

The Power of Summary-Source Alignments

Ori Ernst^{* 1,2}, Ori Shapira², Aviv Slobodkin^{1,2}, Sharon Adar²,
Mohit Bansal^{2,3}, Jacob Goldberger¹, Ran Levy², and Ido Dagan¹

¹Bar-Ilan University ²Amazon ³UNC Chapel Hill
{oriern, obspp18, lovodkin93}@gmail.com
{sharonaz, ranlevy}@amazon.com
{mbansal}@cs.unc.edu
{jacob.goldberger@, dagan@cs.}biu.ac.il

Abstract

Multi-document summarization (MDS) is a challenging task, often decomposed to subtasks of salience and redundancy detection, followed by text generation. In this context, alignment of corresponding sentences between a reference summary and its source documents has been leveraged to generate training data for some of the component tasks. Yet, this enabling alignment step has usually been applied heuristically on the sentence level on a limited number of subtasks. In this paper, we propose extending the summary-source alignment framework by (1) applying it at the more fine-grained proposition span level, (2) annotating alignment manually in a multi-document setup, and (3) revealing the great potential of summary-source alignments to yield several datasets for at least six different tasks. Specifically, for each of the tasks, we release a manually annotated test set that was derived automatically from the alignment annotation. We also release development and train sets in the same way, but from automatically derived alignments. Using the datasets, each task is demonstrated with baseline models and corresponding evaluation metrics to spur future research on this broad challenge.

1 Introduction

Common information needs are most often satisfied by multiple texts rather than by a single text. Processing multiple texts is a challenging feat due to the wealth of content they possess, as well as the anticipated redundancy of information. To address this challenge, various tasks support multi-text processing needs, such as information selection, consolidation and fusion. A prominent application that aims to respond to user information needs is the multi-document summarization (MDS) task. Inherently, MDS either explicitly or implicitly prescribes

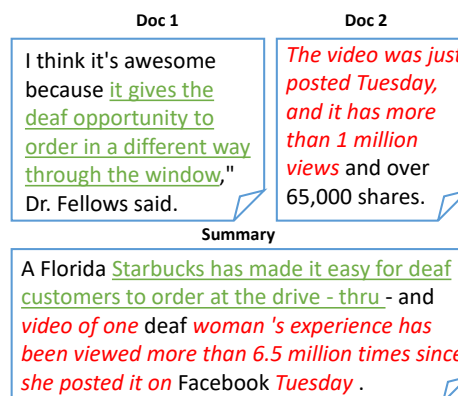


Figure 1: An example of proposition-level multi-document-based alignment. Aligned propositions are in the same color and formatting.

sub-tasks like those listed above (Moryossef et al., 2019; Ernst et al., 2022).

Notably, many works showed the superiority of decomposing a complex task, and more specifically summarization, into its natural subtasks (Moryossef et al., 2019; Ernst et al., 2022; Zhang et al., 2023; Xiao, 2023). However, in most cases there is no training data for each of the subtasks, nor gold test data to evaluate them. That is, while summarization datasets contain gold reference summaries, they do not publish, e.g., gold document salient spans for the *salience detection* task, or gold sentences that fuse information from a few salient document spans for the *fusion* task.

In this paper, we unveil a simple approach to obtain high-quality datasets for a wide variety of multi-text related tasks, via a single annotation process of summary-source alignments. Specifically, we match summary propositions with their proposition-level supporting evidence in the source over an existing MDS dataset (Multi-News; Fabri et al., 2019). An example for alignments is presented in Figure 1.

Aligning all the information segments between a reference summary and its paired document set

* Work was done as an intern at Amazon.

reveals the underlying sub-tasks constructing the summarization process, as illustrated in Figure 2. For instance, the aligned spans in the document set constitute the salient information within them, since they collectively embody the corresponding summary. This characteristic captures a salience detection task (Figure 2 (b)). Similarly, a summary segment can be viewed as a fused version of all its aligned document-set mentions, representing a sentence fusion task (Figure 2 (f)).

Overall, with these alignments we automatically derive datasets for six tasks: (1) salience detection, (2) proposition coreference clustering, (3) evidence detection, (4) text planning, (5) sentence fusion, and (6) in-context passage fusion. In essence, this procedure “reverse engineers” the human summarization process by which the reference summaries were originally created. The resulting data enriches the current inventory of datasets for these individual tasks, while being derived automatically solely from the high-quality alignment annotations. While the tasks addressed here stand on their own merits, such tasks have also been shown to benefit the overall summarization process when used within a pipeline, as mentioned before (more on this in §2).

Our high-quality alignments test set was obtained through a controlled crowdsourcing procedure (Roit et al., 2020), with annotators diligently trained for this task, and contains 100 topics (document-set/summary pairs) with 2256 alignments. We also created large-scale training and development sets by extracting alignments automatically, using the SuperPAL alignment model (Ernst et al., 2021), from the Multi-News train and dev sets (§3). We automatically derive and release train and test datasets from the alignment data for the six mentioned tasks (§4). For each of the tasks, and using its respective dataset, we develop and evaluate two baseline models explicitly targeting the task: one is a trained model while the other is a non-finetuned execution of a ChatGPT LLM (OpenAI, 2023) (§5). Over the six tasks, we generally find that smaller trained models yield better results than the GPT counterpart, leaving room for future advances using our task and dataset suite.

Overall, this work showcases that alignments from an MDS dataset empower a rich collection of multi-document related tasks. These tasks are appealing on their own, and are additionally advan-

tageous as sub-components of MDS solutions.¹

2 Background

Aligning information between source and reference texts has been previously addressed in the realm of summarization. An early effort was conducted for the purpose of summary evaluation through the Pyramid method (Nenkova and Passonneau, 2004). It’s effectiveness, albeit burden of annotation, triggered the pursuit of automatic procedures that mimic the content extraction and alignment (Yang et al., 2016; Gao et al., 2018; Hirao et al., 2018; Zhang and Bansal, 2021), generally via proposition extraction and matching. The alignment approach and manual annotation process for our dataset is reminiscent of the Pyramid method’s, however it is more scalable thanks to controlled crowdsourcing (Roit et al., 2020). While Shapira et al. (2019) also applied crowdsourcing for easing manual alignment, the summary is not exhaustively aligned and high quality is not guaranteed.

Besides summary evaluation, alignments have also been useful for the summarization process itself. To acquire such alignments, the ROUGE metric (Lin, 2004) was leveraged to match between summary and source sentences (Zhang et al., 2018; Cho et al., 2019). This approach was also taken for components of the summarization pipeline, such as detection of salient sentences (Chen and Bansal, 2018a) and sentence fusion (Lebanoff et al., 2019). The heuristic nature of this pairing approach yields noisy alignments since it is both on the sentence level and based on lexical matching.

Many additional works have shown the benefit of decomposed summarization pipelines. Moryossef et al. (2019) and Ernst et al. (2022) split the summarization task to planning and realization phases (corresponding to our derived tasks) showing improved summary outputs. Zhang et al. (2023) find that decomposition of the summarization process can improve faithfulness of the output summary to its source documents. Xiao (2023) shows the benefit of salience detection as an initial phase, both for improving summary quality, and to provide attribution to summary segments.

