

# Discovering Collective Narratives Shifts in Online Discussions

Wanying Zhao<sup>1</sup>, Siyi Guo<sup>2,3</sup>, Kristina Lerman<sup>2,3</sup>, Yong-Yeol Ahn<sup>1</sup>

<sup>1</sup>Center for Complex Networks and Systems Research, Luddy School of Informatics, Computing, and Engineering, Indiana University

<sup>2</sup>Department of Computer Science, University of Southern California

<sup>3</sup>Information Sciences Institute, University of Southern California  
{zhaowany, yyahn}@iu.edu, {siyiguo, lerman}@isi.edu

## Abstract

Narratives are foundation of human cognition and decision making. Because narratives play a crucial role in societal discourses and spread of misinformation and because of the pervasive use of social media, the narrative dynamics on social media can have profound societal impact. Yet, systematic and computational understanding of online narratives faces critical challenge of the scale and dynamics; how can we reliably and automatically extract narratives from massive amount of texts? How do narratives emerge, spread, and die? Here, we propose a systematic narrative discovery framework that fill this gap by combining change point detection, semantic role labeling (SRL), and automatic aggregation of narrative fragments into narrative networks. We evaluate our model with synthetic and empirical data—two Twitter corpora about COVID-19 and 2017 French Election. Results demonstrate that our approach can recover major narrative shifts that correspond to the major events.

## Introduction

Human beings are “storytelling animals” (Gottschall 2012). Narratives help us make sense of our experiences of reality (Richardson 1997; Bruner 1987; Roberts 2001), which enriches our emotions (e.g., Kopfman et al. 1998), influences our beliefs (e.g., Appel and Richter 2007; De Graaf et al. 2012), and affects our decisions (Winterbottom et al. 2008).

The online social media opened up a global public space where narratives emerge, mutate, compete, spread, and die. The online narratives are collectively created by numerous participants (Meraz and Papacharissi 2013; Papacharissi 2016) and can quickly spread globally. The potency of online narratives in society (cf. “Make America Great Again” or “Black Lives Matter”) calls for a thorough investigation into how narratives evolve on social media.

However, it is challenging to study online narrative evolution. First, it is a messy, large-scale process involving millions of people. Second, narratives are often *fragmented* inside short messages. For instance, Twitter’s 140-character limit (expanded to 280 in 2017) leaves little space for users to flush out a narrative (Hermida 2014; Sadler 2018). This

makes studying online narratives, particularly from tweets, challenging. Second, close reading (Moretti 2000) is labor intensive and does not scale well. Third, the concept of narrative is not well-defined and often difficult to operationalize.

Tackling these challenges, we ask:

- **RQ1:** How can we extract narrative fragments hidden in short posts?
- **RQ2:** How can we aggregate and synthesize narrative fragments into narratives that reflect collective attention?
- **RQ3:** How can we systematically detect the collective narrative shifts from a massive collection of social media conversations?

Our computational framework addresses these research questions in the following ways. First, we use change point detection (He, Burghardt, and Lerman 2022) to identify *collective narrative shifts*. Second, we operationalize *narrative fragments* (Dourish and Gómez Cruz 2018) (or *narrative elements* (Jing and Ahn 2021)), using Semantic Role Labeling (SRL) (Shi and Lin 2019). Third, we aggregate narrative fragments into a *narrative network* that captures more complete narratives. Now, let us introduce these three primary components of our method.

**Narrative fragment discovery** The most common approaches to narrative discovery are keyword-based and topic-modeling methods. Keyword-based focuses on sets of coherent individual tokens, or keywords (e.g., Lazard et al. 2015), as a narrative; Topic-modeling discovers “topics”—another proxy of narrative—each of which is a probability distribution over tokens. Latent Dirichlet Allocation (Blei, Ng, and Jordan 2003) and its variations (Yu et al. 2019; Yu and Qiu 2019), including a BERT-based model (Grootendorst 2022), and matrix factorization methods like NMF (Lee and Seung 2001) are commonly used. However, despite the great utility of these methods and the fact that keywords and topics are crucial elements of a narrative, they are still not equivalent to narratives.

One useful way to conceptualize how narratives unfold in social media is to consider each post as a *chronicle*; i.e., from a stream of events with timestamps (posts), which contain actors, motives, and narrative fragments (Dourish and Gómez Cruz 2018), although they may or may not

have direct logical or causal relationships between each other (Sadler 2018, 2021).

This is intimately related to *semantic role labeling (SRL)*, which identifies triplets of Action (a verb), Agent (who initiates the action) and Patient (the recipient of the action) from a sentence (Fillmore 1967). For example, the sentence “I love this coffee shop”, contains a triplet of [“love” (action), “I” (agent), “coffee shop” (patient)]. SRL co-advanced with the advancement of large language models, leading to powerful tools, such as SENNA (Collobert et al. 2011) and AllenNLP (Gardner et al. 2018). SRL has been shown to be able to extract relevant narrative fragments from politician’s tweets during the early stage of the pandemic (Jing and Ahn 2021); Similar methods were applied to extract actors and actions from blog posts and news articles about conspiracies (Tangherlini et al. 2020; Shahsavari et al. 2020). Following these studies, we operationalize a narrative fragment as a SRL triplet—(agent, action, patient).

**Change point detection** To understand the emergence and shifts in collective narratives, we need to study the temporality of collective narratives. The most basic aspect is the identification of significant inflection points where shifts in collective narratives happen.

Many research studied the detection of significant change points. For example, cumulative summation (CUSUM) (Page 1954) is a basic method that detects changes in time by assuming a normality of data distribution. However, this idea (e.g., Willsky and Jones 1976; Barber 2015) tends to be limited to univariate time series data and only detects single change point. Alternative methods built on statistical models, such as hidden Markov model (Raghavan, Galstyan, and Tartakovsky 2013) and Bayesian inference (Nieku et al. 2015; Wilson, Nassar, and Gold 2010) tend to provide higher-quality results but at the cost of computational complexity. Meta Change Point Detection (MtChD) (He, Burghardt, and Lerman 2022) is argued to be a state of the art method that handles sparse and noisy data well. It is computationally efficient and can handle high-dimensional datasets, while maintaining high accuracy. This study thus employs MtChD as the method for change point detection.

**Narrative Network** After identifying temporal contexts (the time periods defined by the change points) and extracting the narrative fragments that are prevalent in each period, we further synthesize these fragments to create *narrative networks*. The idea was originally proposed by Bearman and Stovel (Bearman and Stovel 2000), where they argued that network representation can link pieces of a narrative together. Although the idea was initially used to study individuals’ life stories, it can also be used to study public opinions. Past research has shown that, by constructing a narrative network, one could form a full story-line of the Bridgegate and Pizzagate (Tangherlini et al. 2020). Another study has shown the potential of applying community detection algorithm (Blondel et al. 2008) on narrative networks to uncover major COVID-19 topics (Shahsavari et al. 2020). Here, we construct narrative networks based on the narrative fragments extracted by SRL, then apply backbone filter-

ing (Serrano, Boguná, and Vespignani 2009) to extract and visualize the core of narrative networks. We also compare this approach with simpler statistical methods.

Now, we will describe our methods and lay out our findings from two Twitter datasets about COVID-19 and the 2017 French Election. We find that our method accurately captures the major events related to COVID-19, and identify collective narrative shifts associated with Le Pen and Macron following significant change points, which align with the election-related events. Additionally, we evaluate the robustness of our method with synthetic data and baseline methods.

## Methods

### Data

We use two Twitter datasets: a COVID-19 corpus (Chen et al. 2020) and a 2017 French Election corpus (Ferrara 2017). They cover different topics (politics and pandemic) and two languages (primarily in English and in French, respectively). Furthermore, the French election dataset features two key entities (i.e., Macron vs. Le Pen) while the COVID-19 corpus is not strongly anchored to contrasting entities. Apart from empirical data, we generate three versions of synthetic data to address 1) the robustness of the aggregation method and 2) the robustness of change point detection against noise, and 3) against overlap of collective narratives in time.

**COVID-19 Corpus** The COVID-19 corpus contains tweets about the COVID-19 pandemic collected by using keywords (Chen et al. 2020), capturing 600k tweets in English between January 21 to March 31, 2020.

**2017 French Election Corpus** This corpus includes keyword-identified tweets related to the 2017 French presidential Election (Ferrara 2017), and covers the period from April 26 to May 29, 2017. We focus on original tweets written in top three languages in the corpus—French, English, and Spanish, totaling in 2,438K tweets. We translate non-English tweets to English using the MarianMT model<sup>1</sup>.

