

Examining the State-of-the-Art in News Timeline Summarization

Demian Gholipour Ghalandari^{1,2} and Georgiana Ifrim¹

¹Insight Centre for Data Analytics, University College Dublin, Ireland

²Aylien Ltd., Dublin, Ireland

{demian.gholipour|georgiana.ifrim}@insight-centre.org

Abstract

Previous work on automatic news timeline summarization (TLS) leaves an unclear picture about how this task can generally be approached and how well it is currently solved. This is mostly due to the focus on individual subtasks, such as date selection and date summarization, and to the previous lack of appropriate evaluation metrics for the full TLS task. In this paper, we compare different TLS strategies using appropriate evaluation frameworks, and propose a simple and effective combination of methods that improves over the state-of-the-art on all tested benchmarks. For a more robust evaluation, we also present a new TLS dataset, which is larger and spans longer time periods than previous datasets. The dataset will be made available at <https://github.com/complementizer/news-tls>.

1 Introduction

Timelines of news events can be useful to condense long-ranging news topics and can help us understand how current major events follow from prior events. Timeline summarization (TLS) aims to automatically create such timelines, i.e., temporally ordered time-stamped textual summaries of events focused on a particular topic. While TLS has been studied before, most works treat it as a combination of two individual subtasks, 1) date selection and 2) date summarization, and only focus on one of these at a time (Tran et al., 2013a,b, 2015b). However, these subtasks are almost never evaluated in combination, which leaves an unclear picture of how well TLS is being solved in general. Furthermore, previously used evaluation metrics for the date selection and timeline summarization tasks are not appropriate since they do not consider the temporal alignment in the evaluation. Just until recently, there were no established experimental settings and appropriate metrics for the full TLS task (Martschat and Markert, 2017, 2018).

Date	Summary
2001-11-29	Enron could cost Dutch group \$195m
2001-11-30	1,100 UK jobs go in Enron collapse
2001-12-02	Barclays: Enron bankruptcy will not affect business
2002-01-15	As Enron scandal spreads, US starts to question cash for influence culture
	[...]
2004-07-08	Jury indicts Lay for inflating Enron earnings
2006-05-25	Former Enron bosses found guilty
2006-07-05	Enron founder Lay dies
2008-02-22	US prison beckons British bankers who got cosy with Enron

Table 1: Excerpt of an automatically constructed timeline about the company Enron, using article headlines as summaries. The shaded parts indicate that the date or summary matches entries in a corresponding human-written ground-truth timeline.

In this paper, we examine existing strategies for the full TLS task and how well they actually work. We identify three high-level approaches: 1) *Direct summarization* treats TLS like text summarization, e.g., by selecting a small subset of sentences from a massive collection of news articles; 2) The *date-wise* approach first selects salient dates and then summarizes each date; 3) *Event detection* first detects events, e.g., via clustering, selects salient events and summarizes these individually. The current state-of-the-art method is based on direct summarization (Martschat and Markert, 2018). We therefore focus on testing the two remaining strategies, which have not been appropriately evaluated yet and allow for better scalability.

We propose a simple method to improve date summarization for the date-wise approach. The method uses temporal expressions (textual references to dates) to derive date vectors, which in turn help to filter candidate sentences to summarize particular dates. With this modification, the date-wise approach obtains improved state-of-the-

art results on all tested datasets. We also propose an event-based approach via clustering, which outperforms (Martschat and Markert, 2018) on one of three tested datasets. We use purpose-built evaluation metrics for evaluating timelines introduced by Martschat and Markert (2017). For a more robust evaluation, we also present a new dataset for TLS, which is significantly larger than previous datasets in terms of the number of individual topics and time span.

We summarize our contributions as follows:

1. We compare different TLS strategies side-by-side using suitable evaluation metrics to provide a better picture for how well the full TLS task for news is solved so far.
2. We propose a simple addition to existing methods to significantly improve date-wise TLS, achieving new state-of-the-art results.
3. We present a new TLS dataset that is larger than previous datasets and spans longer time ranges (decades of news timelines).

2 Related Work

Timeline summarization for news articles has received some attention in the last two decades (Swan and Allan, 2000; Allan et al., 2001; Chieu and Lee, 2004; Yan et al., 2011a,b; Kessler et al., 2012; Tran et al., 2013a,b; Li and Li, 2013; Tran et al., 2015a,b; Wang et al., 2015, 2016; Martschat and Markert, 2017, 2018; Steen and Markert, 2019). The task is commonly split into date selection and date summarization subtasks.

2.1 Date Selection

Supervised machine learning has been proposed to predict whether dates appear in ground-truth timelines (Kessler et al., 2012; Tran et al., 2013a). Tran et al. (2015b) use graph-based ranking of dates, which is reported to outperform supervised methods¹.

2.2 Date Summarization

Several approaches construct date summaries by picking sentences from ranked lists. The ranking is based on regression or learning-to-rank to predict ROUGE scores between the sentence and a ground-truth summary (Tran et al., 2013a,b). Tran

¹Despite our best efforts, we could neither obtain code for this method from the authors nor reproduce its reported performance, and therefore did not include it in our experiments.

et al. (2015a) observe that users prefer summaries consisting of headlines to summaries consisting of sentences from article bodies. Steen and Markert (2019) propose abstractive date summarization based on graph-based sentence merging and compression. Other works propose the use of additional data, such as comments on social media (Wang et al., 2015), or images (Wang et al., 2016).