Information alignment has also been treated as a standalone task. Ernst et al. (2021) designed a supervised model that far exceeds the abilities of lexical-based aligners. They also released a

¹All data is publicly available at <https://github.com/oriern/SPARK>.

high quality test set of sub-sentence-level alignments. Collecting data required expert cleaning of crowdsourced annotations. Our more scalable approach yielded a dataset that is an order of magnitude larger (11 vs 100 topics), and further improves alignment accuracy through refined guidelines and removed UI constraints in the annotation tool (Slobodkin et al., 2022).

Recently, Krishna et al. (2023) released USB, a benchmark for summarization-related tasks, also derived from alignments. Alignment was conducted between the leading section of a Wikipedia article and its body, via controlled crowdsourcing, and required editing text spans in the leading section to remain faithful to the article body. In contrast to their *sentence* level alignments, our *proposition*-level alignments eliminate non-aligning noise. Moreover, aligning in the *multi*-document setting, as opposed the *single*-document in USB, introduces challenges arising from cross-document information sharing and size. The differences listed above induce tasks in our work that are mostly different from those addressed in USB, and that can be treated as standalone tasks.

3 Collecting Alignments Data

Our alignments data consists of a high quality test set, collected through careful manual annotation (§3.1), as well as large-scale training and development sets that were automatically compiled (§3.2). All alignments were extracted from the respective data split of Multi-News (Fabbri et al., 2019), a MDS dataset of sets of news articles with professionally prepared summaries.

An instance in our alignment dataset is based upon a document set D and a corresponding reference summary s . The instance consists of a list of aligned pairs $H = \{(h_1^s, h_1^D), (h_2^s, h_2^D), \dots, (h_n^s, h_n^D)\}$ such that (h_i^s, h_i^D) are proposition spans from the summary and document set, respectively, that describe the same piece of information. Since information is expected to repeat across the documents, H likely contains pairs where $h_i^s = h_j^s$. Moreover, the summary should be exhaustively covered, i.e., all propositions in s are expected to appear in H .

3.1 Manually Annotated Test Set

To manually collect the alignments from a document-set/summary pair (“topic”), we follow the annotation protocol of Slobodkin et al. (2022),

Stat	Train	Dev	Test
# topics	44.5K	5567	100
# alignments	1.5M	186K	2256
# clusters	629K	77.5K	1332
# summary sentences	342K	42K	834
avg. cluster size	2.4	2.4	1.7
avg. # clusters per sent	1.8	1.8	1.6
avg. # clusters per topic	14.1	14.1	13.6
avg. # docs per topic	2.71	2.65	2.97

Table 1: Statistics of our alignments data. The test set is manually collected and the dev/train sets are automatically collected, all from the Multi-News MDS dataset. A cluster contains alignments pertaining to document spans that align to the same summary span (referring to the same information).

using controlled crowdsourcing (Roit et al., 2020),² adapting their method to the multi-document setting and altering the annotation guidelines for our purposes. Our annotation yields 2256 alignments from 100 topics. Full statistics are in Table 1.

Annotation interface and procedure. We adopt the web-based annotation tool from Slobodkin et al. (2022), and deploy it on Mechanical Turk³ for crowdsourcing (see Figure 3 in Appendix). The tool shows documents and the corresponding summary side-by-side, and annotators identify matching text segments between a document and the summary. The annotator is instructed to concentrate on an individual summary statement at a time, and to eventually cover the full summary. Each topic is annotated by a single trained annotator. For quality assurance, submissions were randomly checked and direct feedback was given as needed.

Annotation guidelines. A proposition in a summary is determined by a central event (predicate) with its associated arguments. Annotators are guided on how to identify these propositions, including special cases of nested propositions (a proposition being an argument of another), predicates connected via discourse markers, prospective vs. transpired events, and incontiguous propositions. See Appendix C for explanations and details.

²Potential annotators were first picked through a filtering crowdsourcing task, and then went through several increasingly challenging stages of alignment annotations for quality assessment and training on the task. Eventually five annotators were qualified and completed the tasks.

³www.mturk.com

Upon identifying an event in the summary, the annotator locates all aligning spans in the document. A span is defined as the minimal token-set that fully covers the summary proposition without including additional information. A document span may explicitly refer to the summary event, or entail it (the summary event might generalize several instances of coreferring document events).

Inter-annotator agreement. To assess the quality of alignments we measured inter-annotator agreement on a set of instances annotated by all five annotators. A total of 31 summary sentences were annotated against their respective document sets. For every pair of annotators we computed intersection-over-union (*IoU*) of the token indices (only content words) in the document spans that align to the same summary sentence, akin to Ernst et al. (2021). Over the 310 compared pairs (31 sentences with 5 workers), the resulting *IoU* score is 0.717, suggesting high-quality annotation.

3.2 Automatically Collected Data

For training and development sets, we extracted alignments automatically, using SuperPAL (Ernst et al., 2021), from the Multi-News train and dev sets, respectively. Document propositions clustered to the same summary proposition were then clustered together. The train/dev sets have 1.5M/186K alignments from 44.5K/5.5K topics (see Table 1).

4 The Task Suite

Out of the summary-source alignments, annotated manually or automatically, we derive six new datasets for six different tasks, as elaborated below. The tasks are illustrated in Figure 2 and an example topic from our manual dataset is presented in Appendix F. We denote this data suite as “SPARK”, for Summary Proposition Alignment for Reconstructive Knowledgebases.

4.1 Saliency Detection

Saliency detection is the task of marking the important spans within a given source text. It mainly addresses the need within summarization to extract the information around which to summarize the source text (Arumae et al., 2019), either extractively (e.g., Mao et al., 2020) or abstractively (e.g., Chen and Bansal, 2018b). Nevertheless, it can be used as a means to merely highlight central parts of a text for easing on a reader (Self et al., 2013; Sándor and Vorndran, 2014; Ponce et al., 2022).

Task definition and dataset derivation. Given document set D , the task is to mark the spans in D that globally represent the essential information required to obtain a high level overview of D . From our alignments data, this translates to detecting the spans H , i.e., those spans in D that align to the corresponding reference summary s (Figure 2 (b)). Since s is presumably a good portrayal of an overview of D , the spans of H should indeed cover the appropriate information. While the amount and preference of salient information are factors that could be taken into consideration for this task, we rely on the choices of the expert summarizers in the underlying MDS dataset.

From Table 1, we can infer that there are 100 instances for the task (topics) in the test set, and each instance has an average of 22.6 expected spans to identify in the document set.

Evaluation. For evaluation we followed (Tjong Kim Sang and Buchholz, 2000) and used F_1 score on the token-level.

4.2 Proposition Clustering

Information is expressed differently across sources, and this is especially the case in sets of documents on a related topic, as in our setting. When given a list of propositions, grouping together redundant paraphrastic units is a basic need for gathering and organizing content. In summarization and related contexts, redundancy clustering supports generating non-redundant texts that merge overlapping complementary pieces of information. Furthermore, repetition of information typically provides an indication for its importance (Wan and Yang, 2008; Cai et al., 2010; Zhang et al., 2015).

In the broad context of paraphrasing, prior datasets generally address paraphrase pairing (Dolan and Brockett, 2005; Ganitkevitch et al., 2013), while our dataset presents the vaster challenge of paraphrase clustering. For *short* text clustering, prior datasets cluster topically related instances, rather than paraphrastic ones (Phan et al., 2008; Xu et al., 2017; Cohen et al., 2022). Finally, we suggest that paraphrastic matching is better captured at the proposition level, rather than at the sentence level, as a mechanism to prevent misalignment of information.