We further create two separate, keyword-based French Election corpus—a 1) Macron corpus and a 2) Le Pen corpus. To identify the keywords referring to the candidates, we first trained a word2vec (Church 2017) model using the whole corpus. Then, we select top 20 keywords with the smallest cosine distance to each candidate’s name. After that, we iterate the same process for these top 20 keywords until the search was exhaustive. With the potential list of keywords, we manually examine and remove terms that do not refer to one of the candidates.

### Change Point Detection

We use MtChD (He, Burghardt, and Lerman 2022)<sup>2</sup>, a self-supervised change point detection method based on the concept of confusion-based training (Van Nieuwenburg, Liu, and Huber 2017), which considers changes in classification

<sup>1</sup>[https://huggingface.co/docs/transformers/model\\_doc/marian](https://huggingface.co/docs/transformers/model_doc/marian)

<sup>2</sup>[https://github.com/yuziheusc/confusion\\_multi\\_change](https://github.com/yuziheusc/confusion_multi_change)

accuracy as indicators of change points. This method works well on high-dimensional textual data, and it is robust to sparse signals and noise. This is very suitable for the social media data in our study. Specifically, given a randomly selected time point in the data as a trial change point, we train a classifier to predict each data point as before or after the trial change point. If we perform this prediction task at many different time points in the data, we should expect significant changes in the accuracy of this prediction task at true change points. Therefore, the method detects changes as differences in accuracy compared to a null model (predicting the majority class). To construct the classifier, we use TF-IDF with 5,000 most frequent words as the text embedding and compare two types of classification methods: a random forest and a neural network. The change points detected using both methods are very similar. Here we present the results using a random forest. In addition, MtChD also allows us to detect multiple change points by recursively partitioning data on discovered change points. The recursive segmentation has at most three levels and a minimum time length of four days.

### Semantic Role Labeling (SRL)

Because SRL operates on sentences, we first split tweets into sentences using the NLTK tokenizer (Loper and Bird 2002). Then, we adopt a re-implementation<sup>3</sup> of the BERT-based SRL model (Shi and Lin 2019) to extract narrative fragments. This model can further predict semantic roles of verbs captured in a sentence (i.e., the function and semantic meaning of a verb in the sentence as in Propbank standard<sup>4</sup>(Kingsbury and Palmer 2003)), which is missing in similar model provided by AllenNLP<sup>5</sup>. For each sentence, the model generates multiple triplets, each of which represents a narrative fragment in the form of a triplet (Argument 0 **A0**, Verb **V**, Argument 1 **A1**). **A0** represents the agent of an action (Verb); **A1** refers to the patient or theme of the Verb. This triplet altogether conveys “who did what” and “what did what,” which reveals the relationship between entities, such as (“you,” “vote,” “Macron”)—which captures the narrative fragments about ‘your’ voting for Macron (e.g., “You should vote for Macron”).

### Narrative Fragment Aggregation

We adopt a three-step process to aggregate narrative triplets (fragments). First, we map semantic roles from Propbank to VerbAtlas<sup>6</sup>, which helps compressing 5,649 verbs into 466 verb frames. For example, (“You,” “back.01,” “Macron”) and (“You,” “endorse.01,” “Macron”) can be aggregated into (“You,” “FOLLOW\_SUPPORT\_SPONSOR\_FUND,” “Macron”). With this mapping, we convert 422,019 unique narrative triplets (**A0**, Verb, **A1**) into 418,554 triplets (**A0**, (verb) frame, **A1**) in COVID data; and 2,016,058 unique ones into 1,253,182 in the French Election data.

<sup>3</sup><https://github.com/Riccorl/transformer-srl>

<sup>4</sup><https://propbank.github.io/>

<sup>5</sup>[https://docs.allennlp.org/models/main/models/structured\\_prediction/predictors/srl/](https://docs.allennlp.org/models/main/models/structured_prediction/predictors/srl/)

<sup>6</sup><https://verbatlas.org/>

Second, we consider a narrative triplet as a sentence (i.e. “**A0** frame **A1**”) and train its embedding with a pre-trained SentBert<sup>7</sup> model (“all-mpnet-base-v2”) to cluster them. We validate our approach—considering each triplet as a sentence—with the SemEval STS data<sup>8</sup>. We first select sentence pairs with reference similarity (ranges from 0 to 5, determined by human annotators from SemEval) at least four. We consider sentence pairs with high values indicating high semantic similarity. Then, we generate two sets of embeddings by the SentBert model, one on raw SemEval STS sentences, and the other on narrative triplets generated from raw SemEval STS sentences. The test result indicates that the narrative triplets (311/605) outperform the raw sentences (230/605) in predicting sentence similarity.

We then apply BIRCH (Zhang, Ramakrishnan, and Livny 1996) to cluster narrative triplet embeddings. The most frequent triplet in each cluster is considered the representative triplet.

Finally, we cluster **A0** and **A1** in the narrative triplets similarly to the second step. In sum, we convert the COVID data into 261,823 unique narrative triplets, and the French Election data into 738,737 unique triplets.

**Narrative Significance** After aggregating narrative triplets, we calculate the relevance of each narrative triplet within the time frames discovered by MtChD to identify the most important narrative fragments in each time frame. We adopt the log-odd ratios with informative Dirichlet priors (Monroe, Colaresi, and Quinn 2008; Jing and Ahn 2021) to estimate the relevance of each triplet:

$$s_w = \log \frac{f_t^T + f_t^B}{n^T + n^B - f_t^T + f_t^B} - \log \frac{f_t^R + f_t^B}{n^R + n^B - f_t^R + f_t^B}, \quad (1)$$

where  $f_t^C$  stands for frequency of a narrative triplet  $t$  in corpus  $C$ ;  $T$ ,  $R$ , and  $B$  refers to the target, reference, and background corpus respectively.  $n^C$  is the size of the corpus  $C$  in terms of the number of triplets. In the French Election data, for example, given the discrepancy between two candidates is more valuable,  $T$  would be Macron corpus within a time frame,  $R$  would be Le Pen corpus within the same interval, and  $B$  would be the whole French Election corpus within the same interval. In the case of the COVID data,  $R$  could be the data of previous time frame and  $B$  is the whole COVID corpus.

**Narrative Network Construction** We define a narrative network as a weighted, directed network where nodes are arguments (**A0** and **A1**), edges represent verb frames, and the weight of edges is the aforementioned log-odd ratio value. To further illustrate the advantages of applying MtChD, we first construct a global network with overall narrative triplets, and compare it to local networks of significant time periods for both Twitter data. This way, we can better map out how narratives shift across time in the local networks, which is unavailable in the global network.

On top of networks, we apply the multiscale backbone filtering (Serrano, Boguná, and Vespignani 2009) to extract the

<sup>7</sup><https://www.sbert.net/index.html>

<sup>8</sup><https://github.com/facebookresearch/SentEval>

core narrative network. To focus on the most important narratives, we set  $\alpha$  to  $10^{-7}$  for the global network and 0.01 for local networks for the COVID data; for the French Election data, we set  $\alpha$  to 0.001 for Le Pen related local networks, and 0.01 for Macron related local networks.

## Evaluation

We quantitatively evaluate our framework on the two parts: MtChD and SRL. First, we use synthetic data to evaluate MtChD, and then use both synthetic data and the two Twitter corpora to evaluate SRL and clustering performance.

**Evaluation with Synthetic Data** The key idea in this evaluation is to construct a dataset with planted narratives that change over time by using paraphrasing capacity of a generative language model. With these planted narratives and change points that we can manipulate, we test the performance of our method in recovering the change points and core narratives from the synthetic data.

**Synthetic Data Generation** Our synthetic data construction starts with an empirical event data. We consider the 22 major events related to COVID-19 provided by The American Journal of Managed Care (AJMC)<sup>9</sup> as the *reference events*, and paraphrase their summary descriptions into ten additional versions using ChatGPT, resulting in a total of 197 reference narratives. Additionally, we sample *noise events* from Wikipedia. These noise events are major events happened between 2021 and 2022<sup>10</sup>, such as “The 2022 Winter Olympics begin in Beijing, China.” We randomly sample 100 events that are not directly related to COVID-19. Each noise event is represented in the synthetic data as a single narrative without additional paraphrased sentences. With narrative sentences, SRL extracts 104 narrative fragments in total. As we will explain below, using these basic narrative elements, we generate different synthetic dataset for each task.

**Evaluation of Change Point Detection** We use synthetic data to evaluate the robustness of MtChD against noise and temporal overlap by comparing its performance with a widely-used kernel based change point detection method (Truong, Oudre, and Vayatis 2020).

To test the robustness of MtChD against noise, we generate synthetic data to mimic a scenario where there is one change point splitting the overall timeline into two, each of which consists of five days. We assume that the distribution of events remains the same within each interval. Three events before March 17 from the reference events are sampled for the first interval, and another three events that occurred after March 17 are sampled for the second interval. Additionally, we add noise events randomly selected from the noise pool at a specified ratio. We use the same noise pool for both intervals to ensure consistency.

