

When No News is Bad News. News-Based Change Detection during COVID-19

Kristoffer L. Nielbo^{1,2}, Frida Hæstrup¹, Kenneth C. Enevoldsen¹, Peter B. Vahlstrup¹, Rebekah B. Baglini², and Andreas Roepstorff²

¹Center for Humanities Computing Aarhus, Aarhus University, Denmark

²Interacting Minds Centre, Aarhus University, Denmark

Abstract

During the first wave of COVID-19, a peculiar phenomenon could be observed in the flow of news media content both within and between media platforms: as the pandemic spread, news media transformed into *Corona-news*. The information corollary of the content alignment, was that the novelty of news content went down as media focused monotonically on the pandemic event, which resulted in the, from a news media’s perspective, paradoxical situation where the same news was repeated over and over. This information phenomenon, where a decrease in novelty persists, has previously been used to track change in media, but in this paper we use a Bayesian approach to test the claim that a persistent change in novelty can be used to detect change in news media content originating in negative events.

Keywords: newspapers; pandemic response; Bayesian change detection; information theory

1 Introduction

A peculiar behavior could be observed in news media when the first wave of COVID-19 virus spread across the world. In response to the pandemic event, the ordinary rate of change in news content was disrupted because nearly every story became associated with COVID-19. On the one hand, content novelty went down, because nearly every story became more similar to previous stories (i.e., news suddenly became ‘Corona news’), but on the other hand, the COVID-19 association became more prevalent, resulting in, at least initially, an increase in content persistence. A recent study (Nielbo et al., 2021b) argues that this behavior is an example of the *news information decoupling* (NID) principle, according to which information dynamics of news media are (initially) decoupled by temporally extended catastrophes such that the content novelty

decreases as media focus monotonically on the catastrophic event, but the resonant property of said content increases as its continued relevance propagate throughout the news information system. The same study further indicated that NID can be used to detect significant change in news media that originate in catastrophic events. From the perspective of cultural dynamics, the COVID-19 pandemic provides a natural experiment that allows us to study the effect of a global catastrophe on the the dynamics of news media's information. While news media are neither unbiased nor infallible as sources of events, they do reflect preferences, values, and desires of a wide socio-cultural and political user spectrum. As such, news media coverage of COVID-19 functions as a proxy for how cultural information systems respond to unexpected and dangerous events.

Several studies have shown that the associative structure of news media is sensitive to socio-cultural dynamics, for instance the rise and intricacies of modernity as reflected in historical newspapers (Guldi, 2019; van Eijnatten and Ros, 2019), and value-based differential response to negative events such as instability and war (Daems et al., 2019). In a similar vein it has been shown that the ordered one-dimensional representation of the word co-occurrence structure quite accurately captures historically relevant trends in newspapers (Newman and Block, 2006). By embedding the co-occurrence structure in a low-dimensional space, it has been shown that newspaper content reflect fundamental cultural movements and understandings from the 19th century onward (van Eijnatten and Ros, 2019), and that these context-dependent representations are sensitive to cultural bias as reflected in newspapers (Wevers, 2019). In continuation of the 'Culturomics' trend that used Google Books to show how lexical variation is sensitive to events (Michel et al., 2011), a wide range of studies has demonstrated that simple word and concept frequencies are sufficient for robust offline detection of major historical events (Kestemont et al., 2014) and can be used to model the evolution of complicated cultural processes such as the historical interdependencies between media and politics (Bos et al., 2016). Fluctuations of time dependent word frequencies have been shown to discriminate between classes of events that have class-specific fractal signatures, where the social-cultural class display non-stationary and on-off intermittent behavior (Gao et al., 2012). Even within the social-cultural class, different types of events (or stories about events) seem to show fine-grained differences in their degree of self-affinity in newspapers (Wevers et al., 2020).

In line with recent developments in information theory, studies have used information theoretic measures to track the states and dynamics of socio-cultural systems as reflected in lexical data (Barron et al., 2018; Guldi, 2019; Murdock et al., 2015; Nguyen et al., 2020; Nielbo et al., 2019b). One paradigmatic study used relative entropy to study the development of Darwin's thinking in relation to his cultural context (Murdock et al., 2015). Both Shannon entropy and relative entropy have similarly been used in other studies to detect changes in prevalent mental states due to the socio-cultural context (e.g., state censorship, degree of recognition, religious observation) (Nielbo et al., 2019a,b). One specific information theoretic approach applies windowed relative entropy to dense low-dimensional text representations in order to generate signals that capture information *novelty*, \mathcal{N} , as a documents reliable content difference from the past; *transience*, \mathcal{T} , the documents content difference from future documents; and *resonance*, \mathcal{R} , the difference between *novelty* and *transience*, or conceptually, the degree to which future information conforms to a documents *novelty* (Barron et al., 2018; Nguyen et al., 2020). Taking a more dynamic perspective on this approach, one study has shown that discussion boards on social media where the novelty signal displays

both short-range correlations only and a particularly strong association with resonance are more likely to contain trending content (Nielbo et al., 2021a). Using the same approach combined with event detection, has also been shown to reliably predict major change points in historical data (Vrangbæk and Nielbo, 2021).