2.3 Full Timeline Summarization

Chieu and Lee (2004) produce timelines by ranking sentences from an entire document collection. The ranking is based on summed up similarities to other sentences in an n -day window. Nguyen et al. (2014) propose a pipeline to generate timelines consisting of date selection, sentence clustering, and ranking. Martschat and Markert (2018) adapt submodular function optimization, commonly used for multi-document summarization, for the TLS task. The approach searches for a combination of sentences from a whole document collection to construct a timeline and is the current state-of-the-art for full TLS. Steen and Markert (2019) use a two-stage approach consisting of date selection and date summarization to build timelines. Other examples of automatic timeline generation can be found in the social media-related literature, where microblogs are often clustered before being summarized (Wang et al., 2014; Li and Cardie, 2014). We explore a similar framework for evaluating clustering-based TLS.

3 Strategies for Timeline Summarization

Problem Definition

We define the TLS setup and task as follows. Given is a set of news articles A , a set of query keyphrases Q , and a ground-truth (reference) timeline r , with l dates that are associated with k sentences on average, i.e., $m = k * l$ sentences in total. The task is to construct a (system) timeline s that contains m sentences, assigned to an arbitrary number of dates. A simpler and stricter setting can also be used, in which s must contain exactly l dates with k sentences each.

Approach Types

A number of different high-level approaches can be used to tackle this task:

1. **Direct Summarization:** A is treated as one set of sentences, from which a timeline is directly extracted, e.g., by optimizing a sentence




Figure 1: Counts of published articles and textual mentions across dates in an article collection about Enron.

combination (Martschat and Markert, 2018), or by sentence ranking (Chieu and Lee, 2004). Among these, Martschat and Markert (2018)’s solution for the full TLS task has state-of-the-art accuracy but does not scale well.

2. **Date-wise Approach:** This approach selects l dates and then constructs a text summary of k sentences on average for each date.
3. **Event Detection:** This approach first detects events in A , e.g., by clustering similar articles, and then identifies the l most important events and summarizes these separately.

Since no prior work has analyzed the latter two categories for the full TLS task, we discuss and develop such approaches next.

3.1 Date-wise Approach

The approach described here mostly consists of existing building blocks, with a few but important modifications proposed from our side.

Defining the Set of Dates

First, we identify the set of possible dates to include in a timeline. We obtain these from (i) the publication dates of all articles in A and (ii) textual references of dates in sentences in A , such as ‘last Monday’, or ‘12 April’. We use the tool HeidelbergTime² (Strötgen and Gertz, 2013) to detect and resolve textual mentions of dates.

²<https://github.com/HeidelbergTime/heideltime>

Date Selection

Next, we select the l most important dates. We compare the following date selection methods introduced by Tran et al. (2013a):

- **PUBCOUNT:** Ranking dates by the number of articles published on a date.
- **MENTIONCOUNT:** Ranking dates by the number of sentences that mention the date.
- **SUPERVISED:** Extracting date features and using classification or regression to predict whether a date appears in a ground-truth timeline. These features mostly include the publication count and different variants of counting date mentions.

Our experiments show that SUPERVISED works best, closely followed by MENTIONCOUNT (Appendix A.1). Figure 1 shows an example of publication and date mention counts and ground-truth dates over time. Two challenges are evident that date selection methods face: 1) These count signals usually do not perfectly correlate with ground-truth dates, and 2) high values often cluster around important dates, i.e., a ‘correct’ date is often surrounded by other, ‘incorrect’ dates with similarly strong signals.

Candidate Sentences for Dates

To summarize a particular date d , we first need to decide which articles or sentences we use as a source to create a summary from. Previous research has not explored this aspect much due to the separated treatment of subtasks. We propose a simple but effective heuristic to do this. We consider the following two sets to be the primary source of suitable candidate sentences:

- P_d : **Sentences published on or closely after d .** These often contain initial reports of events occurring on d .
- M_d : **Sentences that mention d .** These sentences are from articles published at any point in time, and may retrospectively refer to d , or announce events on d beforehand³.

We evaluate these two options in our experiments, and propose an heuristic that combines these, which we call PM-MEAN. We aim to

³In practice, we include the first 5 sentences in the body of each article published on d and up to 2 days after d into P_d . We include all sentences found in A that mention d into M_d .

find a subset of sentences in $P_d \cup M_d$ that are likely to mention important events happening on d . We convert all the sentences in the collection A to sparse bag-of-words (unigram) vectors with sentence-level TF-IDF weighting. We represent the sets of sentences P_d and M_d using the mean of their respective sentence vectors, \bar{x}_{P_d} and \bar{x}_{M_d} . The core assumption of the method is that the content shared between P_d and M_d is a good source for summarizing events on d . To capture this content, we build a *date vector* x_d , so that we can compare sentence vectors against it to rank sentences. We set the value of x_d for each dimension i in the feature space as follows:

$$x_d^i = \begin{cases} \frac{1}{|P_d|}\bar{x}_{P_d}^i + \frac{1}{|M_d|}\bar{x}_{M_d}^i & \text{if } \bar{x}_{P_d}^i > 0 \text{ and } \bar{x}_{M_d}^i > 0 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Thus the date vector x_d is an average of \bar{x}_{P_d} and \bar{x}_{M_d} weighted by the sizes of P_d and M_d , with any features zeroed out if they are missing in either P_d or M_d . To rank sentences, we compute the cosine similarity between the vector x_s of each candidate sentence $s \in (P_d \cup M_d)$ to x_d . We select the best-scoring candidate sentences by defining a threshold on this similarity. To avoid tuning this threshold, we use a simple knee point detection method (Satopaa et al., 2011) to dynamically identify a threshold that represents the "knee" (or elbow) in the similarity distribution. This set of best-scoring sentences is then used as the input for the final date summarization step.