To test the robustness of MtChD against event overlaps, we adopt a similar process, but with increasing amount of




Figure 1: UMAP of ground truth clusters and noise events, where ground truth narrative have clear clusters and noise events scatter in the space.

overlap between periods. we randomly select one event from the reference that is not one of the six events selected before as the overlap event. We specify the overlap ratio  $a_1$  for interval one, the overlap ratio  $a_2$  for interval two, and the number of days where the overlap events occur. We randomly select narratives from the pool that represent the overlap event, and add them to the data.

**Evaluation on SRL Clustering** We adopt the idea of *coherence* from dynamic topic modeling to evaluate the quality of predicted clusters, which is the product of BIRCH on narrative fragments. In essence, the idea is about quantifying how coherent or homogeneous each topical cluster is. As a metric, it is common to use Umass (Mimno et al. 2011). However, Umass focuses on the extent of *word* overlap within predicted clusters, which does not necessarily capture semantic similarity. Therefore we define an embedding-based coherence measure, which use cosine similarity as a proxy of semantic similarity. We train all narrative fragments with a pre-trained SentBert model (“all-mpnet-base-v2”). And with the embeddings, we assign the average value of cosine similarity of all pair-wise narrative fragments within a predicted cluster. The measurement is applied to both Twitter corpora.

We also evaluate the robustness of BIRCH against noise on synthetic data. All synthetic data generated for this task consist of 197 ground truth narratives from major events related to COVID, and certain amounts of randomly generated noise events (size ranges from 0 to 104). We also test the impact of BIRCH threshold on the performances. Specifically, we experiment with six different threshold values: 1, 1.5, 2, 2.5, 3, and 3.5. Fig 1 shows the clear cluster structures of ground truth narrative, while noise events scatter in the space.

Our evaluation uses two types of criteria: relative and absolute. The **relative criterion** focuses on the pair-wise semantic similarity within a group. Narratives are considered accurately grouped if they belong to the same event and share a common cluster. By contrast, the **absolute criterion** measures the ability of the clustering method to group similar narrative fragments together. We check whether the cluster with the most narratives belonging to a certain event

<sup>9</sup><https://www.ajmc.com/view/a-timeline-of-covid19-developments-in-2020>

<sup>10</sup><https://en.wikipedia.org/wiki/2021>; <https://en.wikipedia.org/wiki/2022>






Figure 2: The global narrative network of the COVID data.

matches the corresponding ground truth cluster.

## Results

We demonstrate how our pipeline<sup>11</sup> works with two Twitter corpora (COVID-19 and French Election), and its robustness against noise and event overlaps with synthetic data.

### COVID Data

**Change Points** Figure 3 shows the significant change points discovered, along with the timeline and narrative networks. The change point 02/28 is the most significant and it is close to the date of the very first US COVID-19 death. The next most significant change points are 02/11 and 03/04. The second change point coincides with California’s declaring a state of emergency (03/04/20). Given that exploring details of the COVID data across all levels is beyond the scope of our study, here we only show results that are about micro-level narrative discovery, partitioning the data on significant change points: 01/30/20, 02/04/20, 02/11/20, 02/16/20, 02/21/20, 02/28/20, 03/04/20, 03/09/20, 03/15/20, and 03/24/20, resulting in 11 time frames. (see Fig. 3). There are 35,650, 34,795, 25,968, 21,206, 21,086, 135,968, 42,872, 28,423, 27,819, 38,296, 39,059 unique narratives observed in the 11 time frames respectively. Please see the COVID event tree<sup>12</sup> for more detail.

**Top Narrative Fragments** The top 15 representative narrative fragments are selected for the time frames (top three are listed in Figure. 3). We find that “anti-chinese jokes,” “China,” “Chinese,” “Wuhan,” “Wuhan virus,” and “Huanan seafood market” were frequently mentioned from 01/21/2020 to 02/20/2020. The related narrative fragments change from “coronavirus outbreak” to “Wuhan doctor + PERFORM + alarm coronavirus dies”, and later about “bioweapons research from seafood market/super lab in

<sup>11</sup>check detailed tables and interactive figures at [https://osf.io/3aguz/?view\\_only=704d701a57224a3989e79b4b7726ceba](https://osf.io/3aguz/?view_only=704d701a57224a3989e79b4b7726ceba)

<sup>12</sup>[https://osf.io/3aguz/?view\\_only=704d701a57224a3989e79b4b7726ceba](https://osf.io/3aguz/?view_only=704d701a57224a3989e79b4b7726ceba)

Wuhan” and “concerns outbreak’s economic impact.” Besides China, other countries such as Israel and South Korea are also the top mentioned entities between 02/16/2020 to 02/27/2020. The United States is first frequently mentioned in 02/11/2020–02/15/2020 about “horror coronavirus + REACH + san antonio” and later about “administration + AUTHORIZE\_ADMIT + infected americans” and “60 cases + CLOSE + us” in 02/28/2020–03/03/2020. Interestingly, after 03/03/2020, narrative fragments related to countries lose their significance, and instead, networks are overtaken by narrative fragments starting with pronouns such as “I,” “You,” and “He.” This may be due to the surge of cases across the glob and signifies the pandemic phase, given that WHO declared COVID-19 a pandemic on 03/11/2020.

**Global Narrative Network** We apply the backbone method on the global narrative network, extracting 924 nodes and 797 edges. We highlight the hubs (with degree at least 10) in red, and highlight edges which represent top 15 overall narratives (see previous section) in orange. As shown in Fig. 2, more than half of the hubs are pronouns, including “They”, “He”, and “We”. Other hubs are more content-specific, including “china”, “people”, and “coronavirus”. Among these hubs, most of the top narratives are related to “People”, “I” and “We”. In fact, there are 52 out of 92 top narrative fragments connected to hubs. The rest of the top narrative fragments tend to be isolated such that their end nodes have a degree of one. Examples include “white people + MOUNT\_ASSEMBLE\_PRODUCE + anti-chinese jokes”, “south korea + INFORM + 160 new cases”, and “grocery stores + BEGIN + clearing out”.

**Local Narrative Networks** According to Fig. 3, starting from 01/21 to 02/27, “China”, “Coronavirus”, and pronoun “I” are the main hubs in local narrative networks. Later, although “Coronavirus” and “I” remain as hubs, “China” gradually loses its significance. Starting from 03/04, pronouns such as “We”, “They”, and “He” along with “I” start to appear as new hubs. After 03/15 they even have higher significance in the local networks compared to the “Coronavirus”. Furthermore, the overlap between local networks and the top 15 overall narratives suggest that on average, eight top narratives are captured by local networks.

### French Election Data

**Change Points** The most significant change point in the French Election data corresponds to the second round of voting (May 7th, 2017), which pitted Emmanuel Macron against Marine Le Pen in the runoff elections. The next-level change points correspond to the start of the #MacronLeaks campaign (05/03) and Macron’s inauguration, which took place on 05/14.

Similarly, the results generated by MtChD includes change points with multiple granularity. The following results are based on the most fine-grained change points, which include 05/03, 05/05, 05/07, 05/08, 05/13, and 05/19, resulting in eight time frames (See Fig. 4). There are 153,339, 68,297, 118,973, 65,281, 79,122, 126,446, 83,031, 63,617 unique narratives found in the eight time frames respectively.