On the specific intersection between news media, COVID-19 and uncertainty, a group of economists has developed an index for economic policy uncertainty based on dense probabilistic representations of newspaper articles (Bess et al., 2020). The index correlates with existing market indices (e.g., VIX and BBO) and can accurately identify phase one of COVID-19 as well as other events associated with increased economic uncertainty. By using similar newspaper representations of front pages during COVID-19 and applying the above mentioned information theoretical approach, one study has argued that the information dynamics of news during COVID-19 were reflective of societal and value-based responses to the pandemic, and further argued that news media’s response to the pandemic is a decoupling of the news content’s novelty and resonance (i.e., news information decoupling) and, furthermore, that the decoupling may reflect political alignment with the current government (Nielbo et al., 2021b).

This study specifically tests the claim of Nielbo et al. (2021b), that NID-like behavior can provide input for change point detection algorithms. Specifically, we propose to test the claim that two change points are observable in news media during the first phase of COVID-19, *Lockdown* and *Opening* respectively, using a Bayesian approach to change point detection.

2 Methods

2.1 Data and Normalization

The data set consists of the linguistic content (title and body text) from front pages¹ of six Danish newspapers (Berlingske, BT, Ekstrabladet, Jyllands-Posten, Kristeligt Dagblad, and Politiken), see Table 1. All newspapers are national and published daily, with the exception of Kristeligt Dagblad that is only published six times per week from Monday to Saturday. Kristeligt Dagblad is kept in the sample because it is a national newspaper with a substantial circulation and represents a specific Danish reader segment. All newspapers were sampled from December 1, 2019 to June 30 2020 resulting in a corpus of 1,271,004 tokens.

In order to normalize linguistic content, numerals and highly frequent function words were removed, and the remaining data were casefolded and lemmatized using a language-specific neural model trained on the Danish Dependency Treebank (Qi et al., 2020). Subsequently, the data were represented as a bag-of-words (BoW) model using latent Dirichlet allocation in order to generate a dense low-rank representation for each front page. Parameter sweep was used for hyperparameter optimization and leave-p-out cross-validation was used for testing generalization. Note that with appropriate modifications to the divergence measures, see equations (4) and (5) in the following section, the approach to change detection in text presented in this paper is applicable to any probabilistic or geometric vector-representation of documents. Novelty and resonance were estimated for windows of seven days, $w = 7$, representing the weekly news cycle.

¹ Front pages are used because they condense the most important news content and are, in comparison to full newspapers, more similar across conditions. Qualitatively similar results can be obtained from full newspapers although the results are subject to considerable more noise.

For validation purpose we used the official dates for the first lockdown March 13 and opening April 15, 2020. While it is uncontroversial to assume that the lockdown date was point-like, the opening was actually more gradual with several restrictions remaining and local lockdowns. The change detection model proposed is however sensitive to this gradual change in the interval of τ (see the next sections).

Source	Type	Political alignment
Berlingske	compact	center-right
BT	tabloid	center-right
Ekstrabladet	tabloid	independent
Jyllands-Posten	compact	center-right
Kristeligt Dagblad	broadsheet	independent evangelical
Politiken	broadsheet	center-left

Table 1: Danish newspaper data set. Column one contains the name of the newspaper, column two the type of newspaper (Broadsheet, Compact, or Tabloid), and three approximate political alignment of the newspaper. It is important to note that the newspapers do not have direct affiliation with political parties today and that the alignment reflect their own classification. In some cases, this may not represent the perception of the readers.

2.2 Novelty and Resonance

Two related information signals were extracted from the temporally sorted BoW model: *Novelty* (\mathcal{N}) as an article $s^{(j)}$'s reliable difference from past articles $s^{(j-1)}, s^{(j-2)}, \dots, s^{(j-w)}$ in window w :

$$\mathcal{N}_w(j) = \frac{1}{w} \sum_{d=1}^w JSD(s^{(j)} | s^{(j-d)}) \quad (1)$$

and *resonance* (\mathcal{R}) as the degree to which future articles $s^{(j+1)}, s^{(j+2)}, \dots, s^{(j+w)}$ conform to article $s^{(j)}$'s novelty:

$$\mathcal{R}_w(j) = \mathcal{N}_w(j) - \mathcal{T}_w(j) \quad (2)$$

where \mathcal{T} is the *transience* of $s^{(j)}$:

$$\mathcal{T}_w(j) = \frac{1}{w} \sum_{d=1}^w JSD(s^{(j)} | s^{(j+d)}) \quad (3)$$