Task definition and dataset derivation. Given a set of proposition-style text units, this task requires producing non-overlapping clusters of units, such that a cluster contains texts that express the




Figure 2: Deriving SPARK task datasets from our alignments, for a given document set (topic): (a) **Alignments** - aligned summary-source propositions are marked here by the same color; (b) **Saliency Detection** - all aligned document propositions are to be selected; (c) **Proposition Clustering** - document propositions aligned with the same summary proposition are to be clustered; (d) **Evidence Detection** - a summary proposition is the input query, and the document propositions aligned with it are to be extracted as evidence; (e) **Text Planning** - document proposition clusters are to be grouped and ordered according to the summary sentence structure; (f) **Sentence Fusion** - document propositions aligning to the same summary sentence are to be fused to generate that sentence; (g) **In-context Fusion** - all document propositions, marked within the documents, are to be fused to generate the full summary.

same meaning or occurrence. Taken from our alignments data, the text units in a cluster are all the spans across the document set that align to the same proposition in the corresponding reference summary (Figure 2 (c)).

In the test data (Table 1), there are 100 instances (topics) for the task, each with an average of 22.6 spans that need to be clustered into 13.3 clusters. A cluster has an average of 1.7 spans, where 577 clusters are singletons.

Evaluation. The traditional clustering metrics are applicable for this task, namely homogeneity (clusters contain only instances that are members of the same gold cluster), completeness (instances that are members of the same gold cluster are also placed together in a predicted cluster), and V-measure (harmonic mean of the first two). In §5 we only report the V-measure for simplicity.

4.3 Evidence Detection

Given a set of documents and a proposition-like phrase, the goal of this task is to find all mentions of the phrase within the documents. For summarization, this could assist in providing attribution for the summary content (Ernst et al., 2022; Hosking et al., 2023; Xiao, 2023). It also relates to fact extraction and verification, where a claim needs to be backed by evidence from within a corpus (e.g., Schuster et al., 2021), and coreference search (Eirew et al., 2022) where corefering mentions of an event are to be detected.

Task definition and dataset derivation. Given document set D and a textual query q , the task is to return all mentions of q within D . With respect to the alignments data, a query is a summary proposition, and the mentions are the document spans that align to it (Figure 2 (d)).

In the test data (Table 1) there are 1332 instances (total number of clusters) for the task. A query requires retrieving an average of 1.7 spans from the document set (with ~ 3 documents).

Evaluation. To evaluate this task we followed the coreference-search evaluation (Eirew et al., 2022) and (Tjong Kim Sang and Buchholz, 2000), and used token-based F_1 .

4.4 Sentence and Paragraph Planning

To produce a coherent text passage, it is necessary to plan the ordering of the information incorporated into the passage. This intermediate task was shown to guide models to generate better results (Moryossef et al., 2019), and has been applied for various generation tasks (Barzilay and Lapata, 2008; Chambers and Jurafsky, 2008; Faille et al., 2020). While most related works perform an evaluation extrinsically on the downstream generation task, we establish a dataset explicitly dedicated to ordering and sentence planning.

Task definition and dataset derivation. Given a list of proposition clusters $\{C_1, \dots, C_k\}$, where a cluster C_i represents a single piece of information, the task comprises two steps. (1) The clusters need to be ordered so that the respective information flows coherently. (2) After ordering, consecutive clusters that should construct a single sentence are to be grouped. Eventually, this two-stage planning task renders a layout for how to generate a passage containing all the information, in terms of passage- and sentence-level construction. As illustrated in Figure 2 (e), based on the alignments data, each proposition cluster is the set of spans in the document set that align to the same summary span. The ordering and grouping decisions are based on the summary structure: (1) the order of the clusters will be in accordance to the order of respective aligned spans in the summary, and (2) each cluster grouping corresponds to summary spans that come from the same summary sentence.

In the test data (Table 1) there are 100 instances (topics) for the task. An instance has an average of 13.3 information units (clusters) that require planning for passages with 8.3 sentences.

Evaluation. To evaluate the ordering of information clusters we used Kendall-Tau correlation between the predicted and the gold ordering (following (Lapata, 2006)). For the cluster grouping, we

used Homogeneity, Completeness and V-measure, viewed as a clustering assignment.

4.5 Sentence Fusion

Fusing various pieces of information into a single coherent sentence is a fundamental task that is required in many generation tasks, and specifically in summarization where the desired sentence should be concise. Traditionally, sentence fusion (Barzilay and McKeown, 2005) may generally involve two merging scenarios: fusing similar information by omitting redundancy but exploiting complement information from different mentions (Thadani and McKeown, 2013; Hirsch et al., 2023), versus combining different information with discourse relation (Geva et al., 2019). These two different fusion types were often addressed separately in prior datasets, while our dataset poses the two challenges simultaneously.

Task definition and dataset derivation. Given one or more clusters of paraphrastic texts, the task is to merge all the texts into a single coherent sentence that reflects the union of information in the texts. Furthermore, the information in the generated sentence should generally be presented in the order of the clusters, if more than one is given. The illustration in Figure 2 (f), portraying the derivation of data from alignments, shows that a cluster of texts consists of the spans from the source documents that align to the same summary sentence. Accordingly, the summary sentence acts as the fused sentence. If the sentence consists of more than one proposition, then all corresponding clusters of alignments act as the input.

In the test data (Table 1) there are 834 instances (summary sentences) for the task. Each sentence is a fusion of an average of 2.7 propositions (~ 1.6 clusters with ~ 1.7 propositions).

Evaluation. Following (Lebanoff et al., 2020; Brook Weiss et al., 2021) we apply lexical similarity between the predicted and the gold sentence using ROUGE F_1 .

4.6 In-context Fusion

Following Slobodkin et al. (2022), another valuable task is generating a passage that consolidates highlights marked within documents. The context around the highlights should assist in resolving anaphora and coreference issues when generating the output. This ability is applicable on its own,

e.g., to help a user prepare an abstract out of specially desired content (Slobodkin et al., 2023b). Likewise, it can naturally be used in a summarization process where content is first selected, followed by the in-context fusion step to generate the output. The independent fuser allows for any conditional selection of highlights (e.g., for query-focused summarization). Our dataset is different from that of Slobodkin et al. (2022) in that we address the multi-document setting, where the input is larger and redundancy is a more prevalent phenomenon.

Task definition and dataset derivation. Given a set of documents and marked spans within the documents, the task is to generate a coherent passage that contains all and only the information in the marked spans. With respect to our alignments data, the highlights are all the spans aligning to the reference summary, and the summary acts as the fused passage to generate (Figure 2 (g)).

In the test data (Table 1) there are 100 instances (topic) for the task. Each document set consists of an average of 22.6 spans that need to be fused into a passage.

Evaluation. We used ROUGE F_1 between the predicted and the gold passage.

5 Baseline Experiments

We next examine the performance of current technology on the six aforementioned datasets. For each task, we consider a dedicated trained model and an execution of `gpt-3.5-turbo`, once in zero-shot mode and once with an in-context example (prompts in Appendix B). Since the large size of a multi-document set limits model architecture options, we resorted to models devised for the multi-document setting. We used our train set to train the dedicated models (§3.2). In addition, the large input sizes afforded us to execute GPT with only one in-context example. We explain the finetuned models in §5.1 and discuss the results in §5.2.

