

Unsupervised Story Discovery from Continuous News Streams via Scalable Thematic Embedding

Susik Yoon
UIUC
susik@illinois.edu

Dongha Lee
Yonsei University
donalee@yonsei.ac.kr

Yunyi Zhang
UIUC
yzhan238@illinois.edu

Jiawei Han
UIUC
hanj@illinois.edu

ABSTRACT

Unsupervised discovery of stories with correlated news articles in real-time helps people digest massive news streams without expensive human annotations. A common approach of the existing studies for unsupervised online story discovery is to represent news articles with symbolic- or graph-based embedding and incrementally cluster them into stories. Recent large language models are expected to improve the embedding further, but a straightforward adoption of the models by indiscriminately encoding all information in articles is ineffective to deal with text-rich and evolving news streams. In this work, we propose a novel *thematic embedding* with an off-the-shelf pretrained sentence encoder to dynamically represent articles and stories by considering their shared temporal themes. To realize the idea for unsupervised online story discovery, a scalable framework USTORY is introduced with two main techniques, *theme- and time-aware dynamic embedding* and *novelty-aware adaptive clustering*, fueled by lightweight story summaries. A thorough evaluation with real news data sets demonstrates that USTORY achieves higher story discovery performances than baselines while being robust and scalable to various streaming settings.

CCS CONCEPTS

• **Information systems** → **Data stream mining**; **Web searching and information discovery**; **Document representation**.

KEYWORDS

News Stream Mining, News Story Discovery, Document Embedding

ACM Reference Format:

Susik Yoon, Dongha Lee, Yunyi Zhang, and Jiawei Han. 2023. Unsupervised Story Discovery from Continuous News Streams via Scalable Thematic Embedding. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*, July 23–27, 2023, Taipei, Taiwan. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3539618.3591782>

1 INTRODUCTION

Background. Given a huge number of news articles (or simply *articles*) being generated through digital mass media in real-time, a set of correlated articles naturally form a *story* as they are consecutively

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGIR '23, July 23–27, 2023, Taipei, Taiwan

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9408-6/23/07...\$15.00
<https://doi.org/10.1145/3539618.3591782>

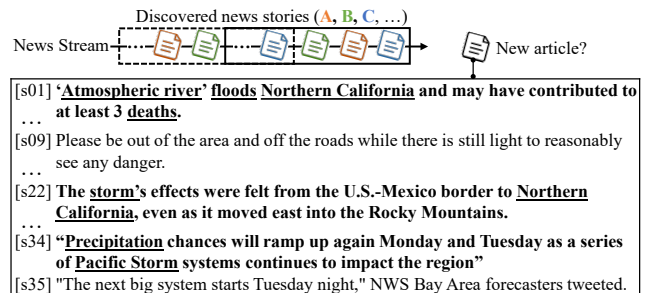


Figure 1: An article [29] about 2023 California Floods where story-indicative sentences and keywords are highlighted.

published while describing relevant events under a unique theme. *Unsupervised online story discovery* helps people digest massive news streams without human annotations by supporting real-time news curation services and downstream tasks such as recommendation, summarization, and event detection [32, 47, 50, 52].

To deal with *text-rich* and *evolving* properties of news streams, existing work for unsupervised online story discovery utilized keywords statistics to represent articles and incrementally cluster them into stories [1, 15, 28, 41]. As recent large language models (LLM) (e.g., BERT [8]) have shown strong performances in retrieval tasks [7, 27, 40], some works utilize LLMs for story discovery but only in an *offline* and *supervised* setting [21, 39] with human-guided annotations which are expensive and easily outdated over time.

Motivation. In this work, for unsupervised online story discovery, we exploit off-the-shelf pretrained sentence encoders (PSE) based on LLMs and embed articles by using sentences as building blocks. This is a more practical and effective approach for dealing with text-rich and evolving news streams, because (1) a target corpus is dynamically changing, (2) true story labels are not readily available, and (3) modeling news articles with an LLM from scratch cause too high computational costs to support online applications.

A straightforward way to embed an article with a PSE is to average individual sentence representations in the article into a single article representation, which we refer to as *indiscriminative embedding*. While effective to some extent, compared with word- or article-level embedding with an LLM, it has clear limitations in considering the thematic contexts of news streams. For instance, in Figure 1, not all of the sentences equally contribute to representing the article for its story 2023 California Floods. While some sentences (e.g., s1, s22, and s34) are highly relevant to the story as described with thematic keywords implying the story, the other sentences (e.g., s9 and s35) are just generic or too specific descriptions that are not necessarily relevant to the story. The indiscriminative embedding with a PSE, however, results in fixed and deterministic article representations regardless of such thematic contexts of news streams and thus fails to maximize the story discovery performance.




Figure 2: Left: Average cosine similarities of news articles [11] embedded by LLMs [8, 34] with various granularity; higher (lower) is better for articles in the same (different) stories. Right: 2D-visualizations of sentences.

Main Idea and Challenges. The key idea we employ to resolve this problem is **thematic embedding** with a PSE, which dynamically embeds articles and stories with their shared themes timely captured in the latest news streams. The temporal themes help focus on only story-relevant parts of articles for story discovery so that articles in the same story can be represented closer while being further from other stories. As demonstrated in Figure 2, the thematic embedding identifies only story-relevant sentences which are more clearly clustered into distinct stories, and this naturally results in better article similarities than indiscriminative sentence-level embedding as well as word- or article-level embeddings.

However, implementing the idea of thematic embedding for unsupervised online story discovery pose considerable technical challenges: (1) unique themes of stories should be automatically identified (i.e., unsupervised) and timely updated (i.e., online) to derive up-to-date representations of articles and stories, (2) stories arbitrarily emerge and expire over time, so an adaptive mechanism is required to continuously cluster articles into stories, and (3) the embedding and clustering should be efficient and scalable to deal with massive news streams and support real-time applications.

Summary. To this end, we propose a novel framework, **USTORY**, for Unsupervised Story discovery with scalable thematic embedding. USTORY can be instantiated with any existing PSEs and employs two main techniques: *theme- and time-aware dynamic embedding* and *novelty-aware adaptive clustering*; the former systematically identifies temporal themes from diverse contexts of news streams and embed articles and stories by considering their theme and time relevance (Section 4). Then, the latter estimates article-story confidence scores to assign articles to existing stories or initiate novel stories. In the meantime, USTORY manages only the minimum but sufficient summary of stories so as to guarantee single-pass processing of the assigned articles (Section 5).

We summarize the main contributions of this work as follows:

- To the best of our knowledge, this is the first work to propose *thematic embedding* with off-the-shelf pretrained sentence encoders for unsupervised online story discovery from news streams.
- We propose a framework USTORY, implementing the idea of thematic embedding, which is *scalable* with single-pass processing and *compatible* with any pretrained sentence encoders. The source code is available at <https://github.com/cliveyn/USTORY>.
- We demonstrate that USTORY *outperforms* existing baselines by up to 51.7% in B^3 -F1 and their PSE-variants by up to 12.5% in B^3 -F1 in three news data sets. The *scalability* and *robustness* of USTORY are also verified through in-depth analyses.

2 RELATED WORK

Early work by Allan et al. [3] introduced topic detection and tracking (TDT) concepts for organizing and mining news articles. News story discovery is one of the most widely studied relevant research topics in TDT, which is crucial for various downstream applications such as news recommendation [32], summarization [47], and fine-grained event mining [52]. Previous efforts in the news story discovery can be classified as retrospective (offline) story discovery or prospective (online) story discovery, where the former generates a story timeline [4] or structure [13] from a given set of articles usually for analyzing domain-specific events while the latter processes continuous and unbounded news streams to discover and track stories in real-time, which is the main scope of this paper.