# Narrative networks

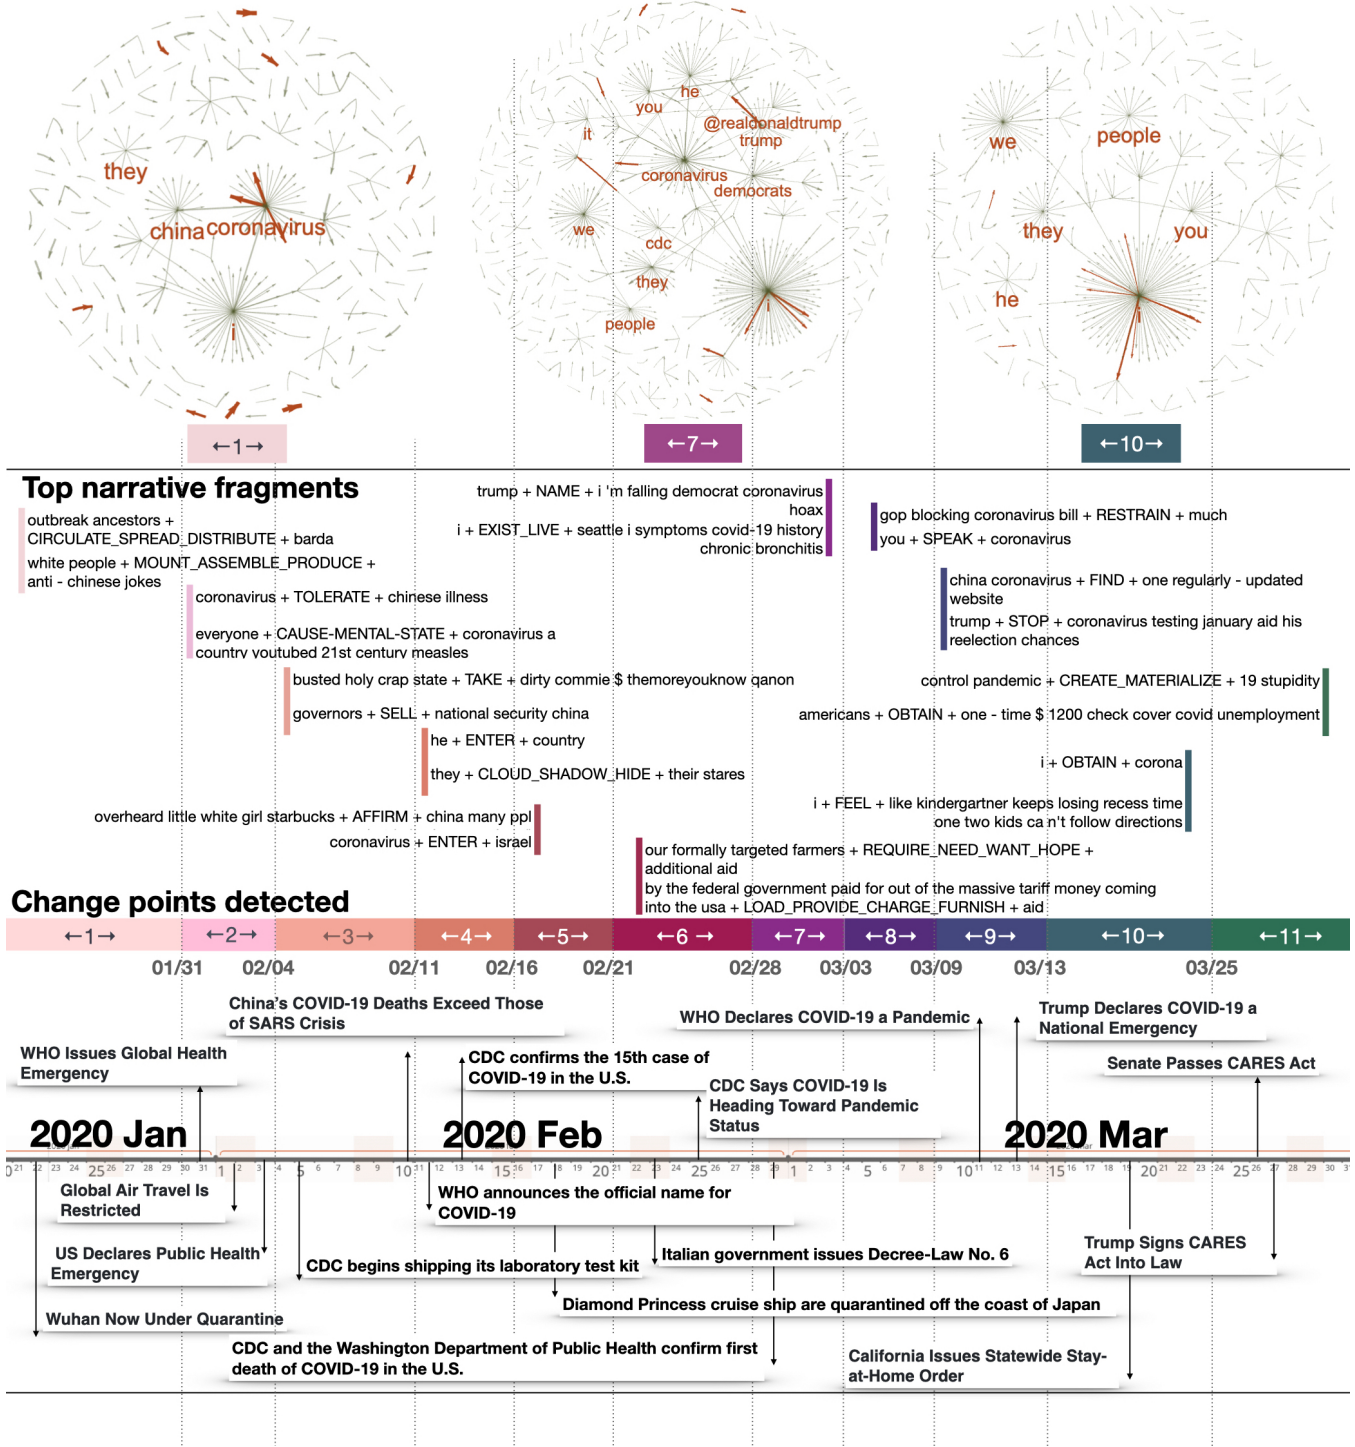


Figure 3: Local narrative networks extracted from the COVID-19 corpus. We set nodes with degree at least 10 in red, and node with degree between five and 10 in black. Orange edges represent the top narrative fragments in the backbone. The bottom subplot shows the timeline related to COVID provided by CDC and AJMC. The other subplots show results of change point detection, top three overall narrative fragments, and selected narrative networks respectively.