Date Summaries

To construct the final timeline, we separately construct a summary for the l highest ranked dates. Prior to our main experiments, we test several multi-document summarization algorithms:

- **TEXTRANK**: Runs PageRank on a graph of pairwise sentences similarities to rank sentences (Mihalcea and Tarau, 2004).
- **CENTROID-RANK**: Ranks sentences by their similarity to the centroid of all sentences (Radev et al., 2004).
- **CENTROID-OPT**: Greedily optimises a summary to be similar to the centroid of all sentences (Ghalandari, 2017).
- **SUBMODULAR**: Greedily optimizes a summary using submodular objective functions that represent coverage and diversity (Lin and Bilmes, 2011).

The only modification to these algorithms in our TLS pipeline is that we prevent sentences not containing any topic keyphrases from query Q to be included in the summary. CENTROID-OPT has the best results (Appendix A.1) and is used in the main experiments.

Timeline Construction

The date-wise approach constructs a timeline as follows: first, rank all potential dates using one of the date selection approaches described, then pick the l highest ranked ones, pick candidate sentences for each date, and summarize each date individually from the according candidate set, using k sentences. We might not be able to summarize a particular date due to the keyword constraint in the summarization step. Whenever this is the case, we skip to the next date in the ranked list, until l is reached.

3.2 Event Detection Approach

When humans are tasked with constructing a timeline, we expect that they reason over important *events* rather than dates. Conceptually, detecting and selecting events might also be more appropriate than selecting dates because multiple events can happen on the same day, and an event can potentially span multiple days.

To explore this, we test a TLS approach based on event detection by means of article clustering. The general approach can be summarized as follows: (1) Group articles into clusters; (2) Rank and select the l most important clusters; (3) Construct a summary for each cluster. Similarly to the date-wise approach, this mostly consists of existing building blocks that we adapt for TLS.

Clustering

For each input collection A , we compute sparse TF-IDF unigram bag-of-words vectors for all articles in A . We apply clustering algorithms to these vectors. To cluster articles, we use Markov Clustering (MCL) with a temporal constraint. MCL (Van Dongen, 2000) is a clustering algorithm for graphs, i.e., a community detection algorithm. It is based on simulating random walks along nodes in a graph. Ribeiro et al. (2017) use this approach for clustering news articles.

We convert A into a graph where nodes correspond to articles so that we can cluster the articles using MCL, with the following temporal constraint: Articles a_1, a_2 are assigned an edge if their publication dates are at most 1 day apart from each other.

The edge weight is set to the cosine similarity between the TF-IDF bag-of-words vectors of a_1 and a_2 . The constraint on the publication dates ensures that clusters do not have temporal gaps. Furthermore, it reduces the number of similarity computations between pairs of articles considerably. We run MCL on this graph and obtain clusters by identifying the connected components in the resulting connectivity matrix⁴.

Assigning Dates to Clusters

We define the cluster date as the date that is most frequently mentioned within articles of the cluster. We identify date mentions using the HeidelbergTime tool.

Cluster Ranking

To construct a timeline, we only need the l most important clusters. We obtain these by ranking and retaining the top- l clusters of the ranked list. We test the following scores to rank clusters by:

- **SIZE:** Rank by the numbers of articles in a cluster.
- **DATEMENTIONCOUNT:** Rank by how often the cluster date is mentioned throughout the input collection.
- **REGRESSION:** Rank using a score by a regression model trained to predict importance scores of clusters.

For the regression-based ranking method, we represent clusters using the following features: number of articles in a cluster; number of days between the publication dates of the first and last article in the cluster; maximum count of publication dates of articles within a cluster; maximum mention count of dates mentioned in articles in a cluster; sum of mention counts of dates mentioned in articles in a cluster. We test two approaches to label clusters with target scores to predict.

- **Date-Accuracy:** This is 1 if the cluster date appears in the ground-truth, else 0.
- **ROUGE:** The ROUGE-1 F1-score⁵ between the summary of the cluster and the ground-truth summary of the cluster date. If the cluster date does not appear in the ground-truth, the score is set to 0.

⁴We use the implementation and default parameters from https://github.com/GuyAllard/markov_clustering

⁵ROUGE-1 obtained a better overall performance than ROUGE-2 for this purpose.

We evaluate these different options (Appendix A.2) and observe that ranking by DATEMENTIONCOUNT works better than the supervised methods, showing that predicting the suitability of clusters for timelines is difficult.

Cluster Summarization

We use the same multi-document summarization method that works best for the date-wise approach (CENTROID-OPT).

Timeline Construction

In summary, the clustering approach builds a timeline as follows: 1) cluster all articles, 2) rank clusters, 3) build a summary with k sentences for the top- l clusters, skipping clusters if a summary cannot be constructed due to missing keywords. Furthermore, we skip clusters if the date assigned to the cluster is already "used" by a previously picked cluster. Conceptually, this implies that we can only recognize one event per day. In initial experiments, this leads to better results than alternatives, e.g., allowing multiple summaries of length k per day.

4 Dataset

Tran et al. introduced the 17 Timelines (T17) (Tran et al., 2013a) and the CRISIS (Tran et al., 2015a) datasets for timeline summarization from news articles. However, we see the need for better benchmarks due to 1) a small number of topics in the T17 and CRISIS datasets (9 and 4 topics respectively), and 2) relatively short time span, ranging from a few months to 2 years.

Therefore, we build a new TLS dataset, called ENTITIES, that contains more topics (47) and longer time-ranges per topic, e.g., decades of news articles. In the following, we describe how we obtain ground-truth timelines and input article collections for this dataset.

Ground-Truth Timelines: We obtain ground-truth timelines from CNN Fast Facts⁶, which has a collection of several hundred timelines grouped in categories, e.g., 'people' or 'disasters'. We pick all timelines of the 'people' category and a small number from other categories.