For online news story discovery, clustering-based approaches have been widely used for embedded articles. One line of studies adopted a *supervised* approach assuming some labeled articles and external knowledge (e.g., entity labels) are available to facilitate the embedding and clustering procedures. Specifically, labeled training data sets were used to learn an adjacency matrix of articles [21–23] or to learn a similarity threshold for cluster assignment [28, 38, 39]. For instance, Miranda et al. [28] introduced a multilingual clustering algorithm for news article streams, where an article-cluster similarity was measured by the weighted sum of TF-IDF subvectors similarities and temporal similarities modeled in a Gaussian distribution. The weights for the various types of similarities were learned from labeled training data sets. Some work fine-tuned an LLM with the labeled training data sets for embedding articles along with external entity knowledge [21, 39].

However, in an online setting, the labeled articles are rarely available and quickly outdated. The other line of studies thus tried to embed and cluster articles in an *unsupervised* manner. ConStream [1] is one of the popular clustering algorithms for document streams and has been widely used as a competitive baseline. It managed keyword frequencies of articles and clusters them with a notion of micro-cluster. A recent work NewsLens [15] found overlapping keywords of articles to build a local topic graph and applied a community detection algorithm to it to detect stories. Staykovski et al. [41] improved NewsLens further by using the sparse TF-IDF embedding (which outperformed the dense doc2vec [16] embedding alternatively proposed in the work) of articles to build a local topic graph. The existing methods, however, focus on explicit keywords statistics of articles, limiting the capability to capture implicit local contexts inside the articles. Furthermore, their fixed article representations do not fully capture evolving global contexts of news streams. In this work, we capture the local contexts of articles with a PSE, while dynamically exploiting it through thematic embedding to adapt to the global contexts of news streams.

Besides, some studies [24, 46] proposed *offline* and *supervised* fine-tuning of language models to specifically model news articles with relevant labeled data sets and tasks. Such fine-tuned models inherently output a fixed and deterministic representation of an input article, i.e., indiscriminative embedding, without considering the global contexts of news streams. However, their pretrained models can also be equipped in our framework for unsupervised online story discovery by being dynamically exploited over news streams through thematic embedding.

3 PROBLEM SETTING

Let $a = [s_1, s_2, \dots, s_{|a|}]$ be a news article (or simply *article*) composed of sentences s and $C = [a_1, a_2, \dots, a_{|C|}]$ be a news story (or simply *story*) composed of correlated articles with a unique theme, such as California Floods, and consecutively published for a certain duration of time. We assume that each article belongs to a single story, and each story has at least M articles. A news stream $\mathcal{A} (= \dots, a_{i-1}, a_i, a_{i+1}, \dots)$ is a continuous and unbounded sequence of articles consecutively published at a timestamp t_{a_i} . A *sliding window* \mathcal{W} of size W slid by S determines a context of the latest articles and ongoing stories in \mathcal{A} . We set $W = 7$ days and $S = 1$ day by default (i.e., if no articles are added to a story for a week, it is considered expired), while they can be alternatively set as the number of articles. Then, Definition 3.1 gives a formal definition of the problem considered in this work.

Definition 3.1. (UNSUPERVISED ONLINE STORY DISCOVERY) Given a news stream \mathcal{A} , the unsupervised online story discovery is to incrementally update a set $\mathbb{C}_{\mathcal{W}}$ of stories from the articles in every sliding window \mathcal{W} without any human supervision or story labels.

4 THEMATIC EMBEDDING

4.1 Motivation

The efforts to encode texts have been long made from symbolic-based models (e.g., bag-of-words) to recent LLMs [8, 25, 33]. As demonstrated in Figure 2, using individual word-level or entire article-level granularity for embedding an article with an LLM may not be optimal since they are either too fine-grained (former) or too coarse-grained (latter) to get its specific story-indicative semantics, which is shared within the same story but not in different stories. Exploiting sentence-level information, on the other hand, effectively balances the abstractness of semantics and naturally meets the input constraints for typical LLMs (e.g., 512 tokens). Recent PSEs, that fine-tune LLMs with benchmark data sets and tasks for specifically embedding sentences, have shown state-of-the-art sentence embedding capabilities across various domains [31, 34]. In this work, we exploit off-the-shelf PSEs by regarding sentences as building blocks for embedding articles.

Exploiting a PSE to embed a long article (i.e., tens of sentences) gives various design choices; for instance, concatenating or mean-pooling of sentences can be straightforward alternatives¹. PSEs can effectively capture the local semantics of individual sentences in an article, but as shown in Figure 1, the sentences are not necessarily relevant to the theme of the article’s story. Thus, USTORY employs *thematic embedding* with a PSE, which first identifies the temporal themes of articles given a particular context of news streams (Section 4.2) and then dynamically embeds the articles and stories considering their theme- and time-relevance (Section 4.3).

4.2 Temporal Theme Identification

Let *corpus* be a set of articles collected over a period of time and *corpora* be a set of the corpus. Then, each corpus must have a *temporal theme* that uniquely represents the corpus especially in the latest temporal context of corpora. We model the temporal theme through a keywords retrieval process because it is efficient

¹Refer to Section 6.3 for the comparison of different alternative embedding strategies.

to mine keywords with a simple tokenization pipeline and it is effective to represent a theme explicitly with a diverse combination of keywords. Specifically, we identify thematic keywords of a corpus in context corpora by conjunctively considering (1) *recency*, (2) *popularity*, and (3) *distinctiveness* of keywords.

For example, let a context corpora \mathbb{D} be the latest articles in news streams and a target corpus d be the articles about 2023 California Floods in \mathbb{D} . The keywords such as ‘Northern California’, ‘Evacuation order’, ‘Pacific Storm’, ‘Die’, and ‘Rescue’ may collectively describe the temporal theme of d , while their compositions and importances will change over time. Suppose that the term ‘Pacific Storm’ has consistently appeared in d for the last 7 days by L times in each day, while ‘Evacuation order’ has appeared more actively for the last 3 days by $2L$ times in each day. While ‘Evacuation order’ was less popular than ‘Pacific Storm’ for the last week (i.e., $2L \times 3$ days $< L \times 7$ days), it would be more valuable to represent the recent theme of d . At the same time, other recent popular keywords such as ‘Die’ or ‘Rescue’ would be devalued if they are also frequently observed in another corpus in \mathbb{D} (e.g., Russia-Ukraine Conflict).

To this end, we naturally incorporate *time-decaying property* into popular ranking functions for discriminative information retrieval (e.g., TF-IDF [2] and BM25L [26])² to identify thematic keywords.

Definition 4.1. (THEMATIC KEYWORDS) Given a target corpus d in a context corpora \mathbb{D} , a set \mathcal{K}_d of the top N thematic keywords that best describe the temporal theme of d at a time t_c is

$$\mathcal{K}_d = \{(k_1, w_{k_1}), (k_2, w_{k_2}), \dots, (k_N, w_{k_N})\}, \text{ where}$$

$$w_k = \text{rec-pop}(k, d, t_c) \cdot \text{dist}(k, \mathbb{D})$$

$$= \sum_{k^j \in \mathcal{T}_d} \exp\left(-\frac{|t_{k^j} - t_c|}{\delta_k}\right) \cdot \log\left(\frac{|d_i \in \mathbb{D}| + 1}{|d_i \in \mathbb{D} : k \in \mathcal{T}_{d_i}| + 1} + 1\right), \quad (1)$$

where k_j is a single- or multi-token term appearing in d ranked by its importance w_{k_j} and \mathcal{T}_d is a set of all term appearances in d .