5.1 Finetuned Models

For **Salience Detection**, we finetuned the Cross-Document Language Model (CDLM; Caciularu et al., 2021), an encoder-only model that was trained specifically to handle multi-document inputs by assigning global attention from selected tokens to the entire document set. In our case, we

added a classification head and input the document set with special tokens marking a candidate span, while the target is a binary decision for whether a span is salient or not. Global attention was assigned to the candidate tokens. At inference time, we must mark candidate spans within the document set for the model to classify for salience. To that end we use Open Information Extraction (Stanovsky et al., 2018), following (Ernst et al., 2021, which was also used to create the train set).

For both the **Proposition Clustering** and **Evidence Detection** tasks we employ the SuperPAL (Ernst et al., 2021) model which was pre-trained to match pairs of similar spans. For each pair, the model outputs a score between 0 (no match) and 1 (match). For Proposition Clustering, the scores were used as input for an Agglomerative Clustering method to group similar spans. For Evidence Detection, we paired each summary span with each candidate *document* span, and selected spans scored above the original SuperPAL threshold (0.5).

For the **Sentence and Paragraph Planning** and **Sentence Fusion** tasks, we finetuned `Flan-T5-XXL` (Chung et al., 2022) as a sequence-to-sequence task with suitable instructions for each of the tasks (detailed in Appendix B). For the Planning task, we grouped together all document spans that are aligned to the same summary span, selected a random representative per group, and numbered each representative. Then the model is required to output an ordered list of list of indices, where each sub-list represents a prospective summary sentence, and the order of the sub-lists outlines the passage structure. For Sentence Fusion, the model receives a group of document spans and is expected to generate the aligned summary span.

For the **In-context Fusion** task, we used the QAMDen model (Caciularu et al., 2023), a recent encoder-decoder transformer made for the multi-document setting, and pre-trained on Multi-News. We finetuned the model for our task by surrounding the highlighted spans with special tokens, akin to Slobodkin et al. (2022, 2023a).

5.2 Results

Results are presented in Table 2. As can be seen, even though we used a much larger model for the zero-shot and in-context modes (`gpt-3.5-turbo`), the finetuned models perform better in all tasks except for Proposition Clustering. Apparently, when the input data is short, without having to input all documents, the GPT

Task Metric	Saliency	Clustering	Evidence	Planning		Sent. Fusion			In-context Fusion		
	F_1	V	F_1	Kendall's τ	V	R-1	R-2	R-L	R-1	R-2	R-L
Finetuned	0.49	0.33	0.36	0.36	0.76	0.45	0.26	0.39	40.54	16.82	22.42
GPT Zero-shot	0.27	0.71	0.22	0.29	0.70	0.43	0.22	0.34	38.45	13.29	19.94
GPT In-Context	0.31	0.83	0.32	0.33	0.67	0.38	0.17	0.29	40.01	13.65	20.43

Table 2: Performance of finetuned, zero-shot GPT, and in-context learning GPT models on all tasks. Overall, a smaller finetuned model yields better results than the GPT counterpart.

Overlap Measure	unigram	bigram	trigram
<i>Alignment Pair</i>	43.39	23.44	16.12
<i>Cluster Max</i>	51.81	31.38	22.92
<i>Full Cluster</i>	54.20	32.26	23.24
<i>In-Cluster</i>	35.72	17.60	11.64

Table 3: Percentage of n-gram overlap of different source span groups with respect to their aligned summary span. Overall, summary spans are partially abstractive. We also measure the n-gram overlap between document spans within the same cluster (*In-Cluster*). This indicates lexical diversity in redundant source spans.

model performs better. In addition, the score differences between the models are quite low on the two Fusion tasks, where the output performance is more subjective. We find that the in-context example assists GPT in most tasks with respect to zero-shot mode. The sentence fusion was harmed since GPT tended to ask for additional details in its output, as it tried to comply with certain characteristics of the example it received in-context. Overall, future research can examine how to push ahead strong large language models either with further fine-tuning or with in-context examples of very large inputs.

6 Source Dataset Characteristics

The alignments extracted from the MDS dataset (Multi-News) sheds light on various characteristics of the source dataset. It helps understand the amount of information redundancy, level of abstractiveness, and spread of content within documents.

Information redundancy. From the alignment data statistics in Table 1, the average size of clusters is 1.7 propositions from 3 documents, which reflects on the low informational redundancy within a document set. Assuming information redundancy impacts apparent importance, this property may affect the ability to recognize salient information and plan passage structure.

Abstractiveness. A cluster of document-set spans and its respective summary span can be ex-

amined for paraphrastic differences to measure abstractiveness within the data. To that end we present in Table 3 the conventional n-gram overlap metric between spans on several levels: (1) *Alignment Pair* is the percentage of summary span n-grams that appear also in the document span; (2) *Cluster Max* is the maximum pair overlap score in a cluster, which indicates general summary-source abstractiveness; (3) *Full Cluster* is the n-gram overlap between the bag of document spans in a cluster with respect to their aligned summary span; (4) *In-Cluster* is the average pairwise overlap between cluster members. *Full Cluster* is only slightly higher than *Cluster Max*, as additional members of the cluster do not contribute much to cover the summary span. This indicates that the summary span is mostly copied from a single document span, and does not merge texts from different places in the documents. To strengthen this insight, *In-Cluster* produces a relatively low score, meaning that abstractiveness is high within a cluster, while one of the cluster members is more lexically similar to the summary span. We can also learn that the relatively low number of clusters per summary sentence (1.6) indicates that summaries only require occasional fusion of different information units. Overall, these insights reinforce that the summaries in the Multi-News MDS dataset have somewhat low abstractiveness, as also observed from Fabbri et al. (2019), though when information repeats within a document set, it is mentioned in noticeably different phrasing.

Spread within Documents. Since a summary is comprehensively aligned with its corresponding document set through our data, we can investigate the individual importance of each document for the summary content (Wolhandler et al., 2022). We find that out of the ~ 3 documents per topic, only $\sim 87\%$ of documents have aligning information with the summary on average, meaning some topics do not require all the documents for the summary. On the cluster level, we find that a summary

span aligns only to $\sim 53\%$ of the documents on average. This stresses the low information redundancy mentioned before, where each summary proposition appears in one or two documents.

7 Conclusion

We advocate the potential utility of proposition-level summary-source information alignment, particularly in the multi-document setting, for exposing a wide range of summarization-related tasks. Specifically, we reveal that these alignments induce datasets for a broad range of appealing tasks arising in summarization, and applicable as standalone tasks. We annotated a high-quality test dataset of alignments, and automatically compiled large-scale train and dev sets. From the alignments data, we automatically derived datasets for six distinct tasks. Our released dataset collection, along with our baselines and analyses, promotes future research on a challenging multi-text task suite.

Limitations

This study obtains alignments out of a MDS dataset in the news domain. To automatically extract alignments, we leveraged SuperPAL, which itself is trained on news data. The model would likely extract less accurate alignments in other domains. Generally, it is worthwhile to perform an alignment study like ours in additional domains and languages. The alignment process and guidelines, as well as the derived tasks may differ in accordance to the source MDS data.

