheroff, A.; Burns, M. Geospatial Correlation Between COVID-19 Health Misinformation on Social Media and Poisoning with Household Cleaners. *medRxiv* **2020**. [CrossRef]

21. Kayes, A.; Islam, M.S.; Watters, P.A.; Ng, A.; Kayesh, H. Automated Measurement of Attitudes Towards Social Distancing Using Social Media: A COVID-19 Case Study. *Preprints* **2020**. [CrossRef]

22. Wang, C.; Pan, R.; Wan, X.; Tan, Y.; Xu, L.; Ho, C.S.; Ho, R.C. Immediate Psychological Responses and Associated Factors During the Initial Stage of the 2019 Coronavirus Disease (COVID-19) Epidemic Among the General Population in China. *Int. J. Environ. Res. Public Health* **2020**, *17*, 1729. [CrossRef] [PubMed]

23. Cullen, W.; Gulati, G.; Kelly, B. Mental Health in the COVID-19 Pandemic. *QJM An Int. J. Med.* **2020**, *113*, 311–312. [CrossRef] [PubMed]

24. Brooks, S.K.; Webster, R.K.; Smith, L.E.; Woodland, L.; Wessely, S.; Greenberg, N.; Rubin, G.J. The Psychological Impact of Quarantine and How to Reduce It: Rapid Review of the Evidence. *Lancet* **2020**, *395*, 912–920. [CrossRef] [PubMed]

25. Dubey, A.D.; Tripathi, S. Analysing the Sentiments towards Work-From-Home Experience during COVID-19 Pandemic. *J. Innov. Manag.* **2020**, *8*. [CrossRef]

26. Duong, V.; Pham, P.; Yang, T.; Wang, Y.; Luo, J. The Ivory Tower Lost: How College Students Respond Differently than the General Public to the COVID-19 Pandemic. *arXiv* **2020**, arXiv:2004.09968.

27. Medford, R.J.; Saleh, S.N.; Sumarsono, A.; Perl, T.M.; Lehmann, C.U. An "Infodemic": Leveraging High-Volume Twitter Data to Understand Public Sentiment for the COVID-19 Outbreak. *medRxiv* **2020**. [CrossRef]

28. Samuel, J.; Ali, G.M.N.; Rahman, M.M.; Esawi, E.; Samuel, Y. COVID-19 Public Sentiment Insights and Machine Learning for Tweets Classification. *Preprints* **2020**. [CrossRef]

29. Batooli, Z.; Sayyah, M. Measuring Social Media Attention of Scientific Research on Novel Coronavirus Disease 2019 (COVID-19): An Investigation on Article-level Metrics Data of Dimensions. *Prepr. Res. Sq.* **2020**. [CrossRef]

30. Kwon, J.; Grady, C.; Feliciano, J.T.; Fodeh, S.J. Defining Facets of Social Distancing during the COVID-19 Pandemic: Twitter Analysis. *medRxiv* **2020**. [CrossRef]

31. Cinelli, M.; Quattrociocchi, W.; Galeazzi, A.; Valensise, C.M.; Brugnoli, E.; Schmidt, A.L.; Zola, P.; Zollo, F.; Scala, A. The COVID-19 Social Media Infodemic. *arXiv* **2020**, arXiv:2003.05004.

32. Park, H.W.; Park, S.; Chong, M. Conversations and Medical News Frames on Twitter: Infodemiological Study on COVID-19 in South Korea. *J. Med. Internet Res. (JMIR)* **2020**, *22*, e18897. [CrossRef] [PubMed]

33. Thelwall, M.; Thelwall, S. Covid-19 tweeting in English: Gender differences. *arXiv* **2020**, arXiv:2003.11090.

34. Alshaabi, T.; Minot, J.; Arnold, M.; Adams, J.L.; Dewhurst, D.R.; Reagan, A.J.; Muhamad, R.; Danforth, C.M.; Dodds, P.S. How the World's Collective Attention is Being Paid to a Pandemic: COVID-19 Related 1-gram Time Series for 24 Languages on Twitter. *arXiv* **2020**, arXiv:2003.12614.

35. Lopez, C.E.; Vasu, M.; Gallemore, C. Understanding the Perception of COVID-19 Policies by Mining a Multilanguage Twitter Dataset. *arXiv* **2020**, arXiv:2003.10359.

36. Dewhurst, D.R.; Alshaabi, T.; Arnold, M.V.; Minot, J.R.; Danforth, C.M.; Dodds, P.S. Divergent Modes of Online Collective Attention to the COVID-19 Pandemic are Associated with Future Caseload Variance. *arXiv* **2020**, arXiv:2004.03516.