The score function $\text{rec-pop}(k, d, t_c)$ is time-decaying term frequency, where each term appearance k^j at time t_{k^j} is counted (for *popularity*) while being exponentially decayed by its temporal difference from t_c (for *recency*). The score function $\text{dist}(k, \mathbb{D})$ is inverse corpus frequency to measure how unique k is in \mathbb{D} (for *distinctiveness*). The decaying factor δ_k controls the degree of decaying and can be set to the total time span of \mathbb{D} (i.e., $\delta_k = \max(t_{k^j}) - \min(t_{k^j}) + 1$ for $k^j \in \mathcal{T}_{\mathbb{D}}$).

4.3 Theme/Time-aware Dynamic Embedding

Article Embedding. A temporal theme can act as key guidance in embedding an article with a PSE; the article is best represented with the theme by focusing on only the theme-relevant parts of the article. Given a certain temporal theme, as a form of thematic keywords, we dynamically represent an article by pooling the representations of sentences in the article weighted by their *theme relevance*, which considers both the frequency and the importance of the thematic keywords found in each sentence.

Definition 4.2. (ARTICLE REPRESENTATION) Given a thematic keywords set \mathcal{K}_d , derived from a target corpus d in context corpora \mathbb{D} ,

²Any frequency-based function can be adopted with the time-decaying property. We used TF-IDF as default which showed better results than the other functions.




Figure 3: Thematic embedding with an article in Figure 1 and its story 2023 California Floods as the target article and story.

a representation $E_{a|d}$ of a target article a given d is

$$E_{a|d} = \sum_{s_i \in a} \frac{\sum_{k_i \in \mathcal{K}_d} |k_i^j \in \mathcal{T}_{s_i}| w_{k_i}}{\sum_{k_i \in \mathcal{K}_d} |k_i^j \in \mathcal{T}_a| w_{k_i}} enc(s_i), \quad (2)$$

where \mathcal{T}_a and \mathcal{T}_s are the term appearance sets of a and its sentence s , respectively, and $enc(s)$ is a representation of s by a PSE.

Story Embedding. As a story is basically a cluster of articles, a typical way to represent it is to average all the incorporated article representations (i.e., a cluster center). However, such a static story embedding does not correctly capture the temporal theme of the story, which gradually evolves with newly added articles. Thus, we dynamically represent a story given a target article (i.e., at a specific time of the article) by pooling the representations of articles in the story weighted by their *time relevance* to the target article.

Definition 4.3. (STORY REPRESENTATION) A representation $E_{C|a}$ of a target story C given a target article a is

$$E_{C|a} = \sum_{a_i \in C} \frac{\exp(-|t_a - t_{a_i}|/\delta_C)}{\sum_{a_j \in C} \exp(-|t_a - t_{a_j}|/\delta_C)} E_{a_i|C}, \quad (3)$$

where the time-decaying property is applied to the temporal distance and the decaying factor δ_C can be set to the total time span of the story (i.e., $\delta_C = \max(t_{a_i}) - \min(t_{a_i})$ for $\forall a_i \in C$).

Figure 3 illustrates an example of thematic embedding. Suppose that an article in Figure 1 is a target article a , and a story C of articles about 2023 California Floods is a target corpus d . Then, the thematic keywords set \mathcal{K}_C is identified as a temporal theme of C . When embedding a given C , the theme-relevant sentences (e.g., s_{01} , s_{22} , and s_{34}) served as key ingredients for representing a . At the same time, when embedding C given a , the articles in C temporally close to a contribute more to represent C . Note that articles and stories can be dynamically represented depending on the target articles and stories, facilitating more effective story discovery.

Algorithm 1: Overall Procedure of USTORY

Input: a news stream \mathcal{A} , a pre-trained sentence encoder enc
Output: a set \mathcal{C}_W of stories in every sliding window W

```

1 for every sliding window  $W$  from  $\mathcal{A}$  do
2   for new articles  $a_j$  in  $W$  do
3      $[enc(s)|s \in a_j] \leftarrow$  sentence representations of  $a_j$ 
4   if  $\mathcal{C}_W$  is empty then
5     /* Novel (initial) Story Discovery (Section 5.2) */
6      $\mathcal{C}_W \leftarrow$  initialize seed stories from new articles;
7   /* Confidence-based Story Assignment (Section 5.3) */
8   for unassigned article  $a_j$  in  $W$  do
9      $C_i^* \leftarrow \operatorname{argmax}_{C_i \in \mathcal{C}_W} conf_{a_j|C_i}$  // The most confident story
10    if  $conf_{a_j|C_i^*} \geq \gamma$  then
11      Assign  $a_j$  to  $C_i^*$ 
12      Update  $PSS_{C_i^*}$  // Story summary (Section 5.4)
13    /* Novel Story Discovery (Section 5.2) */
14   $\mathcal{C}_W \leftarrow$  add seed stories from unassigned articles;
15  Return  $\mathcal{C}_W$ 

```

5 NOVELTY-AWARE ADAPTIVE CLUSTERING

5.1 Overview

Using the thematic embedding in Section 4, USTORY incrementally clusters articles into stories, while adaptively discovering novel stories. The overall procedure of USTORY is illustrated in Figure 4 and outlined in Algorithm 1. For every sliding window of a news stream, USTORY gets sentence representations of new articles from a PSE (Lines 1-3) and conducts two steps. First, if there are no existing stories in the current window, as *novel story discovery* step, new articles are represented with their own theme, and seed stories are identified by cluster center initialization (Lines 4-5). Then, as *confidence-based story assignment* step (Lines 6-10), unassigned articles (including the new articles) are inspected if they can be confidently added to one of the existing stories. Each article-story pair is dynamically embedded and their confidence scores are derived from their thematic similarity. An article is assigned to the most confident story if the score exceeds a threshold. In the meanwhile, a summary of existing stories is utilized and updated. Finally, the novel story discovery step is conducted for the remaining unassigned articles to form new stories (Line 11) and all the discovered current stories are reported (Line 12). The two steps are detailed in the following sections.

5.2 Novel Story Discovery

Initial Article Embedding. When there are no existing stories or confident stories to be assigned to, unassigned articles (e.g., newly published or previously unassigned) are used to find novel seed stories. Since there are no existing themes or adequate themes to be considered, the unique theme of each article itself is identified for the thematic embedding of the article. Specifically, in Definition 4.1, the context corpora \mathbb{D} becomes all articles in a sliding window W and the target corpus d becomes a target article a . Then, the *article-indicative* thematic keywords set \mathcal{K}_a are derived and it leads to the representation $E_{a|\{a\}}$ of the article a by Definition 4.2.




Figure 4: The overall procedure of USTORY.

Seed Stories Discovery. Once the articles are embedded to better reveal their own themes, the articles under similar themes are more likely to be closer than those under different themes. Thus, a typical cluster center initialization technique can be applied to the thematically embedded initial representations of unassigned articles to find novel seed stories. Motivated by the popularly used k-means++ initialization [6], USTORY finds the seed centers with the lowest inertia (i.e., the sum of cosine similarities between each article and its centers) to get the most thematically distinctive seed stories. The number of seeds can be decided by dividing the number of unassigned articles by the minimum number M of articles to initiate a story, which is dependent on user preferences and application requirements³. If there are no such constraints available, an existing data-driven approach for deciding the number of clusters such as LOG-Means [10] or Silhouette Coefficient [35] can be applied.