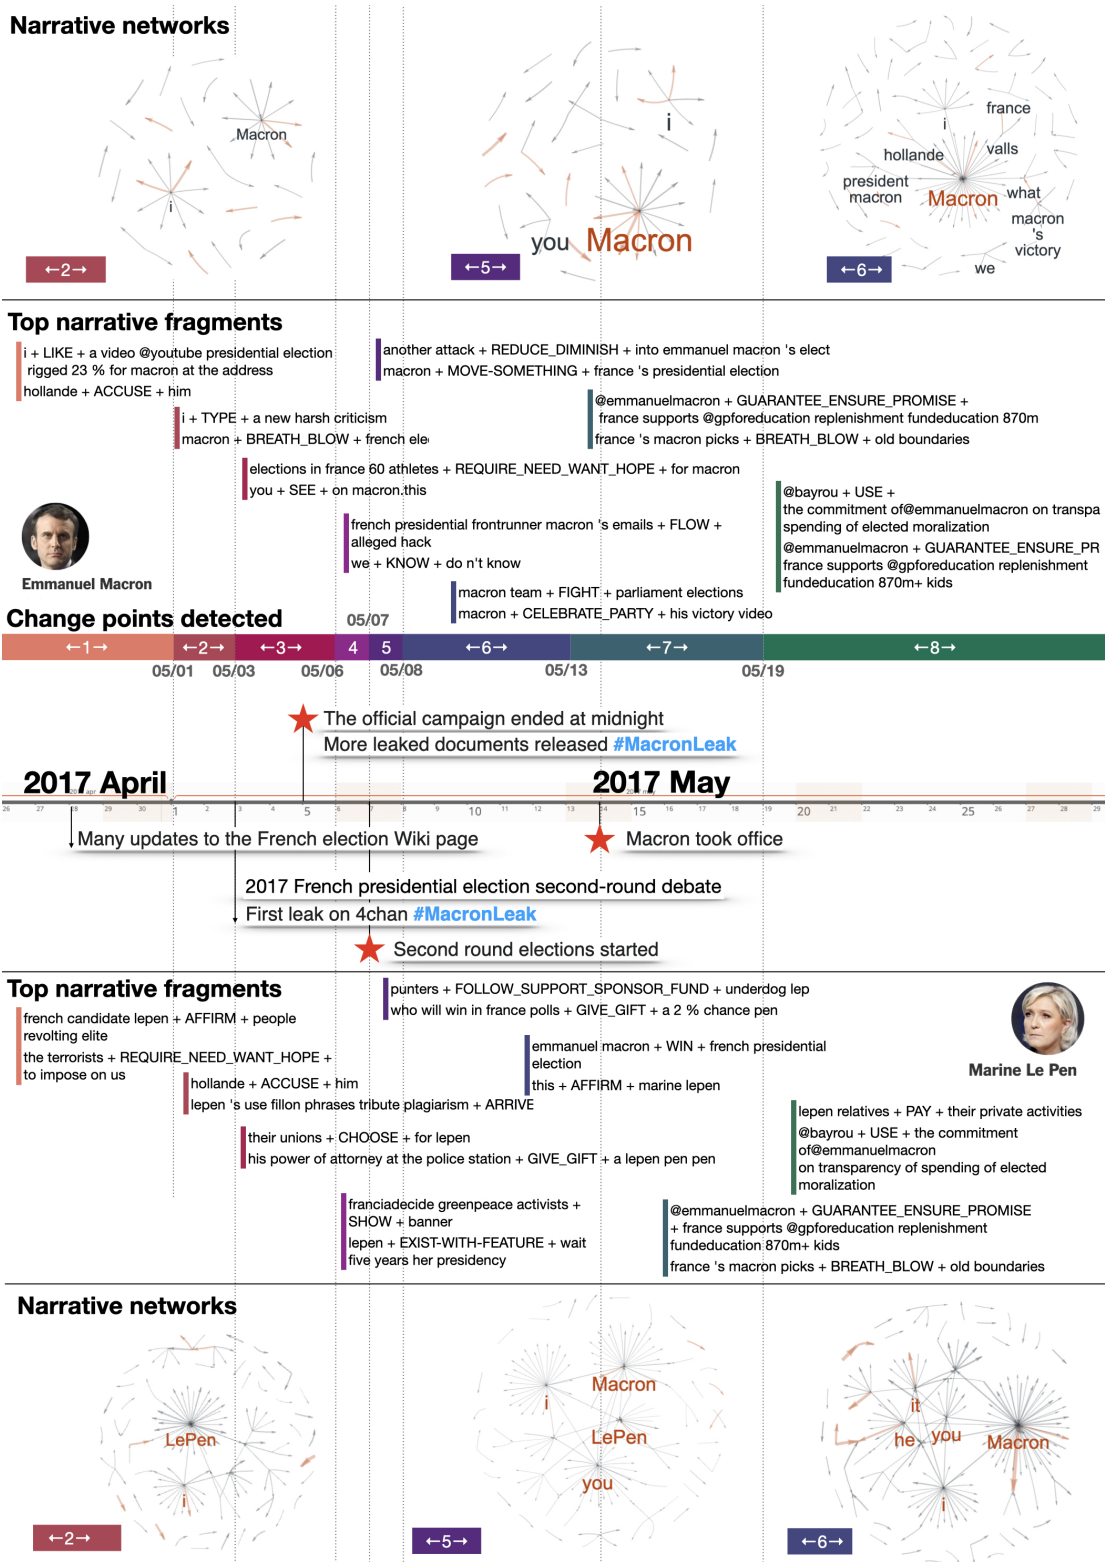


Figure 4: Local narrative networks with the 2017 French election data. Node and edge colors represents the same characteristics as in the previous figure. The bottom subplot shows the timeline from Wikipedia. The other subplots show results of change point detection, the top three narrative fragments, and selected narrative networks respectively.






Figure 5: Robustness of Meta change point detection against noise. The gap refers to the distance between the ground truth change point and the predicted one. The plot shows the predictions by MtChD are pretty stable until noise ratio reaches 0.8; while the performances of the Kernel-based method drastically drop when noise ratio is larger than 0.6.

**Top Overall Narrative Fragments** From the top 15 narrative fragments related to Le Pen, we see those related to her political ideology. For example, “people revolting elite” is the most representative narrative fragment from 04/26 to 04/30, “euro deadweight capital controls option” is repeatedly shown as the top narrative fragments in 05/01–05/02, and in 05/03–05/04. During the same time, narrative fragments with emotions keep appearing, such as “russia sanctions” and “holocaust denial” found in 04/26–04/30, “use filon phrases tribute plagiarism” in 05/01–05/02, and “french police + RESIST + their unions vote for lepen” in 05/03–05/04. Later starting from 05/05, rumors about “invalid damaged lepen ballots” start to spread. And the top narrative fragments after 05/08 shift the attention to the winning of Emmanuel Macron although the original tweets include keywords related to Le Pen.

On the other hand, the top narrative fragments about Macron are mostly positive. For example, the endorsement from Laurence Parisot, Melenchon, and athletes were the main narrative fragments in 04/26–04/30, 05/01–05/02, and 05/03–05/04 respectively. In 05/05–05/06, seven narrative fragments concern with “Macron Leak”, an anti-Macron influence campaign (Vilmer 2021). Although they still remain heated in the next interval, narrative fragments shift to Macron winning the election. And starting from 05/08, the major narrative fragments are about “macron team + Fight + parliament elections” (No.1 in 05/08–05/12), Macron’s political moves such as support “gpforeducation” (No.1 from 05/13–05/18), and “transparency of spending of elected moralization” (No.1 from 05/19–05/29). These narrative fragments reflect important priorities of Macron’s administration, after he was sworn into office (05/14).

**Core Narrative Networks** The local networks related to Le Pen and Macron exhibit big differences. In local networks related to Le Pen, node “Macron” gains more and more attention starting at 05/01, and after 05/07 it attracts more attention compared to node “LePen” and even takes over the core position starting 05/08. However, we don’t see similar phenomenon in local networks related to Macron.

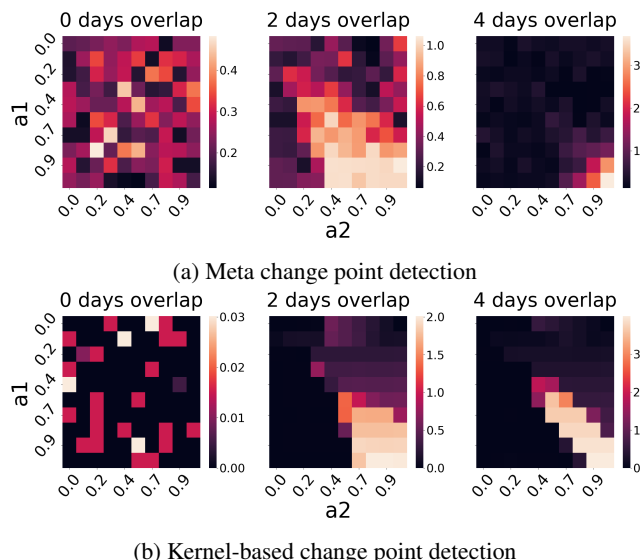


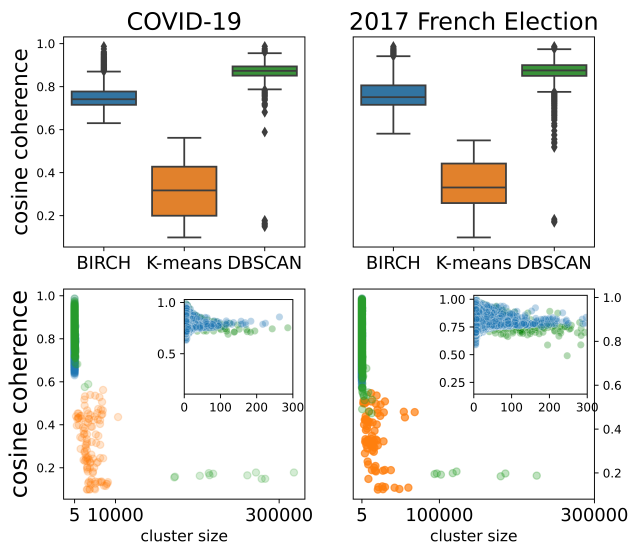
Figure 6: Robustness of change point detection methods against event overlap. We generate synthetic data with one change point data. The x axis represents overlap ratio  $a_1$  for interval one, and y axis is about overlap ratio  $a_2$  for interval two. The value within a cell is the gap between ground truth change point and the predicted change point given a pair of  $a_1$  and  $a_2$ . We show cases of 1) no overlap 2) 2-days overlap and, 3) 4-days overlap for both MtChD and the Kernel-based method. The performances of MtChD and Kernel-based method show association with  $a_1$  and  $a_2$ , with the Kernel-based method showing higher sensitivity to noise.

Node “Macron” always maintains the core position and node “LePen” is never the hubs. In term of the overlap with top narratives, on average nine out the 15 top narratives are captured by local networks related to Le Pen, and on average 11 out of 15 top narratives are captured by local networks related to Macron. The comparison helps us understand the difference between statistical measurement and network-based method in uncovering significant narrative fragments.