Queries: For each ground-truth timeline, we define a set of query keyphrases Q . By default, we use the original title of the timeline as the keyphrase. For people entities, we use the last token of the title to capture surnames only, which increases the

⁶<http://edition.cnn.com/specials/world/fast-facts>

coverage. We manually inspect the resulting sets of keyphrases and correct these if necessary.

Input Articles: For each entity from the ground-truth timelines, we search for news articles using TheGuardian API⁷. We use this source because it provides access to all published articles starting from 1999. We search for articles that have exact matches of the queries in the article body. The timespan for the article search is set so that it extends the ground-truth timeline by 10% of its days before its first and after its last date.

Adjustments and Filtering: The ground-truth timelines are modified to be usable for TLS and to ensure they do not contain data not present in the document collection:

- We remove entries in the ground-truth timelines if they do not specify year, month, and day of an event.
- Ground-truth timelines are truncated to the first and last date of the input articles.
- Entries in the ground-truth timeline are removed if there is no input article published within ± 2 days.

Afterwards, we remove all topics from the dataset that do not fulfill the following criteria:

- The timeline must have at least 5 entries.
- For at least 50% of the dates present in the ground-truth timeline, textual references have to be found in the article collection (e.g., 'on Wednesday' or 'on 1 August'). This is done to ensure that the content of the timelines is reflected to some degree in the article collection.
- There are at least 100 and less than 3000 articles containing the timeline-entity in the input articles. This is done to reduce the running time of experiments.

Dataset Characteristics: Tables 2 and 3 give an overview of properties of the two existing datasets and our new dataset, and mostly show averaged values over tasks in a dataset. An individual task corresponds to one ground-truth timeline that a TLS algorithm aims to simulate. $\#PubDates$ refers to the number of days in an article collection A on which any articles are published. The compression ratio w.r.t. sentences ("comp. ratio (sents)") is m divided by the total number of sentences in A ,

⁷<http://open-platform.theguardian.com/>

and the compression ratio w.r.t dates is l divided by $\#PubDates$. "Avg. date cov" refers to the average coverage of dates in the ground-truth timeline r by the articles in A . This can be counted by using publication dates in A , ("published"), or by textual date references to dates within articles in A ("mentioned"). The fact that there are generally more ground-truth dates covered in textual date references compared to publication dates suggests making use of these date mentions.

T17 has longer (l), and more detailed (k) timelines than the other datasets, CRISIS has more articles per task, and ENTITIES has more topics, publication dates and longer time periods per task.

5 Experiments

5.1 Evaluation Metrics

In our experiments, we measure the quality of generated timelines with the following two evaluation metrics, which are also used by Martschat and Markert (2018):

- **Alignment-based ROUGE F1-score:** This metric compares the textual overlap between a system and a ground-truth timeline, while also considering the assignments of dates to texts.
- **Date F1-score:** This metric compares only the dates of a system and a ground-truth timeline.

We denote the alignment-based ROUGE-1 F1-score as AR1-F and Date F1-score as Date-F1.

5.2 Experimental Settings

Concerning the datasets and task, we follow the experimental settings of Martschat and Markert (2018):

- Each dataset is divided into multiple topics, each having at least one ground-truth timeline. If a topic has multiple ground-truth timelines, we split the topic into multiple tasks. The final results in the evaluation are based on averages over tasks/ground-truth timelines, not over topics.
- Each task includes a set of news articles A , a set of keyphrases Q , a ground-truth timeline r , with number of dates (length) l , average number of summary sentences per date k , and total number of summary sentences $m = l * k$.

Table 2: Dataset Statistics for the TLS task (i)

Dataset	#Topics	#TLs	Avg. #Docs	Avg. #Sents	Avg. #PubDates	Avg. Duration (in days)
T17	9	19	508	20409	124	212
CRISIS	4	22	2310	82761	307	343
ENTITIES	47	47	959	31545	600	4437 (\approx 12 years)

Table 3: Dataset Statistics for the TLS task (ii)

Dataset	Avg. l	Avg. k	Avg. m	Avg. comp. ratio (sents)	Avg. comp. ratio (dates)	Avg. date cov. (published)	Avg. date cov. (mentioned)
T17	36	2.9	108	0.0117	0.43	81%	93%
CRISIS	29	1.3	38	0.0005	0.11	90%	96%
ENTITIES	23	1.2	26	0.0017	0.06	51%	65%

- In each task, we remove all articles from A whose publication dates are outside of the range of dates of the ground-truth timeline r of the task. Article headlines are not used.
- We run leave-one-out cross-validation over all tasks of a dataset.
- We test for significant differences using an approximate randomization test (Marcus et al., 1993) with a p-value of 0.05.

We use the following configurations for our methods:

- A stricter and simpler version of the output size constraint: We produce timelines with the number of dates l and k sentences per date.
- In the summarization step of our methods, we only allow a sentence to be part of a summary if it contains any keyphrase in Q . As opposed to Martschat and Markert (2018), we still keep sentences not matching Q , e.g., for TF-IDF computation, clustering, and computing date vectors.

5.3 Methods Evaluated

We compare the following types of methods to address the full news TLS task.

Direct summarization approaches:

- CHIEU2004: Chieu and Lee (2004) An unsupervised baseline based on direct summarization. We use the reimplementation from Martschat and Markert (2018).
- MARTSCHAT2018: Martschat and Markert (2018) State-of-the-art method on the CRISIS and T17 datasets. It greedily selects a combination of sentences from the entire collection

A maximizing submodular functions for content coverage, textual and temporal diversity, and a high count of date references⁸.