5.3 Confidence-based Story Assignment

When there are existing stories, each unassigned article is evaluated to be assigned to one of the stories. Specifically, the unique temporal themes of existing stories are updated, and then the thematic similarity between each pair of stories and articles is estimated to get the robust article-story confidence score.

Thematic Embedding. A *story-indicative* thematic keywords set \mathcal{K}_C is derived by Definition 4.1, by setting the context corpora \mathbb{D} to a set \mathcal{C}_W of stories in a sliding window W and the target corpus d to the articles in a story C . Then, for each pair of article a and story C , the article representation $E_{a|C}$ and the story representation $E_{C|a}$ are derived by Definition 4.2 and 4.3, respectively.

Article-Story Thematic Similarity. Given a pair of articles and stories, USTORY quantifies their *thematic similarity* by conjunctively considering their semantic themes and symbolic themes. In brief, the former is estimated by the cosine similarity between their thematic representations, and the latter is estimated by the divergence of their thematic keyword distributions. These two types of similarities complement each other to estimate more robust thematic similarities.

For instance, some articles about 2023 California Floods might have a few sentences of slightly different thematic semantics according to the writers' perspectives (e.g., one may focus on the rescue while another may focus more on the victim), but the overall thematic keywords distributions of these articles would be similar as they are describing the same event. On the other hand, some articles about 2023 California Floods and the other articles about Russia-Ukraine Conflict might happen to have a few sentences of

similar thematic semantics describing casualties, but their overall thematic keywords distributions would be different (e.g., for articles in 2023 California Floods 'die' is co-occurred more frequently with 'floods' or 'storm'). Definition 5.1 formally formulates the thematic similarity between an article and a story.

Definition 5.1. (THEMATIC SIMILARITY) A thematic similarity between an article a and a story C is calculated as

$$\text{sim}_{\text{theme}}(a, C) = \max(0, \cos(E_{a|C}, E_{C|a})) \cdot \text{JSD}(P_{a, \mathcal{K}_C} \| P_{C, \mathcal{K}_C}). \quad (4)$$

The first term is a cosine similarity, with the negative values truncated to zero, between the thematic representations of a and C . The second term is the JS-divergence [20]⁴ between the thematic keyword probability distributions of a and C . The keyword probability $P(k_i|a, \mathcal{K}_C)$ in P_{a, \mathcal{K}_C} is estimated as $\frac{|k_i \in \mathcal{T}_a|}{\sum_{k_j \in \mathcal{K}_C} |k_j \in \mathcal{T}_a|}$ where \mathcal{T}_a is the term appearance set of a (and similarly for $P(k_i|C, \mathcal{K}_C)$ in P_{C, \mathcal{K}_C}).

Article-Story Assignment. Finally, USTORY derives a confidence score $\text{conf}_{a, C} \in [0, 1]$ for an article a to be assigned to a story C by comparing their thematic similarity with the other thematic similarities of all possible candidate assignments.

Definition 5.2. (ARTICLE-STORY CONFIDENCE) Given a target article a and a set of candidate stories $C_i \in \mathcal{C}_W$, the article-story confidence score for a to be assigned to C_i is

$$\text{conf}_{a, C_i} = \frac{\exp(T \cdot \text{sim}_{\text{theme}}(a, C_i))}{\sum_{C_j \in \mathcal{C}_W} \exp(T \cdot \text{sim}_{\text{theme}}(a, C_j))}, \quad (5)$$

where T is a temperature for scaling the score distribution.

Then, the article is assigned to the story with the highest confidence score if it exceeds the threshold $\gamma = 1 - (1 - 1/|\mathcal{C}_W|)^T$, indicating an adjusted random assignment probability [17], or otherwise remains unassigned⁵. The unassigned articles are used to find seed stories as introduced in Section 5.2 and repeatedly inspected to be assigned to the updated stories in later sliding windows.

5.4 Scalable Processing with Story Summary

Story Summary. Identifying temporal themes and dynamically embedding articles and stories from scratch in every sliding window cause considerable computation overheads, which is not practical in an online scenario. To realize scalable online processing, USTORY uses a novel data structure, called *pane-based story summary*, motivated by the pane-based aggregation [19] that has been widely adopted for efficient and scalable stream algorithms [42, 48, 49, 51].

⁴Among popular divergence measures [14, 20, 36], we chose JS-divergence [20] since it is bounded within a finite interval $\in [0, 1]$, can be defined even when a keyword exists only in a story, and has a low time complexity, i.e., $O(|\mathcal{K}_C|)$.

⁵The sensitivity analysis and the guidance on T are given in Section 6.4

³Typically, a news story has 5 to 20 articles on its first day in the real data sets [11, 28].

Definition 5.3. (PANE-BASED STORY SUMMARY (PSS)) Let panes p_i be non-overlapping subsets of consequent articles in a news stream and story panes p_i^C of a story C be a set of articles in both C and p_i . Then, a pane-based story summary of C ,

$$PSS_C = \{ \langle p_i : \langle |p_i^C|, tf(\cdot, p_i^C), \sum_{a_j \in p_i^C} E_{a_j|C} \rangle \rangle \}, \quad (6)$$

maps a pane to the triplet of the number of articles, the term frequencies, and the sum of article representations in p_i^C .

The size of a pane determines the granularity of the story summary. While any common divisor of the window size and the slide size can be used, we set the pane size to be the slide size (i.e., the greatest common divisor) to maximize efficiency. In the confidence-based story assignment step in every sliding window, USTORY uses PSS to identify temporal themes and derive dynamic article and story representations, without accessing all previous articles in stories. Note that USTORY updates the triplet in PSS in an *additive* manner whenever articles in stories are updated as the window slides. The sufficiency of PSS and the complexity of USTORY realized by utilizing PSS are analyzed as follows.

Efficiency Analysis. With the help of PSS, USTORY guarantees *single-pass processing* on the assigned articles; once an article is assigned to a story it can be discarded and only PSS is used for the following procedures. Theorem 5.4 and Theorem 5.5 respectively proves the sufficiency of PSS for the story assignment and shows the time and space complexities of USTORY when utilizing PSS.