The quality of our alignments is dependent on the quality of the source MDS dataset (Multi-News) from which we extract the alignments. For example if a reference summary is not fully faithful or comprehensive for some reason, this may have an effect on our alignment assumptions. Our analysis in §6 sheds light on some of these discrepancies.

The baselines we presented are limited to the prompts we used. Other prompts may yield different results.

References

Kristjan Arumae, Parminder Bhatia, and Fei Liu. 2019. [Towards Annotating and Creating Summary Highlights at Sub-sentence Level](#). In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 64–69, Hong Kong, China. Association for Computational Linguistics.

Regina Barzilay and Mirella Lapata. 2008. [Modeling local coherence: An entity-based approach](#). *Computational Linguistics*, 34(1):1–34.

Regina Barzilay and Kathleen R. McKeown. 2005. [Sentence Fusion for Multidocument News Summarization](#). *Computational Linguistics*, 31(3):297–328.

Daniela Brook Weiss, Paul Roit, Ayal Klein, Ori Ernst, and Ido Dagan. 2021. [QA-Align: Representing Cross-Text Content Overlap by Aligning Question-Answer Propositions](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9879–9894, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Avi Caciularu, Arman Cohan, Iz Beltagy, Matthew Peters, Arie Cattan, and Ido Dagan. 2021. [CDLM: Cross-Document Language Modeling](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2648–2662, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Avi Caciularu, Matthew Peters, Jacob Goldberger, Ido Dagan, and Arman Cohan. 2023. [Peek Across: Improving Multi-Document Modeling via Cross-Document Question-Answering](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1970–1989, Toronto, Canada. Association for Computational Linguistics.

Xiaoyan Cai, Wenjie Li, You Ouyang, and Hong Yan. 2010. [Simultaneous Ranking and Clustering of Sentences: A Reinforcement Approach to Multi-Document Summarization](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 134–142, Beijing, China. Coling 2010 Organizing Committee.

Nathanael Chambers and Dan Jurafsky. 2008. [Unsupervised Learning of Narrative Event Chains](#). In *Proceedings of ACL-08: HLT*, pages 789–797, Columbus, Ohio. Association for Computational Linguistics.

Yen-Chun Chen and Mohit Bansal. 2018a. [Fast Abstractive Summarization with Reinforce-Selected Sentence Rewriting](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–686, Melbourne, Australia. Association for Computational Linguistics.

Yen-Chun Chen and Mohit Bansal. 2018b. [Fast Abstractive Summarization with Reinforce-Selected Sentence Rewriting](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–686, Melbourne, Australia. Association for Computational Linguistics.

Sangwoo Cho, Logan Lebanoff, Hassan Foroosh, and Fei Liu. 2019. [Improving the Similarity Measure of](#)

- Determinantal Point Processes for Extractive Multi-Document Summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1027–1038, Florence, Italy. Association for Computational Linguistics.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling Instruction-Finetuned Language Models](#).
- Amir Cohen, Hila Gonen, Ori Shapira, Ran Levy, and Yoav Goldberg. 2022. [McPhraSy: Multi-Context Phrase Similarity and Clustering](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3538–3550, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- William B. Dolan and Chris Brockett. 2005. [Automatically Constructing a Corpus of Sentential Paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Alon Eirew, Avi Caciularu, and Ido Dagan. 2022. [Cross-document Event Coreference Search: Task, Dataset and Modeling](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 900–913, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ori Ernst, Avi Caciularu, Ori Shapira, Ramakanth Pasunuru, Mohit Bansal, Jacob Goldberger, and Ido Dagan. 2022. [Proposition-Level Clustering for Multi-Document Summarization](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1765–1779, Seattle, United States. Association for Computational Linguistics.
- Ori Ernst, Ori Shapira, Ramakanth Pasunuru, Michael Lepioshkin, Jacob Goldberger, Mohit Bansal, and Ido Dagan. 2021. [Summary-Source Proposition-level Alignment: Task, Datasets and Supervised Baseline](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 310–322, Online. Association for Computational Linguistics.
- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. [Multi-News: A Large-Scale Multi-Document Summarization Dataset and Abstractive Hierarchical Model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084, Florence, Italy. Association for Computational Linguistics.
- Juliette Faille, Albert Gatt, and Claire Gardent. 2020. [The Natural Language Pipeline, Neural Text Generation and Explainability](#). In *2nd Workshop on Interactive Natural Language Technology for Explainable Artificial Intelligence*, pages 16–21, Dublin, Ireland. Association for Computational Linguistics.
- Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2013. [PPDB: The Paraphrase Database](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 758–764, Atlanta, Georgia. Association for Computational Linguistics.
- YanJun Gao, Andrew Warner, and Rebecca Passonneau. 2018. [PyrEval: An Automated Method for Summary Content Analysis](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Mor Geva, Eric Malmi, Idan Szpektor, and Jonathan Berant. 2019. [DiscoFuse: A Large-Scale Dataset for Discourse-Based Sentence Fusion](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3443–3455, Minneapolis, Minnesota. Association for Computational Linguistics.
- Tsutomu Hirao, Hidetaka Kamigaito, and Masaaki Nagata. 2018. [Automatic Pyramid Evaluation Exploiting EDU-based Extractive Reference Summaries](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4177–4186, Brussels, Belgium. Association for Computational Linguistics.
- Eran Hirsch, Valentina Pyatkin, Ruben Wolhandler, Avi Caciularu, Asi Shefer, and Ido Dagan. 2023. [Revisiting sentence union generation as a testbed for text consolidation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7038–7058, Toronto, Canada. Association for Computational Linguistics.
- Tom Hosking, Hao Tang, and Mirella Lapata. 2023. [Attributable and Scalable Opinion Summarization](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8488–8505, Toronto, Canada. Association for Computational Linguistics.
- Kundan Krishna, Prakhar Gupta, Sanjana Ramprasad, Byron Wallace, Jeffrey Bigham, and Zachary Lipton. 2023. [USB: A Unified Summarization Benchmark Across Tasks and Domains](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8826–8845, Singapore. Association for Computational Linguistics.
- Mirella Lapata. 2006. [Automatic Evaluation of Information Ordering: Kendall’s Tau](#). *Computational Linguistics*, 32(4):471–484.