Date-wise approaches:

- TRAN 2013 (Tran et al., 2013a): The original date-wise approach, using regression for both date selection and summarization, and using all sentences of a date as candidate sentences.
- PUBCOUNT: A simple date-wise baseline that uses the publication count to rank dates, and all sentences published on a date for candidate selection. We use CENTROID-OPT for summarization.
- DATEWISE: Our date-wise approach after testing different building blocks (see Appendix A.1). It uses supervised date selection, PM-MEAN for candidate selection and CENTROID-OPT for summarization.

Event detection approach based on clustering:

- CLUST: We use DATEMENTIONCOUNT to rank clusters, and CENTROID-OPT for summarization, which are the best options according to our tests (see Appendix A.2).

Note that all methods apart from DATEWISE and CLUST have been proposed previously.

Oracles:

To interpret the alignment-based ROUGE scores better and to approximate their upper bounds, we measure the performance of three different oracle methods:

⁸Multiple variants of this approach were introduced in the paper. We picked the variant called "AsMDS+ $f_{TempDiv}$ + $f_{DateRef}$ " due to its good results.

- **DATE ORACLE:** Selects the correct (ground-truth) dates and uses CENTROID-OPT for date summarization.
- **TEXT ORACLE:** Uses regression to select dates, and then constructs a summary for each date by optimizing the ROUGE to the ground-truth summaries.
- **FULL ORACLE:** Selects the correct dates and constructs a summary for each date by optimizing the ROUGE to the ground-truth summaries.

We give more detail about these in Appendix A.3.

5.4 Results

Table 4 shows the final evaluation results. We reproduced the results of CHIEU2004 and MARTSCHAT2018 reported by Martschat and Markert (2018) using their provided code⁹. The other results are based on our implementations. Table 10 in Appendix A.6 shows several output examples across different methods.

6 Analysis and Discussion

6.1 Performance of TLS Strategies

Among the methods evaluated, DATEWISE consistently outperforms all other methods on all tested datasets in the alignment-based ROUGE metrics. The Date-F1 metric for this method is close to other methods, and not always better, which shows that the advantage of DATEWISE is due to the sentence selection (based on our heuristic date vectors) and summarization. Note that the date selection method is identical to TRAN2013. We conclude from these results that the expensive combinatorial optimization used in MARTSCHAT2018 is not necessary to achieve high accuracy for news TLS.

CLUST performs worse than DATEWISE and MARTSCHAT2018, except on ENTITIES, where it outperforms MARTSCHAT2018. We find that for the other two datasets, CLUST often merges articles from close dates together that would belong to separate events on ground-truth timelines, which may suggest that a different granularity of clusters is required depending on the task.

DATE ORACLE and FULL ORACLE should theoretically have a 100% Date-F1. In practice, their Date-F1 scores turn out lower because, for some dates, no candidate sentences that match query Q

⁹With the exception of CRISIS due to memory issues.

can be found, which causes the dates to be omitted from the oracle timelines.

Based on the performance of different systems, the hardest dataset is ENTITIES, followed by CRISIS.

6.2 What makes TLS difficult?

While the ranking of methods is fairly stable, the performance of all methods varies a lot across the datasets and across individual tasks within datasets. To find out what makes individual tasks difficult, we measure the Spearman correlation between AR1-F and several dataset statistics. The details are included in Appendix A.5. The correlations show that a high number of articles and publication dates and a low compression ratio w.r.t to dates generally decreases performance. This implies that highly popular topics are harder to summarize. The duration of a topic also corresponds to lower performance, but in a less consistent pattern.

The generally low performance across tasks and methods is likely influenced by the following factors:

- The decision for human editors to include particular events in a timeline and to summarise these in a particular way can be highly subjective. Due to the two-stage nature of TLS, this problem is amplified in comparison to regular text summarization.
- Article collections can be insufficient to cover every important event of a topic, e.g., due to the specific set of news sources or the search technique used.

6.3 Running Time

DATEWISE and CLUST are up to an order of magnitude faster to run than MARTSCHAT2018 (Appendix A.4) since their date summarization steps only involve a small subset of sentences in an article collection.

6.4 Adjacent Dates and Redundancy

Automatically constructed timelines often contain a high amount of multiple adjacent dates, while this is not the case in ground-truth timelines. Summaries of such adjacent dates often tend to refer to the same event and introduce redundancy into a timeline. To quantify this, we count the proportion of those "date bigrams" in a chronologically ordered timeline, which are only 1 day apart. The results (see Table 5) show that this is an issue

	T17 Dataset			CRISIS Dataset			ENTITIES Dataset		
	AR1-F	AR2-F	Date-F1	AR1-F	AR2-F	Date-F1	AR1-F	AR2-F	Date-F1
Text Oracle	0.198	0.073	0.541	0.136	0.052	0.297	0.069	0.023	0.20
Date Oracle	0.179	0.057	0.926	0.202	0.063	0.974	0.17	0.047	0.757
Full Oracle	0.312	0.128	0.926	0.367	0.15	0.974	0.232	0.075	0.757
CHIEU2004	0.066	0.019	0.251	0.052	0.012	0.142	0.036	0.01	0.102
MARTSCHAT2018	0.105	0.03	0.544 ●	0.075 ●	0.016	0.281	0.042	0.009	0.167
TRAN2013	0.094	0.022	0.517 ●	0.054	0.011	0.289	0.042	0.012	0.184 †
PUBCOUNT	0.105	0.027	0.481	0.067	0.012	0.233	0.033	0.009	0.107
DATEWISE	0.12 * ● †	0.035 * ●	0.544 * ●	0.089 * ●	0.026 * ●	0.295 ●	0.057 * ● †	0.017 * ● †	0.205 * ● †
CLUST	0.082	0.020	0.407	0.061	0.013	0.226	0.051 †	0.015 †	0.174
DATEWISE (titles)	-	-	-	0.072	0.016	0.287	0.057	0.017	0.194

Table 4: Results on the full TLS task. * indicates a significant improvement over Tran 2013, ● over CLUST, and † over MARTSCHAT2018. DATEWISE (titles) is not included in the significance testing.