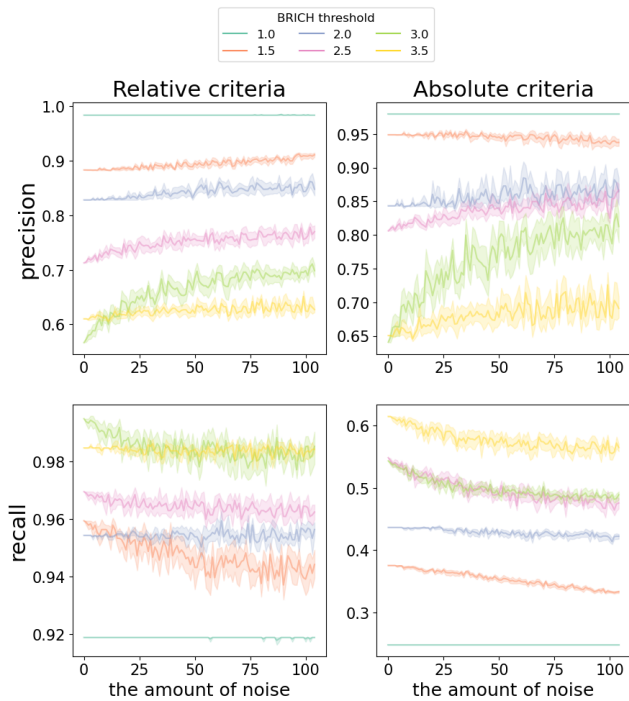
## Evaluation

**Robustness of Meta Change Point Detection** In both experiments, performance is evaluated by the distance between true change point, on day five, and the predicted one. Fig. 5 shows the gap between ground truth and the predicted change point. Meta change point detection (MtChD) remains accurate until noise ratio reaches 0.8, while the Kernel-based method (Truong, Oudre, and Vayatis 2020) shows worse performance after noise ratio of 0.6, showing weaker robustness against random noise.

In addition, we observe a decrease in accuracy with an increasing number of overlapping days in predictions generated by both methods (Fig. 6). Furthermore, as variables  $a_1$  and  $a_2$  increase, we note a widening gap between ground truth and predicted values in the results of MtChD. In extreme cases where  $a_1$  and  $a_2$  equal one, this gap indicates



(a) Empirical



(b) Synthetic

Figure 7: Evaluation on SRL clustering quality. The two subplots on the left in (a) show the cosine coherence value indicating how BIRCH, DBSCAN, and K-means perform in clustering the two Twitter corpora—the COVID-19 and the 2017 French Election. The other two subplots on the right show the association between cosine coherence and the size of a cluster. Subplots in (b) show the precision and recall of BIRCH model in clustering synthetic data under relative and absolute criteria. In addition, we show the changes in performance as the amount of noise rises.

a high reliability, predicting the date earlier than the ground truth by the number of overlapping days. This demonstrates the ability of both MtChD and the Kernel-based method to effectively capture systematic changes in the data. In less extreme cases, the Kernel-based method displays higher sensitivity to noise compared to MtChD. Specifically, in instances of 2-day and 4-day overlaps, predictions by MtChD closely align with ground truth compared to those of the Kernel-based method, given the same pair of  $a_1$  and  $a_2$ . Once again, our method exhibits greater robustness to event overlaps when compared to the Kernel-based method.

**Evaluation of Narrative Clusters** With the evaluation on the Twitter data, we find that both DBSCAN and BIRCH could generate clusters with high average coherence (around 0.8) compared to K-means, but the results of DBSCAN have more outliers with very low coherence (Fig. 7 (a)). By the nature of the coherence measure, it may be reasonable to expect that the larger clusters would have lower coherence. This can have an outsized impact on evaluation results if one method produces larger clusters than the others. To analyze this possibility, we check the relationship between cluster size and coherence. The result shows that all clusters generated by BIRCH have coherence value no less than 0.6, and their sizes are less than 10,000. Although most of the clusters predicted by DBSCAN have coherence value above 0.6 and decent size, yet a couple of the clusters are huge and exhibit low coherence. It suggests huge clusters with extremely low coherence generated by DBSCAN would introduce non-trivial bias for the downstream task—calculating the significance of the aggregated narrative fragments. In sum, these results suggest BIRCH produces more reliable clustering.

In addition, our evaluations (Fig. 7 (b)) on synthetic data demonstrate that precision of BIRCH are all above or at least around 0.6 under both relative and absolute criteria, and the values further increase as more and more noise events are added. Recalls under relative criterion are all above 0.92, and they have slight decline as the number of noise rises. Although recalls under absolute criterion are not as high as that in relative criterion, they remain stable against the increase of noise. In sum, our method is robust against noise.

## Discussion

Here, we propose a framework to automatically extract narratives from a stream of tweets. By combining change point detection with SRL, we investigate the collective narrative shifts. We also evaluate and explore the performance of our framework.

Core narratives captured by both methods align with the timeline provided by AJMC<sup>13</sup> and CDC<sup>14</sup> about COVID, and that provided about French Election 2017<sup>15</sup>, suggesting the substantial feasibility of our method.

<sup>13</sup><https://www.ajmc.com/view/a-timeline-of-covid19-developments-in-2020>

<sup>14</sup><https://www.cdc.gov/museum/timeline/covid19.html>

<sup>15</sup>[https://en.wikipedia.org/wiki/2017\\_French\\_presidential\\_election](https://en.wikipedia.org/wiki/2017_French_presidential_election)

**Temporality Is Essential** Comparing the global narrative network with local narrative networks yield some notable insights. First, most of the hubs are highlighted in at least one of the local networks. However, without the clear time boundary pinpointed by change point detection, the persistence of the hubs is unclear. For example, the entity “China” is the core node before 03/03, but gradually loses attention, as the COVID-19 moves from the local epidemic in China to the pandemic. And although “He”, “You”, and “They” are shown as the hubs in the global network, they did not become hubs until 03/04. Without change point detection, it would be challenging to uncover such narrative shifts.

**Ambient storytelling** One hub that catches our attention concerns the entity “Trump”. Despite its seldom appearance as a hub in local networks, it attracts plenty of attention between 02/28 and 03/14. This can be considered as evidence of *ambient storytelling* (Page, Harper, and Frobenius 2013; Sadler 2021)—collective narrative fragments with no strong sequel or causal link yet building up a complete narrative. Taking a closer look at narrative fragments connect to “Trump” in the global network, we find “trump + STOP + coronavirus testing January aid his reelection chances”, “trump + AFFIRM + coronavirus outbreak all control very small problem”, “trump + DISMISS\_FIRE-SMN + government’s entire pandemic response chain command including white house management infrastructure”, and “democrats + ACCUSE + trump” forming a narrative about Trump’s role in controlling COVID-19 pandemic.

In addition to narrative fragments related to the core entities, we also find another type of ambient storytelling, which has direct links between narrative fragments. For example, the combination of “administration + FOL-LOW\_SUPPORT\_SPONSOR\_FUND + people” and “people + REQUIRE\_NEED\_WANT\_HOPE + stay home” in local network 03/15–03/23 is a clear narrative about officials ordering lockdowns in their jurisdictions.

**Horizontal Storytelling** Other than ambient storytelling, we also find evidence of *horizontal storytelling* (Sadler 2021; Bal and Van Boheemen 2009), where one piece of narrative fragment could imply the whole story. For example, in the local network that depicts narratives related to Le Pen in 05/03–05/04, “students + SPEAK + vote against lepen” is a standalone narrative. Similarly, in the 05/05–05/06 local network about Macron, the narrative fragment “hackers + HIT + macron” is a complete summary of the #MacronLeaks incident (Vilmer 2021).

Among all horizontal storytelling, adding to the previous research on Trump’s tweets (Sadler 2021), we find evidence of narrative fragments using technique of *sideshadowing* (Morson 1994), which attracts viewers’ attention by emphasizing what *could have* happened. For instance, “diy tricks french voters + RECEIVE + invalid damaged lepen ballots” suggests an anti-Macron conspiracy, misleading Le Pen’s supporters to believe that election fraud prevented their candidate from winning. Likewise, “russia sanctions + PRECLUDE\_FORBID\_EXPEL + lepen election win.” is a sensational accusation without a proof, leaving room for imagination that Le Pen were suppressed by foreign behind-

the-scene forces.

**Narrative Network or Narrative Fragments?** One thing that captures our interest is a strong overlap between the top overall narratives and local networks. The average overlap rate is 10/15 for the COVID, and 9/15 and 11/15 for Le Pen- and Macron-subsets respectively. Among the overlapping narratives, few are connected to hubs. In some extreme cases, for example, the largest hub, “Macron” in Le Pen local network 05/07, is not involved in any of the top 15 narratives in the same time frame. This suggests that hubs might have a larger probability of being involved in important narratives than non-hub nodes, yet, whether a narrative could be representative is not determined by its connection to a hub in a narrative network.

At this point, however, we could not draw a definite conclusion that the network-based method outperforms the statistical method in narrative discovery due to the absence of the benchmark dataset about narratives. Nevertheless, we could use audiences’ adoption of narratives as a proxy for evaluation purposes. A previous study has shown that the “persistence” of a tweet, which is quantified as the repeated exposure under a specific topic, has a significant marginal effect on audiences’ adoption (Romero, Meeder, and Kleinberg 2011). Building on this idea, we could consider the frequency of a narrative fragment as the measure of its importance, and further refine it with measures like log-odd ratio with informative Dirichlet priors (Monroe, Colaresi, and Quinn 2008). This way, the top representative narrative fragments would be the potential reference data for narrative networks comparisons.