37. Abd-Alrazaq, A.; Alhuwail, D.; Househ, M.; Hamdi, M.; Shah, Z. Top Concerns of Tweeters During the COVID-19 Pandemic: Infoveillance Study. *J. Med. Internet Res. (JMIR)* **2020**, *22*, e19016. [CrossRef]

38. Wicke, P.; Bolognesi, M.M. Framing COVID-19: How We Conceptualize and Discuss the Pandemic on Twitter. *arXiv* **2020**, arXiv:2004.06986.

39. Jarynowski, A.; Wójta-Kempa, M.; Belik, V. Trends in Perception of COVID-19 in Polish Internet. *medRxiv* **2020**. [CrossRef]

40. Ordun, C.; Purushotham, S.; Raff, E. Exploratory Analysis of Covid-19 Tweets Using Topic Modeling, UMAP, and DiGraphs. *arXiv* **2020**, arXiv:2005.03082.

41. Yang, K.C.; Torres-Lugo, C.; Menczer, F. Prevalence of Low-Credibility Information on Twitter During the COVID-19 Outbreak. *arXiv* **2020**, arXiv:2004.14484.

42. Ahmed, W.; Vidal-Aic, J.; Downing, J.; Seguí, F.L. COVID-19 and the 5G Conspiracy Theory: Social Network Analysis of Twitter Data. *J. Med. Internet Res. (JMIR)* **2020**, *22*, e19458. [CrossRef] [PubMed]

43. Ferrara, E. #COVID-19 on Twitter: Bots, Conspiracies, and Social Media Activism. *arXiv* **2020**, arXiv:2004.09531.

44. Bridgman, A.; Merkley, E.; Loewen, P.J.; Owen, T.; Ruths, D.; Teichmann, L.; Zhilin, O. The Causes and Consequences of COVID-19 Misperceptions: Understanding the Role of News and Social Media. *OSF Prepr.* **2020**. [CrossRef]

45. Ahmed, W.; Vidal-Alaball, J.; Downing, J.; Seguí, F.L. Dangerous Messages or Satire? Analysing the Conspiracy Theory Linking 5G to COVID-19 through Social Network Analysis. *J. Med. Internet Res. (JMIR)* **2020**. [CrossRef]

46. Gallotti, R.; Valle, F.; Castaldo, N.; Sacco, P.; De Domenico, M. Assessing the Risks of "Infodemics" in Response to COVID-19 Epidemics. *medRxiv* **2020**. [CrossRef]

47. Golder, S.; Klein, A.; Magge, A.; O'Connor, K.; Cai, H.; Weissenbacher, D. Extending A Chronological and Geographical Analysis of Personal Reports of COVID-19 on Twitter to England, UK. *medRxiv* **2020**. [CrossRef]

48. Sarker, A.; Lakamana, S.; Hogg-Bremer, W.; Xie, A.; Al-Garadi, M.A.; Yang, Y.C. Self-reported COVID-19 Symptoms on Twitter: An Analysis and a Research Resource. *J. Am. Med. Informat. Assoc.* **2020**. [CrossRef]

49. Li, I.; Li, Y.; Li, T.; Alvarez-Napagao, S.; Garcia, D. What Are We Depressed about When We Talk about COVID19: Mental Health Analysis on Tweets Using Natural Language Processing. *arXiv* **2020**, arXiv:2004.10899.

50. Xu, P.; Dredze, M.; Broniatowski, D.A. The Twitter Social Mobility Index: Measuring Social Distancing Practices from Geolocated Tweets. *arXiv* **2020**, arXiv:2004.02397.

51. Lyu, H.; Chen, L.; Wang, Y.; Luo, J. Sense and Sensibility: Characterizing Social Media Users Regarding the Use of Controversial Terms for COVID-19. *IEEE Trans. Big Data* **2020**. [CrossRef]

52. Schild, L.; Ling, C.; Blackburn, J.; Stringhini, G.; Zhang, Y.; Zannettou, S. "Go Eat A Bat, Chang!": An Early Look on the Emergence of Sinophobic Behavior on Web Communities in the Face of COVID-19. *arXiv* **2020**, arXiv:2004.04046.

53. Rovetta, A.; Bhagavathula, A.S. COVID-19-Related Web Search Behaviors and Infodemic Attitudes in Italy: Infodemiological Study. *JMIR Public Health Surveill.* **2020**, *6*, e19374. [CrossRef] [PubMed]

54. Shahsavari, S.; Holur, P.; Tangherlini, T.R.; Roychowdhury, V. Conspiracy in the Time of Corona: Automatic detection of Covid-19 Conspiracy Theories in Social Media and the News. *arXiv* **2020**, arXiv:2004.13783.