	Ground-truth	MARTSCHAT2018	DATEWISE	CLUST
T17	0.45	0.63	0.62	0.25
CRISIS	0.18	-	0.52	0.06
ENTITIES	0.03	0.18	0.3	0.05

Table 5: Proportion of adjacent dates of timelines produced by different methods, and the ground-truth timelines.

for MARTSCHAT2018 and DATEWISE, but less so for CLUST, which is designed to avoid this behavior. Note that MARTSCHAT2018 includes an objective function to reward diversity within a timeline, while DATEWISE has no explicit mechanism against redundancy among separate dates. Interestingly, when forcing DATEWISE to avoid selecting adjacent dates (by skipping such dates in the ranked list), the performance in all metrics decreases. In this case, high redundancy is a safer strategy for optimizing TLS metrics compared to enforcing a more balanced spread over time. Because of such effects, we advise to use automated evaluation metrics for TLS with care and to conduct qualitative analysis and user studies where possible.

6.5 Use of Titles

While using article titles can make timelines more readable and understandable (Tran et al., 2015a), we do not involve titles in our main experiments, in order to directly compare to MARTSCHAT2018, and due to the lack of titles in T17. The last row in Table 4 shows the results of a separate experiment with DATEWISE in which we build date summaries using titles only. Using only titles generally increases AR Precision at the cost of Recall. AR-F is negatively affected in CRISIS but does not change in ENTITIES. Figure 1 shows parts of a title-based timeline produced by DATEWISE.

7 Conclusion

In this study, we have compared and proposed different strategies to construct timeline summaries of long-ranging news topics: the previous state-of-the-art method based on direct summarization, a date-wise approach, and a clustering-based approach. By exploiting temporal expressions, we have improved the date-wise approach and yielded new state-of-the-art results on all tested datasets. Hence, we showed that an expensive combinatorial search over all sentences in a document collection is not necessary to achieve good results for news TLS. For a more robust and diverse evaluation, we have constructed a new TLS dataset with a much larger number of topics and with longer time-spans than in previous datasets. Most of the generated timelines are still far from oracle timeline extractors and leave large gaps for improvements. Potential future directions include a more principled use of our proposed heuristic for detecting content relevant to specific dates, the use of abstractive techniques, a more effective treatment of the redundancy challenge, and extending the new dataset with multiple sources.

Acknowledgments

This work was funded by the Irish Research Council (IRC) under grant number EBPPG/2018/23, the Science Foundation Ireland (SFI) under grant number 12/RC/2289_P2 and the enterprise partner Aylien Ltd.

References

James Allan, Rahul Gupta, and Vikas Khandelwal. 2001. Temporal summaries of new topics. In *Proceedings of the 24th annual international ACM SIGIR conference on Research and development in in-*

- formation retrieval. *ACM*, 10–18.
- Hai Leong Chieu and Yoong Keok Lee. 2004. Query based event extraction along a timeline. In *Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval*. *ACM*, 425–432.
- Demian Gholipour Ghalandari. 2017. Revisiting the Centroid-based Method: A Strong Baseline for Multi-Document Summarization. In *Proceedings of the Workshop on New Frontiers in Summarization*. 85–90.
- Remy Kessler, Xavier Tannier, Caroline Hagege, Véronique Moriceau, and André Bittar. 2012. Finding salient dates for building thematic timelines. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*. Association for Computational Linguistics, 730–739.
- Jiwei Li and Claire Cardie. 2014. Timeline generation: Tracking individuals on twitter. In *Proceedings of the 23rd international conference on World wide web*. *ACM*, 643–652.
- Jiwei Li and Sujian Li. 2013. Evolutionary hierarchical dirichlet process for timeline summarization. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 556–560.
- Hui Lin and Jeff Bilmes. 2011. A class of submodular functions for document summarization. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*. Association for Computational Linguistics, 510–520.
- Mitchell P Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of English: The Penn Treebank. *Computational linguistics* 19, 2 (1993), 313–330.
- Sebastian Martschat and Katja Markert. 2017. Improving rouge for timeline summarization. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, Vol. 2. 285–290.
- Sebastian Martschat and Katja Markert. 2018. A Temporally Sensitive Submodularity Framework for Timeline Summarization. In *Proceedings of the 22nd Conference on Computational Natural Language Learning*. 230–240.
- Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*. 404–411.
- Kiem-Hieu Nguyen, Xavier Tannier, and Véronique Moriceau. 2014. Ranking multidocument event descriptions for building thematic timelines. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*. 1208–1217.
- Dragomir R Radev, Hongyan Jing, Małgorzata Styś, and Daniel Tam. 2004. Centroid-based summarization of multiple documents. *Information Processing & Management* 40, 6 (2004), 919–938.
- Swen Ribeiro, Olivier Ferret, and Xavier Tannier. 2017. Unsupervised Event Clustering and Aggregation from Newswire and Web Articles. In *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*. 62–67.
- Ville Satopaa, Jeannie Albrecht, David Irwin, and Barath Raghavan. 2011. Finding a” kneedle” in a haystack: Detecting knee points in system behavior. In *2011 31st International Conference on Distributed Computing Systems Workshops*. IEEE, 166–171.
- Julius Steen and Katja Markert. 2019. Abstractive Timeline Summarization. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*. 21–31.
- Jannik Strötgen and Michael Gertz. 2013. Multilingual and Cross-domain Temporal Tagging. *Language Resources and Evaluation* 47, 2 (2013), 269–298. <https://doi.org/10.1007/s10579-012-9179-y>
- Russell Swan and James Allan. 2000. Automatic generation of overview timelines. In *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval*. *ACM*, 49–56.
- Giang Tran, Mohammad Alrifai, and Eelco Herder. 2015a. Timeline summarization from relevant headlines. In *European Conference on Information Retrieval*. Springer, 245–256.
- Giang Tran, Eelco Herder, and Katja Markert. 2015b. Joint graphical models for date selection in timeline summarization. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Vol. 1. 1598–1607.
- Giang Bihn Tran, Mohammad Alrifai, and Dat Quoc Nguyen. 2013a. Predicting relevant news events for timeline summaries. In *Proceedings of the 22nd International Conference on World Wide Web*. *ACM*, 91–92.
- Giang Binh Tran, Tuan A Tran, Nam-Khanh Tran, Mohammad Alrifai, and Nattiya Kanhabua. 2013b. Leveraging learning to rank in an optimization framework for timeline summarization. In *SIGIR 2013 Workshop on Time-aware Information Access (TAIA)*.
- Stijn Marinus Van Dongen. 2000. *Graph clustering by flow simulation*. Ph.D. Dissertation.