The novelty-resonance model was originally proposed in Barron et al. (2018), but here we apply a symmetrized and smooth version by using the Jensen–Shannon divergence (JSD):

$$JSD(s^{(j)} | s^{(k)}) = \frac{1}{2}D(s^{(j)} | M) + \frac{1}{2}D(s^{(k)} | M) \quad (4)$$

with $M = \frac{1}{2}(s^{(j)} + s^{(k)})$ and D is the Kullback-Leibler divergence:

$$D(s^{(j)} | s^{(k)}) = \sum_{i=1}^K s_i^{(j)} \times \log_2 \frac{s_i^{(j)}}{s_i^{(k)}} \quad (5)$$

2.3 Nonlinear Adaptive Filtering

Nonlinear adaptive filtering is applied to the information signals because of their inherent noisiness (Gray, 2007). First, the signal is partitioned into segments (or windows) of length $w = 2n + 1$ points, where neighboring segments overlap by $n + 1$. The time scale is $n + 1$ points, which ensures symmetry. Then, for each segment, a polynomial of order D is fitted. Note that $D = 0$ means a piece-wise constant, and $D = 1$ a linear fit. The fitted polynomial for i th and $(i + 1)$ th is denoted as $y^{(i)}(l_1), y^{(i+1)}(l_2)$, where $l_1, l_2 = 1, 2, \dots, 2n + 1$. Note that the length of the last segment may be shorter than w . We use the following weights for the overlap of two segments.

$$y^{(c)}(l_1) = w_1 y^{(i)}(l + n) + w_2 y^{(i+1)}(l), l = 1, 2, \dots, n + 1 \quad (6)$$

where $w_1 = (1 - \frac{l-1}{n}), w_2 = 1 - w_1$ can be written as $(1 - \frac{d_j}{n}), j = 1, 2$, where d_j denotes the distance between the point of overlapping segments and the center of $y^{(i)}, y^{(i+1)}$. The weights decrease linearly with the distance between point and center of the segment. This ensures that the filter is continuous everywhere, which ensures that non-boundary points are smooth.

2.4 Change Point Detection

In order to model media response to COVID-19, we follow Nielbo et al. (2021a) and state that it is only when a reliable difference in \mathcal{N} can be observed that we inspect the relationship between \mathcal{N} and \mathcal{R} . We propose a simple Bayesian approach to model changes in the mean of \mathcal{N} , a Bayesian mean-shift model. The simplest case of news information decoupling is a temporary state change corresponding to Denmark open \rightarrow lockdown \rightarrow open. This model implies three means, two of which are approximately identical, and two shifts, decoupling start and decoupling end, in \mathcal{N} , corresponding the beginning and the end of the lockdown respectively. Or more formally, we assume that the time series \mathcal{N} contains two change points in time, τ_1 and τ_2 , that can be located anywhere in the signal. While the last part of the the assumption clearly disregards information about the occurrence of the lockdown, it simplifies the model and is sufficient to detect one reliable decoupling, decoupling start and decoupling end, in \mathcal{N} . Aside from the change points, we assume that the series is stable and follow a normal distribution with varied mean, μ_i , and singular variance, σ . Notice that these assumptions follow from a mean-shift model. This gives us the following model for the observed \mathcal{N}_i :

$$\mathcal{N}_t = \begin{cases} \text{Normal}(\mu_1, \sigma) & \text{for } t < \tau_1 \\ \text{Normal}(\mu_2, \sigma) & \text{for } \tau_1 \leq t < \tau_2 \\ \text{Normal}(\mu_3, \sigma) & \text{for } t \geq \tau_2 \end{cases} \quad (7)$$

for which we wish to estimate the location of the change points τ_i , and the value of the means μ_i and variance σ , resulting in the following posterior:

$$P(\mu_i, \sigma, \tau_i | \mathcal{N}_t) = P(\mu_1, \mu_2, \mu_3, \sigma, \tau_1, \tau_2 | \mathcal{N}_t) \quad (8)$$

For estimation of the posterior, we used NUTS sampling with 4000 samples (Salvatier et al., 2015). Again, the estimation was done using slightly conservative priors assuming that the change points, τ_i , can be anywhere in the sequence (with $\tau_2 > \tau_1$) and that the variance, σ , is stable across change points. Note that the half Cauchy prior

distribution has several beneficial properties most importantly fat tail which allows for extreme values (Gelman et al., 2013; Polson and Scott, 2012). These assumptions were modelled using the following priors:

$$\begin{aligned}
\mu_i &\sim \text{Normal}(0, 0.5) \\
\sigma &\sim \text{Half Cauchy}(0.5) \\
\tau_1 &\sim \text{Uniform}(0, \max(\mathcal{N}_t)) \\
\tau_2 &\sim \text{Uniform}(\tau_1, \max(\mathcal{N}_t))
\end{aligned} \tag{9}$$

In order for the model to be valid, that is detect the change points in \mathcal{N} , the following is required for : 1) the relevant dates (lockdown and opening) should be contained in the estimate intervals (i.e., the actually events should detected); and 2) the intervals of the decoupling start and decoupling end should be non-overlapping (i.e., there model should detect a reliable state change).