55. Li, J.; Xu, Q.; Cuomo, R.; Purushothaman, V.; Mackey, T. Data Mining and Content Analysis of the Chinese Social Media Platform Weibo during the Early COVID-19 Outbreak: Retrospective Observational Infoveillance Study. *JMIR Public Health Surveill.* **2020**, *6*, e18700. [CrossRef] [PubMed]

56. Li, S.; Wang, Y.; Xue, J.; Zhao, N.; Zhu, T. The Impact of COVID-19 Epidemic Declaration on Psychological Consequences: A Study on Active Weibo Users. *Int. J. Environ. Res. Public Health* **2020**, *17*, 2032. [CrossRef] [PubMed]

57. Velásquez, N.; Leahy, R.; Restrepo, N.J.; Lupu, Y.; Sear, R.; Gabriel, N.; Jha, O.; Johnson, N. Hate Multiverse Spreads Malicious COVID-19 Content Online Beyond Individual Platform Control. *arXiv* **2020**, arXiv:2004.00673.

58. Zhao, Y.; Xu, H. Chinese Public Attention to COVID-19 Epidemic: Based on Social Media. *medRxiv* **2020**. [CrossRef]

59. Li, L.; Zhang, Q.; Wang, X.; Zhang, J.; Wang, T.; Gao, T.L.; Duan, W.; Tsoi, K.K.f.; Wang, F.Y. Characterizing the Propagation of Situational Information in Social Media during COVID-19 Epidemic: A Case Study on Weibo. *IEEE Trans. Comput. Soc. Syst.* **2020**, *7*, 556–562. [CrossRef]

60. Lampos, V.; Moura, S.; Yom-Tov, E.; Cox, I.J.; McKendry, R.; Edelstein, M. Tracking COVID-19 Using Online Search. *arXiv* **2020**, arXiv:2003.08086.

61. Boberg, S.; Quandt, T.; Schatto-Eckrodt, T.; Frischlich, L. Pandemic Populism: Facebook Pages of Alternative News Media and the Corona Crisis–A Computational Content Analysis. *arXiv* **2020**, arXiv:2004.02566.

62. Jelodar, H.; Wang, Y.; Orji, R.; Huang, H. Deep Sentiment Classification and Topic Discovery on Novel Coronavirus or COVID-19 Online Discussions: NLP Using LSTM Recurrent Neural Network Approach. *arXiv* **2020**, arXiv:2004.11695.

63. Liu, D.; Clemente, L.; Poirier, C.; Ding, X.; Chinazzi, M.; Davis, J.T.; Vespignani, A.; Santillana, M. A Machine Learning Methodology for Real-time Forecasting of the 2019-2020 COVID-19 Outbreak Using Internet Searches, News Alerts, and Estimates from Mechanistic Models. *arXiv* **2020**, arXiv:2004.04019.

64. Hou, Z.; Du, F.; Jiang, H.; Zhou, X.; Lin, L. Assessment of Public Attention, Risk Perception, Emotional and Behavioural Responses to the COVID-19 Outbreak: Social Media Surveillance in China. *medRxiv Prepr.* **2020**. [CrossRef]

65. Stokes, D.C.; Andy, A.; Guntuku, S.C.; Ungar, L.H.; Merchant, R.M. Public Priorities and Concerns Regarding COVID-19 in an Online Discussion Forum: Longitudinal Topic Modeling. *J. Gen. Intern. Med.* **2020**. [CrossRef] [PubMed]

66. Shen, C.; Chen, A.; Luo, C.; Liao, W.; Zhang, J.; Feng, B. Reports of Own and Others' Symptoms and Diagnosis on Social Media Predict COVID-19 Case Counts in Mainland China. *arXiv* **2020**, arXiv:2004.06169.

67. Chen, Q.; Min, C.; Zhang, W.; Wang, G.; Ma, X.; Evans, R. Unpacking the Black Box: How to Promote Citizen Engagement Through Government Social Media During the COVID-19 Crisis. *Comput. Hum. Behav.* **2020**, 106380. [CrossRef]

68. Lucas, B.; Elliot, B.; Landman, T. Online Information Search During COVID-19. *arXiv* **2020**, arXiv:2004.07183.

69. Pekoz, E.A.; Smith, A.; Tucker, A.; Zheng, Z. COVID-19 Symptom Web Search Surges Precede Local Hospitalization Surges. *SSRN Prepr.* **2020**. [CrossRef]

70. Ellis, B.; Wong, W.H. Learning Causal Bayesian Network Structures from Experimental Data. *J. Am. Stat. Assoc.* **2008**, *103*, 778–789. [CrossRef]

71. Koller, D.; Friedman, N. *Probabilistic Graphical Models: Principles and Techniques*; MIT Press: Cambridge, MA, USA, 2009.