Lu Wang, Claire Cardie, and Galen Marchetti. 2015. Socially-Informed Timeline Generation for Complex Events. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.

William Yang Wang, Yashar Mehdad, Dragomir R Radev, and Amanda Stent. 2016. A Low-Rank Approximation Approach to Learning Joint Embeddings of News Stories and Images for Timeline Summarization. In *HLT-NAACL*. 58–68.

Zhenhua Wang, Lidan Shou, Ke Chen, Gang Chen, and Sharad Mehrotra. 2014. On summarization and timeline generation for evolutionary tweet streams. *IEEE Transactions on Knowledge and Data Engineering* 27, 5 (2014), 1301–1315.

Rui Yan, Liang Kong, Congrui Huang, Xiaojun Wan, Xiaoming Li, and Yan Zhang. 2011a. Timeline generation through evolutionary trans-temporal summarization. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 433–443.

Rui Yan, Xiaojun Wan, Jahna Otterbacher, Liang Kong, Xiaoming Li, and Yan Zhang. 2011b. Evolutionary timeline summarization: a balanced optimization framework via iterative substitution. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*. ACM, 745–754.

A Appendices

A.1 Testing Variants of DATEWISE

Table 6 shows results for different variants of the date-wise approach.

Date Selection: While testing different date selection methods, we use PM-MEAN for candidate selection and CENTROID-OPT for summarization. The supervised date selection methods work best, closely followed by MENTIONCOUNT.

Candidate Sentence Selection: We compare different strategies of defining the set of sentences associated with a date prior to summarization. The results show that the PM-method can improve the performance, especially for the Crisis dataset.

Date Summarization: Finally, we test different unsupervised text summarization algorithms to summarize each selected date. CENTROID-OPT works best and is used in our main experiments.

A.2 Testing Variants of CLUST

For the clustering-based TLS approach, we only test different options for ranking clusters. For the summarization step, we use CENTROID-OPT, which works best for the date-wise approach. Table

Date Selector	T17	Crisis	Entities
PubCount	0.49	0.243	0.135
MentCount	0.528	0.295	0.159
Tran 2013 (Reg)	0.535	0.297	0.191
Tran 2013 (Clf)	0.541	0.295	0.172
Candidate Selector	T17	Crisis	Entities
Sents mentioning d	0.11	0.077	0.041
Sents published on d to d + 2 (first 5)	0.112	0.078	0.045
Sents published on d to d + 2 (all)	0.113	0.079	0.041
PM-MEAN	0.118	0.089	0.047
Summarizer	T17	Crisis	Entities
TEXT-RANK	0.113	0.086	0.046
CENTROID-RANK	0.112	0.085	0.046
CENTROID-OPT	0.118	0.089	0.047
SUBMODULAR	0.116	0.088	0.047

Table 6: Variants of Date-wise TLS and their alignment-based ROUGE-1 score.

7 shows somewhat inconsistent results, but overall DATEMENTIONCOUNT obtains the best performance in terms of alignment-based ROUGE.

Cluster Ranking (AR1-F)	T17	Crisis	Entities
SIZE	0.08	0.06	0.048
DATEMENTIONCOUNT	0.081	0.061	0.051
REGRESSION (DATES)	0.08	0.055	0.045
REGRESSION (ROUGE)	0.082	0.055	0.048
Cluster Ranking (Date-F1)	T17	Crisis	Entities
SIZE	0.41	0.22	0.16
DATEMENTIONCOUNT	0.41	0.23	0.15
REGRESSION (DATES)	0.46	0.23	0.16
REGRESSION (ROUGE)	0.44	0.24	0.16

Table 7: Variants of Date Clustering-based TLS.

A.3 Oracles

For the text and full oracles, we use Algorithm 1 for constructing a summary for a date, using ROUGE-1 F1-score as the objective. We include all sentences that mention d , as well as the first 5 sentences of all articles published between d and $d + 5$ days, as candidate sentences for the oracles to summarize dates.