THEOREM 5.4. (SUFFICIENCY OF PSS) *The confidence-based story assignment step requires to identify thematic keywords \mathcal{K}_C and a story representation $E_{C|a}$ for each story C and each unassigned article a (note that $E_{a|C}$ is directly derived from \mathcal{K}_C and a). A pane-based story summary PSS_C is sufficient for deriving \mathcal{K}_C and $E_{C|a}$.*

PROOF. \mathcal{K}_C and $E_{C|a}$ can be accurately derived by only using the triplet information in PSS_C . First, by Definition 4.1, the importance w_k of a thematic keyword k with a story C in a sliding window \mathcal{W} at a current timestamp t_c is computed from $rec\text{-}pop(k, C, t_c)$ and $dist(k, \mathbb{C}_\mathcal{W})$. Since a pane p_i represents the articles in the same slide in \mathcal{W} , i.e., $\forall a \in p_i : t_a = t_{p_i}$, $rec\text{-}pop(k, C, t_c)$ and $dist(k, \mathbb{C}_\mathcal{W})$ are calculated with the term frequency $tf(k, p_i^C)$ in PSS_C .

$$rec\text{-}pop(k, C, t_c) = \sum_{p_i \in \mathcal{W}} \left(\exp\left(-\frac{|t_{p_i} - t_c|}{\delta_k}\right) \cdot tf(k, p_i^C) \right) \text{ and} \quad (7)$$

$$dist(k, \mathbb{C}_\mathcal{W}) = \log\left(\frac{|\mathbb{C}_\mathcal{W}| + 1}{\sum_{p_i \in \mathcal{W}} \mathbb{1}_{tf(k, p_i^C) > 0} + 1}\right).$$

Then, in Definition 4.3, $E_{C|a}$ can be reformulated as the sum of time-decaying article representations divided by the time-decaying count of articles. Thus, it can be calculated with the articles count $|p_i^C|$ and the article representations sum $\sum_{a_j \in p_i^C} E_{a_j|C}$ in PSS_C .

$$E_{C|a} = \frac{\sum_{p_i \in \mathcal{W}} (\exp(-|t_a - t_{p_i}|/\delta_C) \cdot \sum_{a_j \in p_i^C} E_{a_j|C})}{\sum_{p_i \in \mathcal{W}} (\exp(-|t_a - t_{p_i}|/\delta_C) \cdot |p_i^C|)}. \quad (8)$$

□

THEOREM 5.5. (COMPLEXITIES OF USTORY) *Let N_W and N_S be the numbers of articles in a window and a slide, respectively, and N_C be the number of existing stories. Then, the time and space complexities of USTORY are $O(N_C N_W + N_S)$ and $O(N_W + \frac{N_W}{N_S} N_C)$, respectively.*

Table 1: Data sets.

Data set	#Articles (Avg #Sentences)	#Stories (Avg #Articles/day)
Newsfeed [28, 37]	16,136 (21.4)	788 (8)
WCEP18 [11]	47,038 (26.9)	828 (18)
WCEP19 [11]	29,931 (27.6)	519 (18)
USNews (case study)	2,744 (34.9)	-

PROOF. The time complexity of USTORY is specifically divided into that for each step: $O(N_S)$ for embedding articles in every sliding window with a PSE, $O(N_C N_W)$ for the novelty story discovery step, $O(N_C N_W)$ for the confidence-based story assignment step, and $O(N_S)$ for updating PSS. Thus, the total time complexity of USTORY is $O(N_C N_W + N_S)$. Similarly, the space complexity of USTORY is specifically divided into that for managing stories and articles: $O(\frac{N_W}{N_S} N_C)$ for managing PSS where $\frac{N_W}{N_S}$ is the number of panes and $O(N_W)$ for managing articles. Thus, the total space complexity of USTORY is $O(N_W + \frac{N_W}{N_S} N_C)$. □

Since $N_C, N_S \ll N_W$ in practice, the time and space complexities of USTORY are *linear* to N_W , mainly affected by a window size W .

6 EXPERIMENTS

We conducted extensive experiments to evaluate the performance of USTORY, of which results are summarized as follows.

- USTORY outperformed the existing unsupervised online story discovery algorithms and their variants with a PSE in benchmark news data sets in terms of B^3 -F1, AMI, and ARI (Section 6.2).
- The main idea employed in USTORY was demonstrated to be effective through ablation studies on the theme- and time-aware components (Section 6.2) and comparison with alternative embedding strategies (Section 6.3).
- USTORY was scalable to the variation of sliding window sizes and robust to hyperparameters (Section 6.4).
- USTORY discovered more quality embedding space and news stories than baselines from a real news stream (Section 6.5).

6.1 Experiment Setting

News Streams. We used four real news data sets for evaluation, summarized in Table 1. Newsfeed [28, 37] is a multilingual news data set collected from Newsfeed Service [43] in 2014, and we used English news articles with story labels. WCEP [11] is a benchmark news data set collected from Wikipedia Current Event Portal and the Common Crawl archive. We used articles in the stories of at least 50 articles and published in 2018 and 2019 (i.e., WCEP18 and WCEP19). For a qualitative case study, we prepared USNews by collecting news articles through NewsAPI [30] with a query of ‘United States’ for a month. Each data set is simulated as a news stream by feeding articles into sliding windows in chronological order. The window size W and the slide size S of the sliding window are set to 7 days and a day, respectively.

Compared Algorithms. We compared USTORY with the five online story discovery algorithms: *ConStream*[1], *NewsLens* [15], *StoryForest* [22], *Miranda* [28], and *Staykovski* [41]⁶. For a fair comparison, we also prepared their PSE-variants by adopting a PSE

⁶StoryForest and Miranda are adopted for an unsupervised setting.

Table 2: Performance comparison results (the highest scores are highlighted in bold).

		Newsfeed			WCEP18			WCEP19		
		B^3 -F1	AMI	ARI	B^3 -F1	AMI	ARI	B^3 -F1	AMI	ARI
Baselines (ordered by the average B^3 -F1)	ConStream	0.314	0.128	0.069	0.408	0.444	0.222	0.400	0.497	0.292
	NewsLens	0.481	0.309	0.077	0.527	0.490	0.117	0.554	0.529	0.141
	StoryForest	0.696	0.725	0.592	0.673	0.765	0.523	0.697	0.798	0.596
	Miranda	0.706	0.726	0.572	0.694	0.786	0.571	0.698	0.791	0.574
	Staykovski	0.669	0.602	0.358	0.697	0.759	0.487	0.701	0.765	0.487
PSE-variants of the top baselines	Miranda-SenT5	0.732	0.753	0.617	0.710	0.798	0.629	0.717	0.805	0.644
	Staykovski-SenT5	0.684	0.631	0.415	0.735	0.798	0.582	0.704	0.782	0.537
	Miranda-SenRB	0.764	0.785	0.648	0.751	0.835	0.656	0.759	0.837	0.657
	Staykovski-SenRB	0.750	0.720	0.567	0.754	0.824	0.642	0.762	0.830	0.660
Proposed	USTORY-SenT5	0.751*	0.763*	0.638*	0.780*	0.846*	0.694*	0.799*	0.861*	0.733*
	over baselines	Δ 44.4%	Δ 136.0%	Δ 330.1%	Δ 35.9%	Δ 38.6%	Δ 160.5%	Δ 37.4%	Δ 33.0%	Δ 134.4%
	over baselines-SenT5	\blacktriangle 6.2%	\blacktriangle 11.1%	\blacktriangle 28.6%	\blacktriangle 8.0%	\blacktriangle 6.0%	\blacktriangle 14.8%	\blacktriangle 12.5%	\blacktriangle 8.5%	\blacktriangle 25.2%
	USTORY-SenRB	0.789*	0.812*	0.699*	0.810*	0.871*	0.739*	0.825*	0.880*	0.765*
	over baselines	Δ 51.7%	Δ 151.2%	Δ 371.3%	Δ 41.1%	Δ 42.7%	Δ 177.4%	Δ 41.9%	Δ 36.0%	Δ 144.6%
Variants of USTORY-SenRB	over baselines-SenRB	\blacktriangle 4.2%	\blacktriangle 8.1%	\blacktriangle 15.6%	\blacktriangle 7.6%	\blacktriangle 5.0%	\blacktriangle 13.9%	\blacktriangle 8.5%	\blacktriangle 5.6%	\blacktriangle 16.2%
	w/o time-aware	0.780	0.801	0.682	0.801	0.864	0.736	0.817	0.875	0.760
	w/o theme-aware	0.767	0.792	0.665	0.771	0.842	0.679	0.790	0.856	0.703
	w/o both	0.752	0.777	0.651	0.753	0.826	0.677	0.770	0.843	0.709

* denotes statistically significant improvement ($p < 0.05$ with the t-test) over the compared algorithms and Δ , \blacktriangle indicate the average improvement ratio.

for embedding articles by averaging their sentence representations. We used two PSEs for USTORY and the variants of baselines: Sentence-BERT [34]⁷ (i.e., AlgName-SenRB) and Sentence-T5 [31]⁸ (i.e., AlgName-SenT5). Please note that for USTORY and the PSE-variants of baselines, Sentence-BERT (i.e., -SenRB) is used as a default PSE unless otherwise specified.