If so, the network-based method would suffer from limitations of capturing essential narrative fragments, resulting in narrative mismatches. With our data, all mismatched top narratives connect to at least one long entity. For example, we find “nigel Farage formally endorses + EXIST-WITH-FEATURE + a lot say about lepen european union” and “we academics and researchers + SHOW + our support for Emmanuel Macron” in the French Election data, and “billionaire Jack Ma + CREATE\_MATERIALIZE + coronavirus vaccine roughly equivalent average u.s . family” in the COVID data. Such long entities might have a low value of degree and are therefore hard to satisfy the filtering criteria based on backbone disparity.

Since the essential information of a narrative is stored in the edges in a narrative network, the edges carry information about how representative a narrative fragment is and what the relationship between the end nodes is. Therefore, by applying community detection, we overlook the exact meaning of the relationship attached to the edges and considering all edges equal. Especially in the case where competing narratives co-exist in a narrative network, such a move is not trivial. For example, when applying random walk (e.g., PageRank Brin and Page 1998) to capture importance of nodes, it is not reasonable to assign equal probabilities to edges “I + AGREE + vaccine” and “My baby + NON-TOLERATE + vaccine” simply based on the fact that they share both edge weights and next pathway to “vaccine + KILL + children”, because individuals who endorse the narrative “I +

AGREE + vaccine” would hardly consider “vaccine + KILL + children” acceptable, not to mention collectively forming a narrative linking the two fragments. Similarly, it is rarely likely that a collection of narratives with contrasting beliefs would be considered as two communities, instead of one, when applying community detection. That said, it should be noted that we do not intend to claim that applying node-based network methods is incorrect. Instead, we encourage the researchers to try these methods when contrasting narratives is not the main focus of the study (e.g., Shahsavari et al. 2020).

## Comparison to Dynamic Topic Modeling

The main difference between our model and dynamic topic models (DTM), is the way we represent narratives as a collection of narrative fragments. These fragments are sentence-like—in contrast to the set of words produced by topic models—and allow more nuanced interpretation, and the relationship between keywords that could be detected in DTM would remain “intact” from our (Argument 0, Verb, Argument 1) expression. Therefore, in the case where “Trump”, “COVID”, “vaccine”, and “CDC” are shown in the same topic detected by DTM, our method could provide clear expressions such as “Trump + CRITICIZE + COVID vaccine” and “CDC + RECOMMEND + COVID vaccine”, which would not be possible to directly extract from topic modeling.

Another difference between our method and DTM lies in the focus of research questions. For DTM, the focus is more about the content of the topics, while in our case, we assume that the shift of collective attention of the audience would be influenced by the major events that happened recently. Therefore, we do not put focus on an exhaustive search of the topics being discussed, but rather on what the main focus is after each collective attention shift.

## Limitations

Our method has a limitation in that it may fail to capture narratives that rely heavily on elements labeled with such as, ARGUMENT-LOC, as we only consider A0, V, and A1 elements from the results generated by SRL. Therefore, there may be certain details and nuances that our method cannot capture. For example, given a tweet “RIP Dr Li Wenliang who first sounded the alarm in Wuhan—but muzzled by his hospital leaders when he tried to warn.”, from which the narrative fragments extracted would not include detailed information about “in Wuhan” (ARGUMENT-LOC), “when he tried to warn” (ARGUMENT-TMP). However, we assume that if a narrative has been discussed multiple times, the essential information would potentially be captured as A0 or A1 in some form, and if the signal of such a narrative reaches a significant threshold, it would still appear in the top narratives list generated by our method.

Additionally, although our method aligned well with the major events in the COVID corpus, which provides evidence for its effectiveness in identifying important narratives, we acknowledge that the last change point (05/19/2017) detected in the 2017 French Election data does not seem to be

qualitatively related to any major events. We tried to understand the change point by searching related posts published on 05/19/2017, and only found comments about Macron’s challenge in governing<sup>16</sup> and news about Macron visited French troops<sup>17</sup>, which are not strongly related to the top narrative fragments (e.g., “@bayrou + USE + the commitment of @emmanuelmacron on transparency of spending of elected moralization” and “lepen relatives + PAY + their private activities”) detected for that era with our method. While we lack information on the reasons behind this change point, we cannot rule out the possibility that it still reflects a meaningful shift in public attention, such as the one that occurred after Macron took office.

## Conclusion

In summary, this study contributes the following. First, We propose an automated pipeline that identifies the evolution of collective narratives with statistical and network-based methods. The pipeline has successfully identified the major events and the collective opinions related to major entities. We have quantitatively evaluated the robustness of our method with synthetic data. While working with the pipeline, we find evidence in support of various forms of storytelling emphasized in traditional studies of narratives. We also discuss the difference between statistical measurements and network analysis. We acknowledge the strengths of network-based methods in detecting core elements of narratives, and their capability in retaining major representative narratives even after backbone filtering. Nevertheless, we caution the use of network-based methods when analyzing contrasting narratives. Future efforts could be put to test the general applicability of our approach in data with competing narratives and explore more sophisticated network-based methods to study narrative networks.

## Broader perspective, ethics and competing interests

We analyzed tweets discussing COVID-19 pandemic and French Elections, but we expect our narrative discovery methods to generalize to other datasets. All data is public and collected following Twitter’s terms of service, with the study considered exempt by the authors’ IRB. To minimize risk to users, all identifiable information was removed and analysis was performed on aggregated data. Only public entities and politicians are mentioned explicitly by name in the data. Authors declare no competing interests.

## Acknowledgments

We thank Dr. Jisun An (Indiana University), Dr. Haewon Kwak (Indiana University), and Dr. Lewen Wei (University of New South Wales) for comments on the manuscript. This project was funded in part by DARPA under contract HR001121C0168.

<sup>16</sup><https://www.nytimes.com/2017/05/19/opinion/mr-macrons-next-challenge-governing.html>

<sup>17</sup><https://www.aljazeera.com/news/2017/5/19/emmanuel-macron-visits-french-troops-in-northern-mali>