72. Rubin, D.B. Causal Inference Using Potential Outcomes: Design, Modeling, Decisions. *J. Am. Stat. Assoc.* **2005**, *100*, 322–331. [CrossRef]

73. Pearl, J. An Introduction to Causal Inference. *Int. J. Biostat.* **2010**, *6*. [CrossRef]

74. Pearl, J. *Causality*; Cambridge University Press: Cambridge, UK, 2009.

75. Twitter. Available online: https://twitter.com/ (accessed on 12 May 2020).

76. Dowd, J.B.; Andriano, L.; Brazel, D.M.; Rotondi, V.; Block, P.; Ding, X.; Liu, Y.; Mills, M.C. Demographic Science Aids in Understanding the Spread and Fatality Rates of COVID-19. *Proc. Natl. Acad. Sci. USA* **2020**, *117*, 9696–9698. [CrossRef] [PubMed]

77. Guo, Y.R.; Cao, Q.D.; Hong, Z.S.; Tan, Y.Y.; Chen, S.D.; Jin, H.J.; Tan, K.S.; Wang, D.Y.; Yan, Y. The Origin, Transmission and Clinical Therapies on Coronavirus Disease 2019 (COVID-19) Outbreak-An Update on the Status. *Mil. Med. Res.* **2020**, *7*, 1–10. [CrossRef] [PubMed]

78. Yang, X.; Yu, Y.; Xu, J.; Shu, H.; Liu, H.; Wu, Y.; Zhang, L.; Yu, Z.; Fang, M.; Yu, T.; et al. Clinical Course and Outcomes of Critically Ill Patients with SARS-CoV-2 Pneumonia in Wuhan, China: A Single-centered, Retrospective, Observational Study. *Lancet Respir. Med.* **2020**, *8*, 475–481. [CrossRef]

79. Wang, W.; Tang, J.; Wei, F. Updated Understanding of the Outbreak of 2019 Novel Coronavirus (2019-nCoV) in Wuhan, China. *J. Med. Virol.* **2020**, *92*, 441–447. [CrossRef] [PubMed]

80. WHO. *Report of the WHO-China Joint Mission on Coronavirus Disease 2019 (COVID-19)*; World Health Organization: Geneva, Switzerland, 2020.

81. Li, C.; Ji, F.; Wang, L.; Hao, J.; Dai, M.; Liu, Y.; Pan, X.; Fu, J.; Li, L.; Yang, G.; et al. Asymptomatic and Human-to-Human Transmission of SARS-CoV-2 in a 2-Family Cluster, Xuzhou, China. *Emerg. Infect. Dis.* **2020**, *26*, 1626–1628. [CrossRef] [PubMed]

82. World Bank Open Data—Population Ages 65 and Above. Available online: https://data.worldbank.org/ (accessed on 12 May 2020).

83. Distribution of Households by Household Type from 2003 Onwards—EU-SILC Survey. Available online: https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=ilc_lvph02&lang=en (accessed on 12 May 2020).

84. Social Media Stats-February 2020. Available online: https://gs.statcounter.com/ (accessed on 12 May 2020).

85. National Responses to the COVID-19 Pandemic—Lockdown Data. Available online: https://en.wikipedia.org/wiki/National_responses_to_the_COVID-19_pandemic (accessed on 12 May 2020).

86. Sanh, V.; Debut, L.; Chaumond, J.; Wolf, T. DistilBERT, A Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter. *arXiv* **2019**, arXiv:1910.01108.

87. Zheng, X.; Aragam, B.; Ravikumar, P.; Xing, E.P. DAGs with NO TEARS: Continuous Optimization for Structure Learning. In Proceedings of the 32nd International Conference on Neural Information Processing Systems, Montreal, QC, Canada, 3–8 December 2018; pp. 9492–9503. [CrossRef]

88. Chickering, D.M. Learning Bayesian Networks is NP-complete. In *Learning from Data*; Springer: Berlin/Heidelberg, Germany, 1996; pp. 121–130. [CrossRef]

89. Chickering, D.M.; Heckerman, D.; Meek, C. Large-sample Learning of Bayesian Networks is NP-hard. *J. Mach. Learn. Res.* **2004**, *5*, 1287–1330.

90. Wise, T.; Zbozinek, T.D.; Michelini, G.; Hagan, C.C.; Mobbs, D. Changes in Risk Perception and Self-reported Protective Behaviour During the First Week of the COVID-19 Pandemic in the United States. *R. Soc. Open Sci.* **2020**, *7*, 200742. [CrossRef]

91. Zhong, B.L.; Luo, W.; Li, H.M.; Zhang, Q.Q.; Liu, X.G.; Li, W.T.; Li, Y. Knowledge, Attitudes, and Practices Towards COVID-19 Among Chinese Residents During the Rapid Rise Period of the COVID-19 Outbreak: A Quick Online Cross-sectional Survey. *Int. J. Biol. Sci.* **2020**, *16*, 1745. [CrossRef]

92. Merchant, R.M.; Lurie, N. Social Media and Emergency Preparedness in Response to Novel Coronavirus. *J. Am. Med. Assoc. (JAMA)* **2020**, *323*. [CrossRef] [PubMed]