Algorithm 1: Greedy summarization oracle.

```

1 Input: Candidate sentences  $C$ , reference summary  $R$ ,
  summary length  $k$ 
2 Output: Summary sentences  $S$ 
3  $S \leftarrow \{\}$ 
4 while  $|S| < k$  and  $|C| > 0$  do
5    $s^* \leftarrow \arg \max_{s \in C} ROUGE(S, R)$ 
6    $S \leftarrow S \cup \{s^*\}$ 
7    $C \leftarrow C \setminus \{s^*\}$ 
8 Return  $S$ 

```

Dataset	Method	l	k	#articles	#dates	Comp. ratio (sents)	Comp. ratio (dates)	duration
T17	MARTSCHAT 2018	-0.116	0.381	-0.421	-0.586 *	0.298	0.57 *	-0.376
T17	DATEWISE	-0.196	0.616 *	-0.354	-0.714 *	0.319	0.638 *	-0.577 *
T17	CLUST	0.283	0.429	-0.247	-0.411	0.504 *	0.576 *	-0.197
Crisis	DATEWISE	0.19	-0.3	0.393	0.096	-0.271	0.147	0.028
Crisis	CLUST	-0.087	0.038	-0.184	0.029	0.061	-0.037	0.013
Entities	MARTSCHAT 2018	-0.05	-0.012	-0.657 *	-0.682 *	0.644 *	0.649 *	-0.338 *
Entities	DATEWISE	-0.038	-0.056	-0.394 *	-0.406 *	0.348 *	0.39 *	-0.103
Entities	CLUST	0.028	-0.044	-0.461 *	-0.501 *	0.422 *	0.515 *	-0.358 *

Table 8: Correlations between Task Properties and Method Performance.

A.4 Running Time

In Table 9 we compare the running time of DATEWISE and MARTSCHAT2018 on the T17 and ENTITIES datasets¹⁰. The implementations of both our methods and of MARTSCHAT2018 make use of parallel computation to obtain pairwise similarities between sentences or documents where required. We do not parallelize our methods in any other way. We could not run MARTSCHAT2018 on the CRISIS dataset since it requires too much memory, which demonstrates the need for more scalable state-of-the-art methods. DATEWISE and CLUST are considerably faster on both datasets, due to their "divide-and-conquer" nature: The summarization step is applied to only l smaller portions of articles and sentences, instead of the entire set. Note that part of the time is required to run the evaluation tool to compute alignment-based ROUGE.

Dataset	Method	Avg. seconds per topic
T17	MARTSCHAT2018	176
T17	DATEWISE	16
T17	CLUST	15.9
ENTITIES	MARTSCHAT2018	106.3
ENTITIES	DATEWISE	29.5
ENTITIES	CLUST	34.7

Table 9: Running time comparison between current state-of-the-art method MARTSCHAT2018 and the methods we implemented.

A.5 Correlations between Performance and Dataset Characteristics

Detailed results of correlations between different methods and different dataset characteristics are shown in Table 8.

¹⁰On a machine with 16 3.70GHz Intel CPUs and 32GB memory.

A.6 Output Examples

Table 10 shows parts of timelines produced by different methods for a selection of dates that all methods have selected. The topics are taken from the ENTITIES dataset. The examples demonstrate different levels of detail in describing particular events.

DATEWISE (titles only)		
Topic	Date	Summary
Steve Jobs	2009-01-14	Apple boss Steve Jobs to take extended leave
Steve Jobs	2011-08-25	Steve Jobs resigns as Apple CEO
Charles Taylor	2010-08-09	Mia Farrow contradicts Naomi Campbell in Charles Taylor trial
Charles Taylor	2012-04-26	Charles Taylor found guilty of abetting Sierra Leone war crimes
DATEWISE		
Topic	Date	Summary
Steve Jobs	2009-01-14	The boss of the Apple computer empire , Steve Jobs , today disclosed that his health problems have become " more complex " , prompting him to take extended leave from his role as chief executive until the end of June .
Steve Jobs	2011-08-25	Steve Apple made a followup statement : Apple 's Board of Directors today announced that Steve Jobs has resigned as Chief Executive Officer , and the Board has named Tim Cook , previously Apple 's Chief Operating Officer , as the company 's new CEO .
Charles Taylor	2010-08-09	Campbell 's former agent Carole White and the actor Mia Farrow – both of whom were present at a dinner hosted by Nelson Mandela and attended by Taylor and Campbell – have given evidence in a court in the Hague today , some of which appeared to contradict the testimony given by the model last week .
Charles Taylor	2012-04-26	On Thursday Charles Taylor , warlord turned president of Liberia , was convicted of aiding and abetting war crimes by the Sierra Leone special court in The Hague .
CLUST		
Topic	Date	Summary
Steve Jobs	2009-01-14	In his message to staff , Jobs said : " Unfortunately , the curiosity over my personal health continues to be a distraction not only for me and my family , but everyone else at Apple as well .
Steve Jobs	2011-08-25	Steve Jobs has resigned as chief executive of Apple .
Charles Taylor	2010-08-09	Campbell said she was told by her former agent Carole White and the actor Mia Farrow that the diamonds came from Taylor , but otherwise she had no idea who sent them .
Charles Taylor	2012-04-26	Today , as they watched Taylor be convicted of aiding and abetting war crimes on all counts , they have seen justice done .
MARTSCHAT2018		
Topic	Date	Summary
Steve Jobs	2009-01-14	The boss of the Apple computer empire , Steve Jobs , today disclosed that his health problems have become " more complex " , prompting him to take extended leave from his role as chief executive until the end of June .
Steve Jobs	2011-08-25	Steve Apple made a followup statement : Apple 's Board of Directors today announced that Steve Jobs has resigned as Chief Executive Officer , and the Board has named Tim Cook , previously Apple 's Chief Operating Officer , as the company 's new CEO .
Charles Taylor	2010-08-09	Farrow denies that she or White told Campbell that the diamonds had come from Taylor .
Charles Taylor	2012-04-26	The first African head of state to be tried in an international court , Taylor will on Thursday hear the verdict of the Special Court for Sierra Leone in his five - year trial on charges of war crimes and crimes against humanity , including murder , rape , sexual slavery and using child soldiers .

Table 10: Partial timelines produced by different methods, for a fixed selection of and topics and dates.