2.5 News Information Decoupling

Finally, in order to describe the information states before and after an event and confirm if a change point reflects a decoupling (i.e., that novelty decreases while resonance increases), we fit resonance on novelty to estimate the \mathcal{NR} slope β_1 before and during the event in question (e.g., the Danish lockdown):

$$\mathcal{R}_i = \beta_0 + \beta_1 \mathcal{N}_i + \epsilon_i, \quad i = 1, \dots, n. \tag{10}$$

where β_0 is the intercept and ϵ is a random variable representing the errors of the fit.

3 Results

To establish a baseline for novelty and resonance, we computed the per newspaper linear slope for resonance on novelty (\mathcal{NR}) from December 01, 2019 to February 26, 2020 (the first incidence of COVID-19 in Denmark was registered on February 27, 2020). As can be observed from Figure 1, the slopes are remarkably similar, indicating a medium to strong association between novelty and resonance ($M = 0.56$, $SD = 0.06$) before the national outbreak of COVID-19. In the normal state of affairs, novelty and resonance therefore seems to be coupled such that novel news items resonate more than overused and repetitive items and vice versa. This general news dynamic confirms the intuition that news media, all things being equal, maintain their relevance by propagating news.

Figure 2 displays a prototypical example of NID during the first phase of COVID-19 (Nielbo et al., 2021b). Although COVID-19 news items date back to December 2019, 'Wuhan', newspaper content is not impacted until the period after the first national outbreak (in this case in Denmark). 'Virus'. From the phase 1 lockdown 'Lockdown' to the opening, 'Opening', the newspaper shows a valley in novelty and, initially, a peak in resonance until both processes approximately return to normal after the opening. Figure 3 shows the same trend, but with a noticeable difference in the time it takes *Berlingske* to return to a normal state of affairs and (i.e., 'Opening' then is less pronounced).

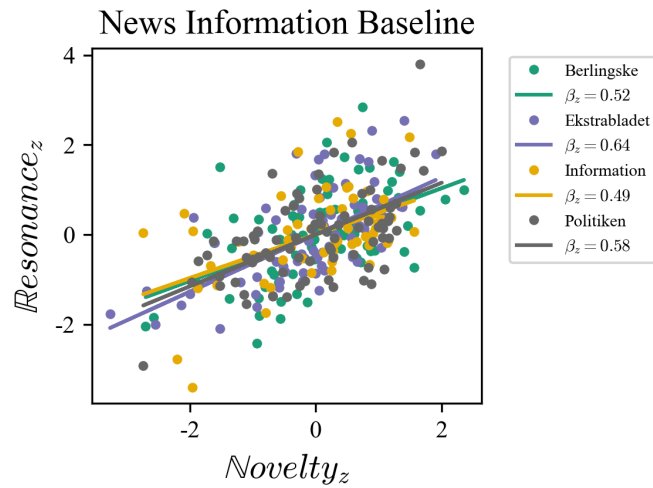


Figure 1: \mathcal{NR} slope baseline for four national newspapers that represents the left-right political spectrum. Data are sampled before COVID-19 phase 1 in Denmark initiated. Information has been included in the graph, which is a left-wing broadsheet newspaper.




Figure 2: Novelty (upper panel) and resonance (lower panel) for the center-left newspaper *Politiken* before and during COVID-19 phase 1. Trend lines in the upper and lower panel are estimated using a nonlinear adaptive filter for three different levels of smoothing to emphasize the robustness of the trend, see Equation 6.




Figure 3: Novelty (upper panel) and resonance (lower panel) for the center-right newspaper *Berlingske* before and during COVID-19 phase 1. Trend lines in the upper and middle panel are estimated using a nonlinear adaptive filter, see Equation 6.

To validate the observed behavior, we tested for two change points in novelty using a Bayesian model. The first change point, ‘NID Start’ should separate pre-lockdown from lockdown centered on week 11 (March 9-15), and the second lockdown, ‘NID End’ from post opening (centered on week 16, April 13-19). Table 2 shows the estimated change points for six national newspapers, two of which are *Tabloid* newspapers (Class) and the remainder *Broadsheet/Compact*. From the model, it can be observed that all broadsheet/compact newspapers seem to support the NID principle in novelty. The first change point is placed in weeks 10-11, the second, however, is more a matter of contention. The opening change point lies within April and displays a month’s delayed response. Finally, it can be observed that tabloid press shows no indication of NID behavior. Table 3 shows the posterior distributions for novelty, their means and highest density intervals, for broadsheet/compact newspapers, which clearly indicates that they do conform to NID.

Source	Type	NID Start	NID End	NID
Berlingske	C	03.07 [03.03, 03.09]	04.28 [04.09, 05.08]	True
BT	T	04.10 [12.30, 09.01]	07.25 [04.22, 09.03]	False
Ekstrabladet	T	01.28 [01.02, 03.17]	05.08 [01.16, 07.22]	False
Jyllands-Posten	C	03.10 [03.08, 03.14]	05.25 [05.21, 06.06]	True
Kristligt Dagblad	B	03.07 [03.05, 03.12]	04.15 [04.11, 04.17]	True
Politiken	B	03.13 [03.12, 03.13]	04.08 [04.05, 04.08]	True

Table 2: Estimated temporal change points at 94% high density intervals for novelty. Column one contains the name of the newspaper, columns two its type (*Broadsheet*, *Compact Tabloid*), NID Start and End is the beginning and end of the lockdown as represented in the newspaper, and the final column indicated if the specific source supported the NID principle.