## References

- Appel, M.; and Richter, T. 2007. Persuasive effects of fictional narratives increase over time. *Media Psychology*, 10(1): 113–134.
- Bal, M.; and Van Boheemen, C. 2009. *Narratology: Introduction to the theory of narrative*. University of Toronto Press.
- Barber, J. 2015. A generalized likelihood ratio test for coherent change detection in polarimetric SAR. *IEEE Geoscience and Remote Sensing Letters*, 12(9): 1873–1877.
- Bearman, P. S.; and Stovel, K. 2000. Becoming a Nazi: A model for narrative networks. *Poetics*, 27(2-3): 69–90.
- Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan): 993–1022.
- Blondel, V. D.; Guillaume, J.-L.; Lambiotte, R.; and Lefebvre, E. 2008. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10): P10008.
- Brin, S.; and Page, L. 1998. The anatomy of a large-scale hypertextual web search engine. *Computer networks and ISDN systems*, 30(1-7): 107–117.
- Bruner, J. 1987. Life as narrative. *Social research*, 11–32.
- Chen, E.; Lerman, K.; Ferrara, E.; et al. 2020. Tracking social media discourse about the covid-19 pandemic: Development of a public coronavirus twitter data set. *JMIR public health and surveillance*, 6(2): e19273.
- Church, K. W. 2017. Word2Vec. *Natural Language Engineering*, 23(1): 155–162.
- Collobert, R.; Weston, J.; Bottou, L.; Karlen, M.; Kavukcuoglu, K.; and Kuksa, P. 2011. Natural language processing (almost) from scratch. *Journal of machine learning research*, 12: 2493–2537.
- De Graaf, A.; Hoeken, H.; Sanders, J.; and Beentjes, J. W. 2012. Identification as a mechanism of narrative persuasion. *Communication research*, 39(6): 802–823.
- Dourish, P.; and Gómez Cruz, E. 2018. Datafication and data fiction: Narrating data and narrating with data. *Big Data & Society*, 5(2).
- Ferrara, E. 2017. Disinformation and social bot operations in the run up to the 2017 French presidential election. *First Monday*, 22(8).
- Fillmore, C. J. 1967. The case for case. *Universals in Linguistic Theory*.
- Gardner, M.; Grus, J.; Neumann, M.; Tafjord, O.; Dasigi, P.; Liu, N. F.; Peters, M. E.; Schmitz, M.; and Zettlemoyer, L. 2018. AllenNLP: A Deep Semantic Natural Language Processing Platform. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, 1–6.
- Gottschall, J. 2012. *The storytelling animal: How stories make us human*. Houghton Mifflin Harcourt.
- Grootendorst, M. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794*.
- He, Y.; Burghardt, K. A.; and Lerman, K. 2022. Leveraging change point detection to discover natural experiments in data. *EPJ Data Science*, 11(1): 49.
- Hermida, A. 2014. Twitter as an ambient news network. *Twitter and society*, 89: 359–372.
- Jing, E.; and Ahn, Y.-Y. 2021. Characterizing partisan political narrative frameworks about COVID-19 on Twitter. *EPJ data science*, 10(1): 53.
- Kingsbury, P.; and Palmer, M. 2003. Propbank: the next level of treebank. In *Proceedings of Treebanks and lexical Theories*, volume 3. Citeseer.
- Kopfman, J. E.; Smith, S. W.; Ah Yun, J. K.; and Hodges, A. 1998. Affective and cognitive reactions to narrative versus statistical evidence organ donation messages.
- Lazard, A. J.; Scheinfeld, E.; Bernhardt, J. M.; Wilcox, G. B.; and Suran, M. 2015. Detecting themes of public concern: a text mining analysis of the Centers for Disease Control and Prevention’s Ebola live Twitter chat. *American journal of infection control*, 43(10): 1109–1111.
- Lee, D. D.; and Seung, H. S. 2001. Algorithms for non-negative matrix factorization. In *14th Annual Neural Information Processing Systems Conference, NIPS 2000*. Neural information processing systems foundation.
- Loper, E.; and Bird, S. 2002. NLTK: The Natural Language Toolkit. In *Proceedings of the ACL-02 Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics*, 63–70.
- Meraz, S.; and Papacharissi, Z. 2013. Networked gatekeeping and networked framing on# Egypt. *The international journal of press/politics*, 18(2): 138–166.
- Mimno, D.; Wallach, H.; Talley, E.; Leenders, M.; and McCallum, A. 2011. Optimizing semantic coherence in topic models. In *Proceedings of the 2011 conference on empirical methods in natural language processing*, 262–272.
- Monroe, B. L.; Colaresi, M. P.; and Quinn, K. M. 2008. Fightin’ words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, 16(4): 372–403.
- Moretti, F. 2000. Conjectures on world literature. *New left review*, 1: 54.
- Morson, G. S. 1994. *Narrative and freedom: The shadows of time*. Yale University Press.
- Niekum, S.; Osentoski, S.; Atkeson, C. G.; and Barto, A. G. 2015. Online bayesian changepoint detection for articulated motion models. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 1468–1475. IEEE.
- Page, E. S. 1954. Continuous inspection schemes. *Biometrika*, 41(1/2): 100–115.
- Page, R.; Harper, R.; and Frobenius, M. 2013. From small stories to networked narrative: The evolution of personal narratives in Facebook status updates. *Narrative Inquiry*, 23(1): 192–213.
- Papacharissi, Z. 2016. Affective publics and structures of storytelling: Sentiment, events and mediality. *Information, communication & society*, 19(3): 307–324.

- Raghavan, V.; Galstyan, A.; and Tartakovsky, A. G. 2013. Hidden Markov models for the activity profile of terrorist groups. *The Annals of Applied Statistics*, 2402–2430.
- Richardson, B. 1997. *Unlikely stories: Causality and the nature of modern narrative*. University of Delaware Press.
- Roberts, G. 2001. *The history and narrative reader*. Psychology Press.
- Romero, D. M.; Meeder, B.; and Kleinberg, J. 2011. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proceedings of the 20th international conference on World wide web*, 695–704.
- Sadler, N. 2018. Narrative and interpretation on Twitter: Reading tweets by telling stories. *New Media & Society*, 20(9): 3266–3282.
- Sadler, N. 2021. *Fragmented narrative: Telling and interpreting stories in the Twitter Age*. Routledge.
- Serrano, M. Á.; Boguná, M.; and Vespignani, A. 2009. Extracting the multiscale backbone of complex weighted networks. *Proceedings of the national academy of sciences*, 106(16): 6483–6488.
- Shahsavari, S.; Holur, P.; Wang, T.; Tangherlini, T. R.; and Roychowdhury, V. 2020. Conspiracy in the time of corona: Automatic detection of emerging COVID-19 conspiracy theories in social media and the news. *Journal of computational social science*, 3(2): 279–317.
- Shi, P.; and Lin, J. 2019. Simple bert models for relation extraction and semantic role labeling. *arXiv preprint arXiv:1904.05255*.
- Tangherlini, T. R.; Shahsavari, S.; Shahbazi, B.; Ebrahimzadeh, E.; and Roychowdhury, V. 2020. An automated pipeline for the discovery of conspiracy and conspiracy theory narrative frameworks: Bridgegate, Pizzagate and storytelling on the web. *PLOS ONE*, 15(6).
- Truong, C.; Oudre, L.; and Vayatis, N. 2020. Selective review of offline change point detection methods. *Signal Processing*, 167: 107299.
- Van Nieuwenburg, E. P.; Liu, Y.-H.; and Huber, S. D. 2017. Learning phase transitions by confusion. *Nature Physics*, 13(5): 435–439.
- Vilmer, J.-B. J. 2021. Fighting information manipulation: the french experience. In Jayakumar, S.; Ang, B.; and Anwar, N. D., eds., *Disinformation and Fake News*, 75–89.
- Willsky, A.; and Jones, H. 1976. A generalized likelihood ratio approach to the detection and estimation of jumps in linear systems. *IEEE Transactions on Automatic control*, 21(1): 108–112.
- Wilson, R. C.; Nassar, M. R.; and Gold, J. I. 2010. Bayesian online learning of the hazard rate in change-point problems. *Neural computation*, 22(9): 2452–2476.
- Winterbottom, A.; Bekker, H. L.; Conner, M.; and Mooney, A. 2008. Does narrative information bias individual’s decision making? A systematic review. *Social science & medicine*, 67(12): 2079–2088.
- Yu, D.; Xu, D.; Wang, D.; and Ni, Z. 2019. Hierarchical topic modeling of Twitter data for online analytical processing. *IEEE Access*, 7: 12373–12385.
- Yu, J.; and Qiu, L. 2019. ULW-DMM: An Effective Topic Modeling Method for Microblog Short Text. *IEEE Access*, 7: 884–893.
- Zhang, T.; Ramakrishnan, R.; and Livny, M. 1996. BIRCH: an efficient data clustering method for very large databases. *ACM sigmod record*, 25(2): 103–114.

## Paper Checklist

1. For most authors...
  - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
  - (b) Do your main claims in the abstract and introduction accurately reflect the paper’s contributions and scope? **Yes**
  - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes, please see Methods section**
  - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes, please see Methods and Discussion sections**
  - (e) Did you describe the limitations of your work? **Yes, please see Discussion section**
  - (f) Did you discuss any potential negative societal impacts of your work? **No, because our proposed method is true to the data. It does not provide arbitrary filtering or manipulation of narratives in the data. The result it generates would only reflect the key components of the provided data.**
  - (g) Did you discuss any potential misuse of your work? **Yes, please see Discussion section**
  - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **NA**
  - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**
2. Additionally, if your study involves hypotheses testing...
  - (a) Did you clearly state the assumptions underlying all theoretical results? **NA. Our paper does not involves hypotheses testing.**
  - (b) Have you provided justifications for all theoretical results? **NA**
  - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA**
  - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **NA**

- (e) Did you address potential biases or limitations in your theoretical framework? NA
  - (f) Have you related your theoretical results to the existing literature in social science? NA
  - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? NA
3. Additionally, if you are including theoretical proofs...
- (a) Did you state the full set of assumptions of all theoretical results? NA
  - (b) Did you include complete proofs of all theoretical results? NA
4. Additionally, if you ran machine learning experiments...
- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes, please check Method Section and URLs in footnote. The URLs include link to GitHub repository and supplemental material in OSF.](#)
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? NA, because we don't have training data. The model we proposed is a pipeline combining varied well-established models, such as Semantic Role Labeling provided by AllenNLP.
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? NA, we do not have prediction evaluation. We have evaluation on pipeline with synthetic data and empirical in Method and Results sections
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No, because the model is not computationally expensive, should be able to work on any devices.](#)
  - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? [Yes, please check evaluation on pipeline with synthetic data and empirical in Method and Results sections.](#)
  - (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? [Yes, please check Discussion section for limitation.](#)
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
- (a) If your work uses existing assets, did you cite the creators? [Yes](#)
  - (b) Did you mention the license of the assets? [Yes](#)
  - (c) Did you include any new assets in the supplemental material or as a URL? [Yes, the code for our model](#)
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes](#)
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes, all personally identifiable information were removed before we obtained the data.](#)
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
- (a) Did you include the full text of instructions given to participants and screenshots? NA, we do not have crowdsourcing in this research.
  - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? NA
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? NA
  - (d) Did you discuss how data is stored, shared, and de-identified? NA