Implementation Details. All compared algorithms were implemented with *Python 3.8.8* and evaluated on a Linux server with AMD EPYC 7502 32-Core CPU and 1TB RAM. We tokenized sentences and terms in articles by *SpaCy 3.2.0* with *en_core_web_lg* pipeline [9] and counted (1,2)-gram terms by *Sikit-learn 0.24.2*. For each algorithm, we used the default hyperparameters following the original work or tuned them to get the best results; specifically, the keywords size $N \in [1, 15]$ and the temperature $T \in [1, 5]$ for USTORY ($N = 10$ and $T = 2$ were used by default); the standard score threshold for the micro-cluster assignment $\in [0, 3]$ for *ConStream*; the number of overlapping keywords for creating an edge $\in [1, 5]$ and the similarity threshold for merging communities $\in [0, 1]$ for *NewsLens*; the number of keywords for an article $\in [1, 20]$ to form a keyword graph for *StoryForest*; the similarity threshold for cluster assignment $\in [1, 5]$ for *Miranda*; and the similarity threshold for creating an edge $\in [0, 0.5]$ and that for merging clusters $\in [0.5, 1]$ for *Staykovski*. The minimum number M of articles to form a valid story was set to the average number of articles in a story in a day, which was 8 for Newsfeed and 18 for WCEP18 and WCEP19.

Evaluation Metrics: We used B-cubed F1 score (B^3 -F1) [5] for evaluating article-wise cluster quality and used Adjusted Mutual Information (AMI) [45] and Adjusted Rand Index (ARI) [12] for evaluating the mutual information and similarity of clustering results, respectively, adjusted for chance. For each metric, the average score over all sliding windows is reported to evaluate the overall performances over the entire news streams.

⁷<https://huggingface.co/sentence-transformers/all-roberta-large-v1>

⁸<https://huggingface.co/sentence-transformers/sentence-t5-large>

6.2 Story Discovery Accuracy

Table 2 shows the overall evaluation results of all algorithms.

Comparison with Baselines. USTORY achieved higher B^3 -F1, AMI, and ARI scores than baselines for all cases. For instance, USTORY-SenRB achieved 44.9% higher B^3 -F1, 76.6% higher AMI, and 231.1% higher ARI than baselines when averaged over all cases. We also compared the PSE-variants of the top two baselines, Miranda and Staykovski (note that the PSE-variants of the other algorithms consistently performed worse than them). The PSE-variants improved all scores from their original versions, which shows the generic embedding power of a PSE. USTORY, however, still outperformed them consistently in all cases for both SenT5 and SenRB. Specifically, USTORY-SenRB achieved 6.8% higher B^3 -F1, 6.2% higher AMI, and 15.2% higher ARI than the SenRB-variants of the top two baselines when averaged over all cases. This clearly shows that USTORY exploits a PSE more effectively with the help of thematic embedding.

Ablation Study of USTORY. We verified the efficacy of the theme and time-aware components employed in USTORY by preparing the three variants:

- *w/o time-aware* does not consider the recency in thematic keywords and the time relevance in embedding articles and stories.
- *w/o theme-aware* does not consider the theme relevance in embedding articles and stories.
- *w/o both* does not consider both the time- and theme-aware components described above.

USTORY consistently achieved higher scores than the above three variants, which indicates that both components jointly help to discover stories. More specifically, the theme-aware component contributed more than the time-aware component to the performance. These results show that temporally close articles in a story do not necessarily share thematically similar content only, so it is more critical to identify and filter out theme-irrelevant parts in each article for online story discovery.



Figure 5: B^3 -F1 scores of USTORY with four strategies.

Figure 6: Varying sliding window sizes.

6.3 Embedding Strategies Comparison

We evaluated the story discovery performance of USTORY with four embedding strategies to exploit a PSE:

- *IndSentMean*: mean pooling of sentence representations
- *IndSentConc*: a representation of concatenated sentences
- *ThemSentMean* (default in USTORY): weighted pooling of the thematically prioritized sentence representations
- *ThemSentConc*: a representation of the concatenation of thematically prioritized sentences.

Note that USTORY automatically identifies temporal themes of news streams to prioritize sentences for ThemSentMean and ThemSentConc. We also varied the number of sentences used for embedding articles for further understanding of the compared strategies. Figure 5 shows the B^3 -F1 results, while the results of AMI and ARI showed similar trends. The proposed ThemSentMean strategy in USTORY showed the highest scores throughout the varying number of sentences. This indicates that thematic embedding effectively exploits all information in the article to discover stories. In general, the Mean strategies, IndSentMean and ThemSentMean, outperformed the Concatenate strategies, IndSentConc and ThemSentConc, because the concatenated sentences can not incorporate all information in the article due to the input sequence limit of a PSE and also fail to capture local semantics. Meanwhile, it is notable that ThemSentConc still took advantage of the thematic prioritization of sentences and resulted in higher performance than IndSentConc.

6.4 Scalability and Sensitivity Analysis

Scalability of USTORY. We analyzed the scalability of USTORY and the top two baselines on varying window sizes (W), i.e., the number of articles in a window, which is the most critical factor affecting scalability. Figure 6 shows the average wall clock time for the compared algorithms to process each sliding window and





Figure 7: Processing time breakdown of USTORY.

Figure 8: Varying keywords size N and temperature T .

their B^3 -F1 results. Note that AMI and ARI results showed similar trends. While the top two baselines resulted in higher processing time due to the high computation cost for community detection (Staykovski) and clustering (Miranda), USTORY took a much shorter processing time for each sliding window. The increasing rate of processing time over larger windows for USTORY was lower than Miranda and comparable to Staykovski in most cases. USTORY also consistently achieve higher B^3 -F1 scores than the other algorithms. This demonstrates the efficacy of story summary (PSS) for efficient and accurate story discovery.

Processing Time Breakdown of USTORY. We further broke down the processing time of USTORY into the four main steps in Section 5. As shown in Figure 7, the initial article embedding step took the most processing time, followed by the story assignment step, the story summary (PSS) update step, and the seed stories discovery step. To further improve the scalability of USTORY, especially for time-critical applications, the initial embedding step can be alternatively substituted with the IndSentMean strategy discussed in Section 6.3, which are expected to be faster but still effective for initialization purposes. Meanwhile, the last three steps in Newsfeed accounted for more portions of the total processing time compared with their proportions in the other data sets. This also conforms to the complexity analysis result, as Newsfeed has more concurrent stories in sliding windows.

Sensitivity of USTORY. Figure 8 shows the sensitivity of USTORY on the two main hyperparameters, keywords size N and the temperature T , for the three performance metrics. The performances of USTORY converge early at some points and become near optimal around the default values (i.e., $N = 10$ and $T = 2$). These trends show that setting the adequate number of keywords for identifying temporal themes (e.g., $N \geq 3$) and flattening the article-story confidence score distribution to at least some degree (e.g., $T \geq 2$) can lead to a robust story discovery performance.