Source	\mathcal{N}_{pre}	\mathcal{N}_{NID}	\mathcal{N}_{post}
Berlingske	0.36 [0.35, 0.37]	0.29 [0.27, 0.31]	0.34 [0.34, 0.35]
Jyllands-Posten	0.29 [0.28, 0.30]	0.23 [0.22, 0.24]	0.27 [0.26, 0.28]
Kristligt Dagblad	0.27 [0.26, 0.28]	0.19 [0.18, 0.21]	0.26 [0.25, 0.27]
Politiken	0.27 [0.26, 0.28]	0.15 [0.14, 0.17]	0.26 [0.25, 0.26]

Table 3: Estimates of mean \mathcal{N} values at 94% high density intervals before during and after the lockdown for the four broadsheet newspapers that supported the NID principle, see Table 1. All newspapers show a reliable reduction of \mathcal{N} during the lockdown.

That novelty decreases during a catastrophic event is nevertheless only half the story. For NID to be supported by the data, resonance should increase during the lockdown such that the medium to strong association between novelty and resonance is momentarily weakened. Following Nielbo et al. (2021b) we inspected the time-windowed linear fits of resonance on novelty, \mathcal{NR} slopes, in order to confirm this, see Figure 4. All broadsheet newspapers display a slope decrease during the lockdown, thereby conforming to the NID principle 4. Tabloids on the other hand, follow an inverse pattern, such that the \mathcal{NR} slope increases during the lockdown period.

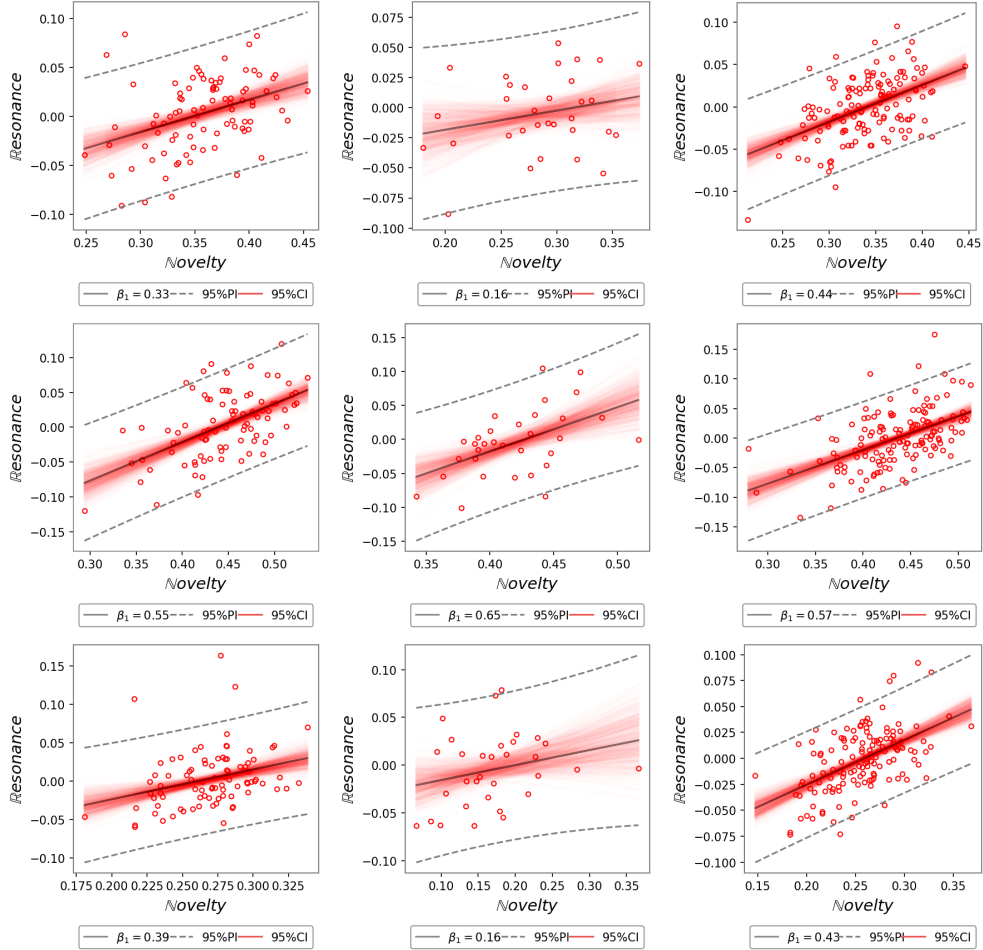


Figure 4: \mathcal{NR} slopes before during and after the lockdown for Berlingske (upper row), Ekstrabladet (middle row), and Politiken (lower row) during COVID-19 phase 1. Punctured lines represents the 95% prediction intervals, and semi-transparent lines the 95% confidence intervals using a bootstrap approach with 500 samples.