- Logan Lebanoff, Franck Dernoncourt, Doo Soon Kim, Lidan Wang, Walter Chang, and Fei Liu. 2020. [Learning to Fuse Sentences with Transformers for Summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4136–4142, Online. Association for Computational Linguistics.
- Logan Lebanoff, Kaiqiang Song, Franck Dernoncourt, Doo Soon Kim, Seokhwan Kim, Walter Chang, and Fei Liu. 2019. [Scoring Sentence Singletons and Pairs for Abstractive Summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2175–2189, Florence, Italy. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A Package for Automatic Evaluation of Summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Yuning Mao, Yanru Qu, Yiqing Xie, Xiang Ren, and Jiawei Han. 2020. [Multi-document Summarization with Maximal Marginal Relevance-guided Reinforcement Learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1737–1751, Online. Association for Computational Linguistics.
- Amit Moryossef, Yoav Goldberg, and Ido Dagan. 2019. [Step-by-Step: Separating Planning from Realization in Neural Data-to-Text Generation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2267–2277, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ani Nenkova and Rebecca Passonneau. 2004. [Evaluating Content Selection in Summarization: The Pyramid Method](#). In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 145–152, Boston, Massachusetts, USA. Association for Computational Linguistics.
- OpenAI. 2023. [Models - OpenAI API](#). openai.com. Accessed: 2023-06-01.
- Xuan-Hieu Phan, Le-Minh Nguyen, and Susumu Horiguchi. 2008. [Learning to Classify Short and Sparse Text & Web with Hidden Topics from Large-scale Data Collections](#). In *Proceedings of the 17th International Conference on World Wide Web, WWW '08*, page 91–100, New York, NY, USA. Association for Computing Machinery.
- Héctor R. Ponce, Richard E. Mayer, and Ester E. Méndez. 2022. [Effects of Learner-Generated Highlighting and Instructor-Provided Highlighting on Learning from Text: A Meta-Analysis](#). *Educational Psychology Review*, pages 989–1024.
- Paul Roit, Ayal Klein, Daniela Stepanov, Jonathan Mamou, Julian Michael, Gabriel Stanovsky, Luke Zettlemoyer, and Ido Dagan. 2020. [Controlled Crowdsourcing for High-Quality QA-SRL Annotation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7008–7013, Online. Association for Computational Linguistics.
- Ágnes Sándor and Angela Vorndran. 2014. [Highlighting Salient Sentences for Reading Assistance](#), pages 43–55. Springer Fachmedien Wiesbaden, Wiesbaden.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. [Get Your Vitamin C! Robust Fact Verification with Contrastive Evidence](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 624–643, Online. Association for Computational Linguistics.
- Jessica Zeitz Self, Rebecca Zeitz, Chris North, and Alan L. Breitler. 2013. [Auto-Highlighter: Identifying Salient Sentences in Text](#). In *2013 IEEE International Conference on Intelligence and Security Informatics*, pages 260–262.
- Ori Shapira, David Gabay, Yang Gao, Hadar Ronen, Ramakanth Pasunuru, Mohit Bansal, Yael Amsterdamer, and Ido Dagan. 2019. [Crowdsourcing Lightweight Pyramids for Manual Summary Evaluation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 682–687, Minneapolis, Minnesota. Association for Computational Linguistics.
- Aviv Slobodkin, Avi Caciularu, Eran Hirsch, and Ido Dagan. 2023a. [Don't Add, don't Miss: Effective Content Preserving Generation from Pre-Selected Text Spans](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12784–12800, Singapore. Association for Computational Linguistics.
- Aviv Slobodkin, Niv Nachum, Shmuel Amar, Ori Shapira, and Ido Dagan. 2023b. [SummHelper: Collaborative Human-Computer Summarization](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 554–565, Singapore. Association for Computational Linguistics.
- Aviv Slobodkin, Paul Roit, Eran Hirsch, Ori Ernst, and Ido Dagan. 2022. [Controlled Text Reduction](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5699–5715, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Gabriel Stanovsky, Julian Michael, Luke Zettlemoyer, and Ido Dagan. 2018. [Supervised Open Information Extraction](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 885–895,

New Orleans, Louisiana. Association for Computational Linguistics.

Kapil Thadani and Kathleen McKeown. 2013. [Supervised sentence fusion with single-stage inference](#). In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 1410–1418, Nagoya, Japan. Asian Federation of Natural Language Processing.

Erik F. Tjong Kim Sang and Sabine Buchholz. 2000. [Introduction to the CoNLL-2000 Shared Task Chunking](#). In *Fourth Conference on Computational Natural Language Learning and the Second Learning Language in Logic Workshop*.

Xiaojun Wan and Jianwu Yang. 2008. [Multi-document Summarization Using Cluster-based Link Analysis](#). In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '08*, page 299–306, New York, NY, USA. Association for Computing Machinery.

Ruben Wolhandler, Arie Cattan, Ori Ernst, and Ido Dagan. 2022. [How “Multi” is Multi-Document Summarization?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5761–5769, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Min Xiao. 2023. [Multi-doc Hybrid Summarization via Salient Representation Learning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 5: Industry Track)*, pages 379–389, Toronto, Canada. Association for Computational Linguistics.

Jiaming Xu, Bo Xu, Peng Wang, Suncong Zheng, Guan-hua Tian, Jun Zhao, and Bo Xu. 2017. [Self-Taught Convolutional Neural Networks for Short Text Clustering](#). *Neural Networks*, 88:22–31.

Qian Yang, Rebecca Passonneau, and Gerard de Melo. 2016. [PEAK: Pyramid Evaluation via Automated Knowledge Extraction](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1).

Shiyue Zhang and Mohit Bansal. 2021. [Finding a Balanced Degree of Automation for Summary Evaluation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6617–6632, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Xingxing Zhang, Mirella Lapata, Furu Wei, and Ming Zhou. 2018. [Neural Latent Extractive Document Summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 779–784, Brussels, Belgium. Association for Computational Linguistics.

Yang Zhang, Yunqing Xia, Yi Liu, and Wenmin Wang. 2015. [Clustering Sentences with Density Peaks for Multi-document Summarization](#). In *Proceedings of*

the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 1262–1267, Denver, Colorado. Association for Computational Linguistics.

Zixuan Zhang, Heba Elfardy, Markus Dreyer, Kevin Small, Heng Ji, and Mohit Bansal. 2023. [Enhancing Multi-Document Summarization with Cross-Document Graph-based Information Extraction](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1696–1707, Dubrovnik, Croatia. Association for Computational Linguistics.

A Further Details on Baseline Implementations

We describe here details regarding the baseline models outlined in §5.1. Specifically, for each task, we describe heuristics applied in case an LLM model outputs an answer in the wrong format (relevant for the GPT and finetuned Flan baselines).

Salience Detection and Evidence Detection. In these tasks we asked the model to extract (fully copy) spans from the source text. However, in many cases the model slightly changed the extracted span by adding or omitting a word. Since our evaluation for these two tasks is token-level F_1 , we need to locate the extracted spans in the source documents. To do so, we extracted proposition candidates (using OpenIE; Stanovsky et al., 2018) from the source, and for each predicted span, we found the OpenIE proposition with the highest lexical overlap.

Proposition Clustering. In this task we asked the model to cluster spans by predicting a cluster index for each span. However, in some of the cases the model did not provide an index for an input spans. In such cases, we assign a random (existing) cluster index to this text span.

Sentence and Paragraph Planning. The model is tasked with outputting a list of lists describing the order of information (span clusters) within the final paragraph, each represented by its index. In some cases, the model omits an index of one of the spans, adds a non-existent index, or even repeats an existing index more than once. To cope with this, we removed non-existing indices, kept only the first occurrence of a repeating index, and appended a randomly ordered list of missing indices.

Sentence Fusion and In-Context Fusion. As these two tasks generate free text and are evaluated by ROUGE, mis-formatting is not relevant.

B Model Prompts

Table 4 presents the prompts used on the gpt-3.5-turbo-0301 model for the six tasks, in zero-shot and in-context learning modes. Table 5 shows the prompts used for finetuning a flan-t5-xxl model for the Planning and Sentence Fusion tasks.

C Full Annotation Guidelines

This section describes the complete annotation guidelines for the crowdsourced alignment procedure.