Figure 9: The timeline and visualizations of stories discovered by USTORY in USNews (best viewed in color).




Figure 10: Visualizations of the clusters about Russia-Ukraine Conflict (in black circles) and North Korea Missile Test (mis-clustered by the other algorithms, in red stars).

6.5 Case Study

We conducted a qualitative case study on USNews with USTORY and the PSE-variants of the top two baselines. Figure 9 visualizes the ten discovered stories and their article embeddings with USTORY and the PSE-variants of baselines. The leftmost sub-figure shows a timeline of the sizes of stories with their titles (manually given by the authors for convenience) and their top five thematic keywords (automatically identified by USTORY).

USTORY successfully identified and tracked the important real news stories. For instance, the long-term stories such as Russia-Ukraine Conflict and Covid 19, which have been constant issues in recent years, were continuously discovered by USTORY throughout the whole period. At the same time, the short-term stories such as North Korea Missile Test and Capitol Attack Reflects that have been discussed for some days were also timely discovered by USTORY. The resulting embedding space also shows that the articles are distinctively clustered by their stories. Following the two embedding strategies discussed in Section 6.3, the right two sub-figures respectively visualizes the article representations embedded by ThemSentMean (used by default in USTORY) and IndSentMean (used in the PSE-variants of baselines) in a 2D-space through t-SNE [44]. As expected, the discovered stories are clearly distinguished with the ThemSentMean embedding, while they are more cluttered and mixed with the IndSentMean embedding.

We further analyzed the differences between the compared algorithms in detail by drilling down into a specific story, Russia-Ukraine Conflict. Figure 10 shows the 2D-visualizations of the articles in the clusters related to Russia-Ukraine Conflict discovered by each algorithm. It is clear that USTORY discovered a more compact cluster of articles about Russia-Ukraine Conflict. There are some articles about North Korea Missile Test mis-clustered together *only by the PSE-variants of baselines*. This is because some articles about North Korea Missile Test contain general contents on war, weapons, or conflicts, which are also frequently found in the articles about Russia-Ukraine Conflict. However, USTORY successfully distinguished the two stories in the embedding space by focusing on only the theme-relevant parts of articles; the example prioritized sentences are given in Table 3.

Table 3: The top-3 and bottom-3 sentences prioritized by USTORY in an example article about Russia-Ukraine Conflict.

s#	Sentence
s04	President Joe Biden has warned that Russia will face severe economic consequences if Russian President Vladimir Putin were to launch an invasion of Ukraine.
s15	Russia says it feels threatened by the prospect of the U.S. deploying offensive missile systems in Ukraine, even though Biden has assured Putin he has no intention of doing so.
s01	U.S. suggests it's open to limiting military exercises, missile deployments in talks with Russia.
...	...
s08	The Geneva talks, to be followed by other sessions next week in Brussels and Vienna, are aimed at averting a crisis.
s14	"Both sides would need to make essentially the same commitment."
s27	These actions could also restrict export of products made abroad if they contain more than a specified percentage of U.S. content.

7 DISCUSSION AND CONCLUSION

While we followed an unsupervised approach in this work, it would be promising to extend the idea of thematic embedding in a weakly supervised setting by assuming that story-indicative external knowledge is available (e.g., category [27], entity [39], taxonomy [18], or timeline [4]). They can auxiliary help identify temporal themes and also allow a more flexible time scope of ongoing stories; USTORY can refer to an extra set of landmark stories with their lightweight summaries to capture sparse stories with longer timelines.

In conclusion, this paper proposed USTORY, a scalable framework for unsupervised online story discovery from news streams. To deal with text-rich and evolving news streams, USTORY employs a novel idea of thematic embedding with a pretrained sentence encoder and conducts novelty-aware adaptive clustering by managing compact story summaries. Extensive experiments on real news data sets showed that USTORY outperforms the existing baselines and their variants, while being scalable and robust to various experiment settings. We believe that USTORY can be used for many downstream real-time applications to help people understand unstructured and massive new streams.

ACKNOWLEDGMENTS

The first author was supported by the National Research Foundation of Korea (Basic Science Research Program: 2021R1A6A3A14043765). The research was supported in part by US DARPA (FA8750-19-2-1004 and HR001121C0165), National Science Foundation (IIS-19-56151, IIS-17-41317, IIS 17-04532, 2019897, and 2118329), and the Institute of Information and Communications Technology Planning and Evaluation (IITP) in Korea (2020-0-01361).