Source	Type	\mathcal{NR}_{pre}	\mathcal{NR}_{NID}	\mathcal{NR}_{post}
Berlingske	C	0.33 [0.17, 0.51]	0.16 [-0.07, 0.38]	0.44 [0.32, 0.58]
BT	T	0.49 [0.29, 0.66]	0.55 [0.28, 0.83]	0.26 [0.08, 0.43]
Ekstrabladet	T	0.55 [0.38, 0.72]	0.65 [0.26, 1]	0.57 [0.42, 0.71]
Jyllands-Posten	C	0.42 [0.24, 0.63]	0.31 [0.04, 0.56]	0.39 [0.25, 0.51]
Kristligt Dagblad	B	0.57 [0.34, 0.78]	0.43 [0.06, 0.78]	0.76 [0.55, 0.95]
Politiken	B	0.39 [0.14, 0.61]	0.16 [-0.05, 0.37]	0.43 [0.32, 0.54]

Table 4: \mathcal{NR} coefficients at 95% confidence intervals before during and after the lockdown for all newspapers in the sample. Column two contains the newspaper type (Broadsheet, Compact Tabloid).

4 Discussion

This study has sought to validate the news information decoupling (NID) principle on a sample of six national newspapers from Denmark during the first phase of COVID-19. Using a Bayesian approach to change point detection, we showed that content novelty in broadsheet newspapers does indeed display statistically reliable points of change coinciding with the COVID-19 lockdown and opening. NID was further corroborated by the \mathcal{NR}_{pre} slopes that indicated a decoupling of resonance from novelty during the lockdown. Several direct observations can be made from the findings. First, the estimated change points for the ‘Pre-lockdown \rightarrow Lockdown’ are spread over a two week interval, which indicates that a lockdown could be reasonably predicted already from the first COVID-19 incident in Denmark. Second, in a similar vein, the ‘Lockdown \rightarrow Opening’ change point interval is spread over a full month from April 8 to May 8, which may reflect disagreement about if and when the lockdown ended. The Danish government, during the first phase of COVID-19, was center-left and the model’s uncertainty in determining the opening may reflect political alignment of the newspapers (Nielbo et al., 2021b), where center-right newspapers (e.g., Berlingske and Jyllands-Posten) were more sceptical towards the government’s implementation of an opening than the center-left (e.g., Politiken). In other words, the center-right may have been more reluctant to acknowledge the opening as a return to normal. Third, tabloid newspapers did not show any indication of a news decoupling. On the contrary, their \mathcal{NR}_{pre} slopes momentarily increased during the lockdown. This increase in slopes does, however, not provide any useful information, because, as shown by our change point detection model, the periodization is not meaningful for the two tabloid newspapers.

As already mentioned, validation of the NID principle is still needed for multilingual data and its value for crisis management should be further tested. For change detection, the scope of the principle needs additional testing; does NID generalize beyond a small set of negative events to, for instance, temporally extended significant events (e.g., moon landing, fall of the Berlin Wall). Finally, several contrast already hinted at need to be tested, e.g., left vs. right-wing newspapers, tabloid vs. broadsheet newspapers, silly season and other seasonal effects, are all interesting venues for media and journalism researchers.

We advise caution in developing rich (domain-specific) interpretations of the model’s behavior, because the behavior primarily shows that something is happening (i.e., change detection), but not what is happening (i.e., change characterization) during the lockdown. There are however two possible interpretations of NID that we would like to foreground in this context: an unmediated and a mediated interpretation

of the observed NID behavior. According to the unmediated interpretation, NID is a context-independent response to negative events, where news focus monotonically on the event in question (e.g., COVID-19 lockdown), thereby lowering the overall uncertainty or unpredictability of the regular news cycle. The unmediated account does not imply that NID is a deterministic response to negative events as such. It is possible that NID only captures a class of events (e.g., temporally extended catastrophic events) similar to how the persistence of information processes can classify event types (Gao et al., 2012; Wevers et al., 2020). Against this account is the observed differences between newspaper types, the broadsheet/compact vs. tabloid contrast. If NID is truly unmediated, we are likely to expect a more uniform response from all newspaper types. According to the alternative account, NID is a mediated response to negative events given certain cultural and societal conditions. Mediating factors could be societal coordination mechanisms reflected in complex social variables such as trust in government, social uncertainty, and economic equality. Denmark belongs to the Nordic group of universal welfare states that is characterized by high levels of trust in the government, and low levels of social uncertainty and inequality. It may be that NID reflects the close coordination between media, population and government under these specific conditions. The differential NID behavior observed in the contrast between political alignment, the centre-left vs centre-right contrast, may well support this interpretation. That centre-right newspapers were more reluctant than centre-left newspapers to acknowledge the return to a normal state of affairs, may very well indicate a negotiation of trust in government. To properly validate and decide between these interpretations, we need to develop additional models for relevant contrasts (e.g., news data from other Nordic countries).