C.1 Summary-related Guidelines

As mentioned in §3.1, we guide annotators to separate summary sentences into separate events, focusing on one event at a time. An event is identified as a predicate alongside all its arguments, with instructions for annotators to include all associated arguments, even if repeated across events, e.g., “Jane came by and left”, where “Jane” is part of both the “came by” and “left” events.

We also aim to address facts represented in various grammatical forms, but for the sake of simplification for the annotators, we highlight the following two forms: • **SECONDARY VERB**: This involves nested events, where a smaller event serves as an argument for a larger one, e.g., “John insisted on inviting her”. Annotators are guided to merge these into a single event if both appear in the source document, but to align only the nested event if it alone is present. Annotators are also guided to distinguish between prospective and transpired events. For instance, “John insisted on inviting her” (prospective) should not align with “John invited her” (transpired) since they convey different events. Moreover, in instances of nested spans containing distinct events, like “She said she arrived and went to bed” annotators should align the primary event with each nested event separately if both are documented. • **CONNECTING WORDS**: For events linked by discourse markers, which we refer to in our guidelines as *connecting words*, annotators are trained to identify when these words indicate a genuine connection, such as “He ate because he was hungry” versus when they merely place events side by side, like in “He went home and ate an apple”. In the former, both events are to be combined into a single alignment if present in the document, while in the latter, events should be aligned separately.

C.2 Document-related Guidelines

On the document side, we provide the annotators with the following guidelines on how to align a span to a summary proposition. • **PARAPHRASING**: We guide our workers to not depend exclusively on phrases with common words, since the matching document phrases are frequently a paraphrase of their summary counterparts. • **CONSECUTIVENESS**: We instruct our workers to avoid highlighting unnecessary details, and keep the highlights non-consecutive if necessary. • **ENTAILMENT**: As described in §3, we instruct the annotators to also align document spans that either entail the summary event, or are entailed by it. E.g., “John ate an apple” versus “John ate fruit”. • **MISSING DETAILS**: In cases where some details of the summary event are missing from the currently inspected document, we guide our annotators to leave those unhighlighted on the summary side, and align only the details that do appear. • **EXHAUSTIVENESS**: We also train our workers to identify all document mentions of the current summary event, and align each one separately.

D Obtaining Gold Alignment Clusters

Since the alignment annotation is conducted one document at a time, the coreferring propositions from across documents (those aligning to the same summary proposition) need to be clustered together. Considering that a summary proposition may be marked slightly differently each time (with different boundaries), we allow a 0.5 (tuned threshold) intersection-over-union (of tokens) to consider summary spans as referring to the same proposition. To validate this threshold, we manually examined 10 topics that contain 94 clusters. We found that only one cluster merged irrelevant propositions, and only 3 pairs of clusters should have been merged to a larger cluster. Accordingly, this enables almost perfect clustering of document spans for our data. For the Evidence Detection task, we aggregated the cluster query as the union of all aligned summary spans in this cluster.

E Model Training

Saliency Detection. We trained the CDLM model for 2 epochs with learning rate of 1e-5 and batch size of 3 instances on 3 A100 GPUs for one hour (meaning an effective batch size was 9).

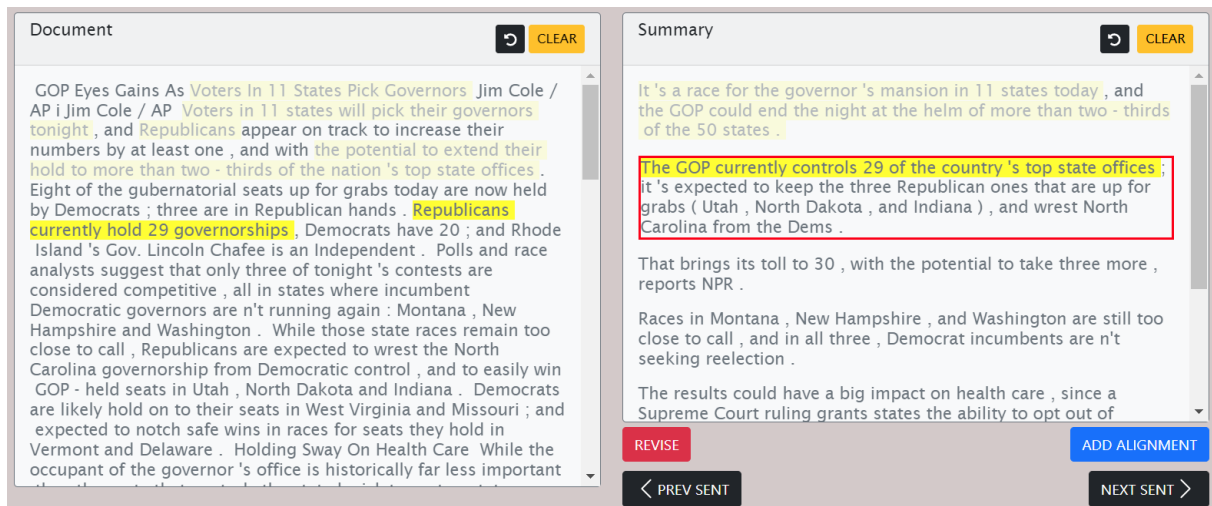


Figure 3: The alignment annotation interface. The annotator marks a span (proposition) in the summary (right) along with all matching spans in the current document (left). To minimize cognitive load, a summary is shown next to a single document at a time, and the procedure is conducted separately for all documents in the document set. Also visual focus is placed on one summary sentence at a time (red rectangle) to orient the process.

In-Context Fusion. We trained the PEEK model for 50,000 steps with learning rate of $3e-5$ and batch size of 16 instances on 1 A100 GPUs for one hour.

Sentence Fusion and Planning. We finetuned a Flan-T5-XXL as a Sequence-to-Sequence task using LoRA and applied 8-bit quantization for optimization. We trained using a sample of 10K examples derived from our train-set, for one epoch using a learning rate of $5e-5$ and using AdamW optimizer.

F Data Example

We show an example of one topic of our SPARK data suite, starting from the alignment annotation (Figure 4) followed by its derived instances (Figures 5, 6, 7, 8, 9).

Summary

Two sources tell Variety the Super Bowl LIII halftime act has been chosen, and it's Maroon 5. Another source tells Us, "The offer has been extended and they've pretty much accepted," while ETOnline says the band is the "frontrunner" to perform. The band itself, fronted by The Voice coach Adam Levine, has not commented, but Levine has in the past made it clear he'd love to perform there. As for the NFL, it says in a statement, "It's a Super Bowl tradition to speculate about the performers for the Pepsi Halftime Show. We are continuing to work with [longtime sponsor] Pepsi on our plans but do not have any announcements to make on what will be another epic show." The big game will be played at Atlanta, Georgia's Mercedes-Benz Stadium on Feb. 3, 2019.

Document 1

Football fans, get ready to rock. While there's no telling which two teams will face off when Super Bowl LIII hits Atlanta on February 3, a source tells Us that Maroon 5 have been tapped to grace the halftime stage.

"The offer has been extended and they've pretty much accepted," says the insider in the latest issue of Us Weekly. It's a dream come true for Adam Levine and his six bandmates. As the Voice coach told Howard Stern in 2015, "We very actively want to play the Super Bowl."