REFERENCES

- [1] Charu C Aggarwal and S Yu Philip. 2010. On Clustering Massive Text and Categorical Data Streams. *Knowledge and Information Systems* 24, 2 (2010).
- [2] Akiko Aizawa. 2003. An information-theoretic perspective of tf-idf measures. *Information Processing & Management* 39, 1 (2003), 45–65.
- [3] James Allan, Jaime G Carbonell, George Doddington, Jonathan Yamron, and Yiming Yang. 1998. Topic Detection and Tracking Pilot Study Final Report.
- [4] Jeffery Ansah, Lin Liu, Wei Kang, Selasie Kwashie, Jixue Li, and Jiuyong Li. 2019. A graph is worth a thousand words: Telling event stories using timeline summarization graphs. In *Proceedings of the World Wide Web Conference*.
- [5] Amit Bagga and Breck Baldwin. 1998. Entity-based Cross-document Coreferencing using the Vector Space Model. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- [6] Bahman Bahmani, Benjamin Moseley, Andrea Vattani, Ravi Kumar, and Sergei Vassilvitskii. 2012. Scalable K-Means+. *Proceedings of the VLDB Endowment* 5, 7 (2012).
- [7] Zhuyun Dai and Jamie Callan. 2019. Deeper text understanding for IR with contextual neural language modeling. In *Proceedings of the international ACM SIGIR conference on research and development in information retrieval*, 985–988.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- [9] en_core_web_lg. 2023. https://spacy.io/models/en#en_core_web_lg
- [10] Manuel Fritz, Michael Behring, and Holger Schwarz. 2020. LOG-means: Efficiently Estimating the Number of Clusters in Large Datasets. *Proceedings of the VLDB Endowment* 13, 12 (2020), 2118–2131.
- [11] Demian Gholipour Ghalandari, Chris Hokamp, John Glover, Georgiana Ifrim, et al. 2020. A Large-Scale Multi-Document Summarization Dataset from the Wikipedia Current Events Portal. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- [12] Lawrence Hubert and Phipps Arabie. 1985. Comparing Partitions. *Journal of Classification* 2, 1 (1985), 193–218.
- [13] Hyungsul Kim, Ahmed El-Kishky, Xiang Ren, and Jiawei Han. 2019. Mining News Events from Comparable News Corpora: A Multi-Attribute Proximity Network Modeling Approach. In *Proceedings of the IEEE International Conference on Big Data*.
- [14] Solomon Kullback and Richard A Leibler. 1951. On Information and Sufficiency. *The Annals of Mathematical Statistics* 22, 1 (1951), 79–86.
- [15] Philippe Laban and Marti A Hearst. 2017. newsLens: building and visualizing long-ranging news stories. In *Proceedings of the Events and Stories in the News Workshop in conjunction with the Annual Meeting of the Association for Computational Linguistics*.
- [16] Quoc Le and Tomas Mikolov. 2014. Distributed Representations of Sentences and Documents. In *Proceedings of the International Conference on Machine Learning*.
- [17] Dongha Lee, Jiaming Shen, SeongKu Kang, Susik Yoon, Jiawei Han, and Hwanjo Yu. 2022. TaxoCom: Topic Taxonomy Completion with Hierarchical Discovery of Novel Topic Clusters. *Proceedings of the World Wide Web Conference* (2022).
- [18] Dongha Lee, Jiaming Shen, Seonghyeon Lee, Susik Yoon, Hwanjo Yu, and Jiawei Han. 2022. TopicExpan: Topic Taxonomy Expansion via Hierarchy-Aware Topic Phrase Generation. (2022).
- [19] Jin Li, David Maier, Kristin Tuft, Vassilis Papadimos, and Peter A Tucker. 2005. No Pane, No Gain: Efficient Evaluation of Sliding-window Aggregates over Data Streams. *ACM SIGMOD Record* 34, 1 (2005), 39–44.
- [20] Jianhua Lin. 1991. Divergence Measures based on The Shannon Entropy. *IEEE Transactions on Information Theory* 37, 1 (1991), 145–151.
- [21] Mathis Linger and Mhamed Hajaiej. 2020. Batch Clustering for Multilingual News Streaming. In *Proceedings of the International Workshop on Narrative Extraction from Texts held in conjunction with the European Conference on Information Retrieval*.
- [22] Bang Liu, Fred X Han, Di Niu, Linglong Kong, Kunfeng Lai, and Yu Xu. 2020. Story Forest: Extracting Events and Telling Stories from Breaking News. *ACM Transactions on Knowledge Discovery from Data* 14, 3 (2020), 1–28.
- [23] Bang Liu, Di Niu, Kunfeng Lai, Linglong Kong, and Yu Xu. 2017. Growing story forest online from massive breaking news. In *Proceedings of the ACM on Conference on Information and Knowledge Management*, 777–785.
- [24] Jialu Liu, Tianqi Liu, and Cong Yu. 2021. NewsEmbed: Modeling News through Pre-trained Document Representations. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- [25] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A Robustly Optimized Bert Pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [26] Yuanhua Lv and ChengXiang Zhai. 2011. When Documents are Very Long, BM25 Fails!. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [27] Yu Meng, Yunyi Zhang, Jiaxin Huang, Chenyan Xiong, Heng Ji, Chao Zhang, and Jiawei Han. 2020. Text Classification Using Label Names Only: A Language Model Self-Training Approach. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- [28] Sebastião Miranda, Arturs Znotins, Shay B Cohen, and Guntis Barzdins. 2018. Multilingual Clustering of Streaming News. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- [29] NBCNews. 2023. <https://www.nbcnews.com/news/weather/atmospheric-river-floods-northern-california-rcna63902>
- [30] NewsAPI. 2023. <https://newsapi.org>
- [31] Jianmo Ni, Gustavo Hernández Ábrego, Noah Constant, Ji Ma, Keith B Hall, Daniel Cer, and Yinfei Yang. 2021. Sentence-t5: Scalable Sentence Encoders from Pre-trained Text-to-text Models. *arXiv preprint arXiv:2108.08877* (2021).
- [32] Tao Qi, Fangzhao Wu, Chuhan Wu, Peijie Sun, Le Wu, Xiting Wang, Yongfeng Huang, and Xing Xie. 2022. ProFairRec: Provider fairness-aware news recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [33] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research* 21 (2020), 1–67.
- [34] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese Bert-Networks. *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- [35] Peter J Rousseeuw. 1987. Silhouettes: a Graphical Aid to the Interpretation and Validation of Cluster Analysis. *J. Comput. Appl. Math.* 20 (1987), 53–65.
- [36] Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. 2000. The Earth Mover’s Distance as a Metric for Image Retrieval. *International Journal of Computer Vision* 40, 2 (2000), 99–121.
- [37] Jan Rupnik, Andrej Muhic, Gregor Leban, Primoz Skraba, Blaz Fortuna, and Marko Grobelnik. 2016. News across languages-cross-lingual document similarity and event tracking. *Journal of Artificial Intelligence Research* 55 (2016), 283–316.
- [38] João Santos, Afonso Mendes, and Sebastião Miranda. 2022. Simplifying Multilingual News Clustering Through Projection From a Shared Space. *arXiv preprint arXiv:2204.13418* (2022).
- [39] Kailash Karthik Saravanakumar, Miguel Ballesteros, Muthu Kumar Chandrasekaran, and Kathleen McKeown. 2021. Event-Driven News Stream Clustering using Entity-Aware Contextual Embeddings. In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics*.
- [40] Suzanna Sia, Ayush Dalmia, and Sabrina J Mielke. 2020. Tired of Topic Models? Clusters of Pretrained Word Embeddings Make for Fast and Good Topics too!. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- [41] Todor Staykovski, Alberto Barrón-Cedeno, Giovanni Da San Martino, and Preslav Nakov. 2019. Dense vs. Sparse Representations for News Stream Clustering. In *Proceedings of the International Workshop on Narrative Extraction from Texts held in conjunction with the European Conference on Information Retrieval*.
- [42] Kanat Tangwongsan, Martin Hirzel, and Scott Schneider. 2019. Optimal and General Out-of-order Sliding-window Aggregation. *Proceedings of the VLDB Endowment* 12, 10 (2019), 1167–1180.
- [43] Mitja Trampuš and Blaz Novak. 2012. Internals of an aggregated web news feed. In *Proceedings of 15th Multiconference on Information Society*, 221–224.
- [44] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing Data Using t-SNE. *Journal of Machine Learning Research* 9, 11 (2008).
- [45] Nguyen Xuan Vinh, Julien Epps, and James Bailey. 2010. Information Theoretic Measures for Clusterings comparison: Variants, Properties, Normalization and Correction for Chance. *Journal of Machine Learning Research* 11 (2010), 2837–2854.
- [46] Chuhan Wu, Fangzhao Wu, Yang Yu, Tao Qi, Yongfeng Huang, and Qi Liu. 2021. NewsBERT: Distilling Pre-trained Language Model for Intelligent News Application. In *Findings of the Association for Computational Linguistics: EMNLP*.
- [47] Susik Yoon, Hou Pong Chan, and Jiawei Han. 2023. PDSum: Prototype-driven Continuous Summarization of Evolving Multi-document Sets Stream. *arXiv preprint arXiv:2302.05550* (2023).
- [48] Susik Yoon, Jae-Gil Lee, and Byung Suk Lee. 2019. NETS: Extremely Fast Outlier Detection from a Data Stream via Set-based Processing. *Proceedings of the VLDB Endowment* 12, 11 (2019).
- [49] Susik Yoon, Jae-Gil Lee, and Byung Suk Lee. 2020. Ultrafast local outlier detection from a data stream with stationary region skipping. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [50] Susik Yoon, Youngjun Lee, Jae-Gil Lee, and Byung Suk Lee. 2022. Adaptive Model Pooling for Online Deep Anomaly Detection from a Complex Evolving Data Stream. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [51] Susik Yoon, Yooju Shin, Jae-Gil Lee, and Byung Suk Lee. 2021. Multiple dynamic outlier-detection from a data stream by exploiting duality of data and queries. In *Proceedings of the 2021 International Conference on Management of Data*.
- [52] Yunyi Zhang, Fang Guo, Jiaming Shen, and Jiawei Han. 2022. Unsupervised Key Event Detection from Massive Text Corpora. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.