From a more theoretical perspective, we would like to propose an interpretation of the combination of information-theoretical and change detection models as providing a potential state variable or indicator of societal uncertainty. To properly understand the behavior of complex socio-cultural systems in response to a catastrophe like COVID-19, we have to continuously monitor the states of said system. Ideally, we want to monitor all intrinsic variables related to uncertainty, and potentially trust and inequality given our above mentioned un-/mediated interpretations. This strategy however is not feasible and instead we suggest to rely on the fundamental embedding theorem of chaos theory, which states that the detailed dynamics of a system that has an underlying attractor can be readily studied by reconstructing a suitable phase space of a scalar time series recorded from the system (Packard et al., 1980; Sauer et al., 1991; Takens, 1981). Chaos theory offers an elaborate scheme for generating aperiodic, highly irregular data from a deterministic system that can be characterized by only very few state variables instead of a random system with infinite numbers of degrees of freedom (Gao and Xu, 2021). While the evolution of a complex social system may not be modeled as a dynamical system with a single attractor, we can assume that the dynamics of a large-scale social system can be approximated by switching between a large number of attractors, some of which may be simple, such as fixed points that may be associated with the dynamics of cultural information, while others may be complicated, including chaotic attractors (Gao et al., 2007; Ott, 2002). In order to understand societal uncertainty in face of the pandemic, we need to find an adequate continuous variable related to cultural information that is shared by members of the society during the event in question. We propose that a model of uncertainty in news media could very well provide such an adequate variable. On this account, the approach to news information dynamics presented here provides a valuable tool for

media-based indices that can supplement existing economic and policy-based indices of uncertainty (Bess et al., 2020).

5 Online Resources

All data are proprietary and have been collected through Infomedia's API: <https://infomedia.dk/>. For inquiries regarding models and derived data, please contact chcaa@cas.au.dk. The source code for methods is available on Github: <https://bit.ly/3beahFd>. More details on NID detection can be found at NeiC's NDHL website: <https://bit.ly/3bfeW9C>.

6 Acknowledgments

This research was supported the "HOPE - How Democracies Cope with COVID-19"-project funded by The Carlsberg Foundation with grant CF20-0044, NeiC's Nordic Digital Humanities Laboratory project, and DeiC Type-1 HPC with project DeiC-AU1-L-000001. The authors would like to thank Berlingske Media, JP/Politikens Hus, and Kristeligt Dagblad for providing access to proprietary data.

References

- Alexander T. J. Barron, Jenny Huang, Rebecca L. Spang, and Simon DeDeo. Individuals, institutions, and innovation in the debates of the french revolution. *Proceedings of the National Academy of Sciences*, 115(18):4607–4612, 2018. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1717729115.
- Mikkel Bess, Erik Grenestam, Alessandro Tang-Andersen Martinello, and Jesper Pedersen. Uncertainty and the real economy: Evidence from denmark. *Working Paper - Danmarks Nationalbank*, 165, 2020.
- Patrick Bos, Huub Wijfjes, Maaïke Piscaer, and Gerrit Voerman. Quantifying “pillarization”: Extracting political history from large databases of digitized media collections. *Proceedings of the 3rd HistoInformatics Workshop*, page 10, 2016.
- Joke Daems, Thomas D'haeninck, Simon Hengchen, Teclé Zere, and Christophe Verbruggen. ‘workers of the world’? a digital approach to classify the international scope of belgian socialist newspapers, 1885–1940. *Journal of European Periodical Studies*, 4(1):99–114, 2019. ISSN 2506-6587. doi: 10.21825/jeps.v4i1.10187.
- Jianbo Gao and Bo Xu. Complex Systems, Emergence, and Multiscale Analysis: A Tutorial and Brief Survey. *Applied Sciences*, 11(12):5736, June 2021. ISSN 2076-3417. doi: 10.3390/app11125736.
- Jianbo Gao, Yinhe Cao, Wen-wen Tung, and Jing Hu. *Multiscale Analysis of Complex Time Series: Integration of Chaos and Random Fractal Theory, and Beyond*. Wiley-Interscience, 1 edition edition, 2007. ISBN 978-0-471-65470-4.
- Jianbo Gao, Jing Hu, Xiang Mao, and Matjaž Perc. Culturomics meets random fractal theory: insights into long-range correlations of social and natural phenomena over the past two centuries. *Journal of The Royal Society Interface*, 9(73):1956–1964, 2012.

- Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis, Third Edition*. Chapman and Hall/CRC, 3 edition edition, 2013. ISBN 978-1-4398-4095-5.
- Katharine Lynn Gray. Comparison of Trend Detection Methods. *Graduate Student Theses, Dissertations, & Professional Papers*, 228:98, 2007.
- Jo Guldi. The measures of modernity: The new quantitative metrics of historical change over time and their critical interpretation. *International Journal for History, Culture and Modernity*, 7(1):899–939, 2019. ISSN 2666-6529, 2213-0624. doi: 10.18352/hcm.589.
- Mike Kestemont, Folgert Karsdorp, and Marten Düring. Mining the twentieth century’s history from the time magazine corpus. In *Proceedings of the 8th Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities (LaTeCH)*, pages 62–70. Association for Computational Linguistics, 2014. doi: 10.3115/v1/W14-0609.
- J.-B. Michel, Y. K. Shen, A. P. Aiden, A. Veres, M. K. Gray, The Google Books Team, J. P. Pickett, D. Hoiberg, D. Clancy, P. Norvig, J. Orwant, S. Pinker, M. A. Nowak, and E. L. Aiden. Quantitative analysis of culture using millions of digitized books. *Science*, 331(6014):176–182, 2011. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1199644.
- Jaimie Murdock, Colin Allen, and Simon DeDeo. Exploration and exploitation of victorian science in darwin’s reading notebooks. *arXiv:1509.07175*, 2015.
- David J. Newman and Sharon Block. Probabilistic topic decomposition of an eighteenth-century american newspaper. *Journal of the American Society for Information Science and Technology*, 57(6):753–767, 2006. ISSN 15322882, 15322890. doi: 10.1002/asi.20342.
- Dong Nguyen, Maria Liakata, Simon DeDeo, Jacob Eisenstein, David Mimno, Rebekah Tromble, and Jane Winters. How we do things with words: Analyzing text as social and cultural data. *Frontiers in Artificial Intelligence*, 3:62, 2020. ISSN 2624-8212. doi: 10.3389/frai.2020.00062.
- K.L. Nielbo, P.B. Vahlstrup, and A. Bechmann. Trend reservoir detection: Minimal persistence and resonant behavior of trends in social media. *Proceedings of Computational Humanities Research*, 1, 2021a.
- Kristoffer L Nielbo, Katrine F Baunvig, Bin Liu, and Jianbo Gao. A curious case of entropic decay: Persistent complexity in textual cultural heritage. *Digital Scholarship in the Humanities*, 34(3), 2019a. ISSN 2055-7671. doi: 10.1093/llc/fqy054.
- Kristoffer L. Nielbo, M. L. Perner, C. P. Larsen, Jonas Nielsen, and D. Laursen. Automated compositional change detection in saxo grammaticus’ gesta danorum. In *DHN*, pages 320–332, 2019b.
- Kristoffer L. Nielbo, Rebekah B. Baglini, Peter B. Vahlstrup, Kenneth C. Enevoldsen, Anja Bechmann, and Andreas Roepstorff. News information decoupling: An information signature of catastrophes in legacy news media. *arXiv:2101.02956 [cs]*, 2021b.
- Edward Ott. *Chaos in Dynamical Systems*. Cambridge University Press, 2 edition, 2002. ISBN 978-0-521-01084-9. doi: 10.1017/CBO9780511803260.

- N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw. Geometry from a time series. *Phys. Rev. Lett.*, 45(9):712–716, 1980. doi: 10.1103/PhysRevLett.45.712. Publisher: American Physical Society.
- Nicholas G. Polson and James G. Scott. On the half-cauchy prior for a global scale parameter. *Bayesian Analysis*, 7(4), 2012. ISSN 1936-0975. doi: 10.1214/12-BA730.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. Stanza: A python natural language processing toolkit for many human languages. *arXiv:2003.07082 [cs]*, 2020.
- John Salvatier, Thomas Wiecki, and Christopher Fonnesbeck. Probabilistic programming in python using PyMC. *arXiv:1507.08050 [stat]*, 2015.
- Tim Sauer, James A. Yorke, and Martin Casdagli. Embedology. *Journal of Statistical Physics*, 65(3):579–616, 1991. ISSN 1572-9613. doi: 10.1007/BF01053745.
- Floris Takens. Detecting strange attractors in turbulence. In David Rand and Lai-Sang Young, editors, *Dynamical Systems and Turbulence, Warwick 1980*, pages 366–381. Springer Berlin Heidelberg, 1981. ISBN 978-3-540-38945-3.
- Joris van Eijnatten and Ruben Ros. The eurocentric fallacy. a digital-historical approach to the concepts of ‘modernity’, ‘civilization’ and ‘europe’ (1840–1990). *International Journal for History, Culture and Modernity*, 7(1):686–736, 2019. ISSN 2666-6529, 2213-0624. doi: 10.18352/hcm.580.
- E.E.H Vrangbæk and K.L. Nielbo. Composition and change in de civitate dei: A case study of computationally assisted methods. *Studia Patristica*, 2021.
- Melvin Wevers. Using word embeddings to examine gender bias in dutch newspapers, 1950-1990. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 92–97. Association for Computational Linguistics, 2019. doi: 10.18653/v1/W19-4712.
- Melvin Wevers, Jianbo Gao, and Kirstoffer L. Nielbo. Tracking the consumption junction: Temporal dependencies between articles and advertisements in dutch newspapers. *Digital Humanities Quarterly*, 014(2), 2020.