Though at the time the 39-year-old frontman denied rumors that the group was in talks to play the big game that year, he also told Stern they had been on a "short list" of groups for a while.

⋮

Document 2

Looks like it's not a "Secret" anymore!

A source close to the situation says Maroon 5 is scheduled to perform Super Bowl LIII halftime on Feb 3rd in Atlanta, Georgia.

A second source close to the NFL tells ET that Maroon 5 is the halftime performer for Super Bowl LIII.

The band -- made up of frontman Adam Levine with Jesse Carmichael, James Valentine, PJ Morton, Matt Flynn, Sam Farrar and Mickey Madden -- is currently on their Red Pill BluesTour, but conveniently have a break between a New Year's Eve show in Las Vegas and the Australian leg of their tour, which kicks off on Feb. 19.

Levine has previously expressed interest in performing the halftime show, telling Howard Stern in 2015 that he and his bandmates "very actively want to play the Super Bowl."

A spokesperson for the NFL released a statement to ET regarding who is being considered to perform on Super Bowl Sunday. "It's a Super Bowl tradition to speculate about the performers for the Pepsi Halftime Show," the statement begins. "We are continuing to work with Pepsi on our plans, but do not have any announcements to make on what will be another epic show."

⋮

Document 3

Maroon 5 will be the halftime performers at the 2019 Super Bowl, multiple sources confirm to Variety.

⋮

Reached for comment, a NFL rep tells Variety, "It's a Super Bowl tradition to speculate about the performers for the Pepsi Halftime Show. We are continuing to work with [longtime sponsor] Pepsi on our plans but do not have any announcements to make on what will be another epic show."

A rep for Maroon 5 could not immediately be reached for comment.

Figure 4: The manual alignment annotation on topic31 from our data. The documents have been shortened for presentation purposes.

Document 1

Football fans, get ready to rock. While there's no telling which two teams will face off when **Super Bowl LIII hits Atlanta on February 3, a source tells Us that Maroon 5 have been tapped to grace the halftime stage.**

"The offer has been extended and they've pretty much accepted," says the insider in the latest issue of Us Weekly. It's a dream come true for Adam Levine and his six bandmates. As the Voice coach told Howard Stern in 2015, "We very actively want to play the Super Bowl."

Though at the time the 39-year-old frontman denied rumors that the group was in talks to play the big game that year, he also told Stern they had been on a "short list" of groups for a while.

Currently on tour promoting their 2017 album, Red Pill Blues, the band has a break from shows in between a Las Vegas New Year's Eve gig and the February 19 start of the Australian leg of their tour - leaving their schedules wide open for the February 3 game.

Before embarking on the tour in May, Levine opened up to Ellen DeGeneres about bringing his wife, Behati Prinsloo, and daughters Dusty Rose, 23 months, and Gio Grace, 7 months, on the road with him for the first time. "We'll probably have to have a family bus," he told the talk show host. "Touring with the kids will be fun because it will be an amazing experience for them. They get to see the world."

At last year's Super Bowl in Minneapolis, Minnesota, Justin Timberlake took the stage to deliver a medley of his hits and a tribute to the late singer Prince.

For more on Maroon 5 and Super Bowl LIII, pick up the new issue of Us Weekly, on stands now.

Reporting by Jennifer Peros

Sign up now for the Us Weekly newsletter to get breaking celebrity news, hot pics and more delivered straight to your inbox!

Want stories like these delivered straight to your phone? Download the Us Weekly iPhone app now!

Document 2

Looks like it's not a "Secret" anymore!

A source close to the situation says Maroon 5 is scheduled to perform Super Bowl LIII halftime on Feb 3rd in Atlanta, Georgia.

A second source close to the NFL tells ET that Maroon 5 is the halftime performer for Super Bowl LIII.

The band -- made up of frontman Adam Levine with Jesse Carmichael, James Valentine, PJ Morton, Matt Flynn, Sam Farrar and Mickey Madden -- is currently on their Red Pill BluesTour, but conveniently have a break between a New Year's Eve show in Las Vegas and the Australian leg of their tour, which kicks off on Feb. 19.

Levine has previously expressed interest in performing the halftime show, telling Howard Stern in 2015 that **he and his bandmates "very actively want to play the Super Bowl."**

A spokesperson for the NFL released a statement to ET regarding who is being considered to perform on Super Bowl Sunday: "It's a Super Bowl tradition to speculate about the performers for the Pepsi Halftime Show," the statement begins. **"We are continuing to work with Pepsi on our plans, but do not have any announcements to make on what will be another epic show."**

ET has reached out to the band's reps, as well as Pepsi, who sponsors the halftime show.

Last year at the U.S. Bank Stadium in Minneapolis, Minnesota, Justin Timberlake did the honors of performing at halftime during the Super Bowl LI game, which the Philadelphia Eagles won 41-33 over the New England Patriots.

See highlights from his show, which included a tribute to Prince, below:

RELATED CONTENT:

Kelly Clarkson Fans Petition She Headline Super Bowl Halftime Show After Slaying US Open Performance

Tom Brady and Gisele Bundchen Have PDA-Packed Sunday One Week After Super Bowl Loss

Justin Timberlake Nearly Brings Super Bowl 'Selfie Kid' to Tears During 'Ellen' Surprise

Related Gallery

Document 3

Maroon 5 will be the halftime performers at the 2019 Super Bowl, multiple sources confirm to Variety.

While the group is stylistically more similar to recent performers like Justin Timberlake and Bruno Mars than more controversial ones such as Beyonce and Lady Gaga, Maroon 5's most recent hit was a tag-team with Cardi B - "Girls Like You" - and it recently collaborated with Kendrick Lamar on the song "Don't Wanna Know," so some left-field guest appearances on Super Bowl Sunday are certainly a possibility.

One of the great challenges in choosing a halftime performer isn't just its mainstream appeal, but in finding an artist who hasn't already performed. Over the years, Bruce Springsteen, Paul McCartney, U2, The Who, Katy Perry and even the Black Eyed Peas have all performed, and of course the 2004 set from Janet Jackson (featuring Timberlake) was marred by controversy due to her "wardrobe malfunction." Prince's 2007 performance - which saw him playing a galvanizing version of "Purple Rain" in a rainstorm - is widely regarded as the all-time best halftime set.

The NFL has faced an additional challenge this year due to its stance regarding former San Francisco 49ers quarterback Colin Kaepernick's decision to take a knee during the National Anthem as a sign of protest, which has spurred outrage in some quarters, from the president on down. Kaepernick has not played since the 2016 season and his contract has not been picked up by another team, largely due to the controversy. Reports emerged last year that Jay-Z had turned down an offer to perform during halftime, which he seemingly confirmed in the lyrics to his recent song with Beyonce, "Ape\$h-": "I said no to the Super Bowl: you need me, I don't need you."

Reached for comment, **a NFL rep tells Variety, "It's a Super Bowl tradition to speculate about the performers for the Pepsi Halftime Show. We are continuing to work with [longtime sponsor] Pepsi on our plans but do not have any announcements to make on what will be another epic show."**

A rep for Maroon 5 could not immediately be reached for comment.

Figure 5: An example of a **Salience Detection** instance derived from the alignments in Figure 4. All aligned document propositions are salient. These highlighted documents can also serve as input to the **In-context Passage Fusion** task, where the output would be the original reference summary.

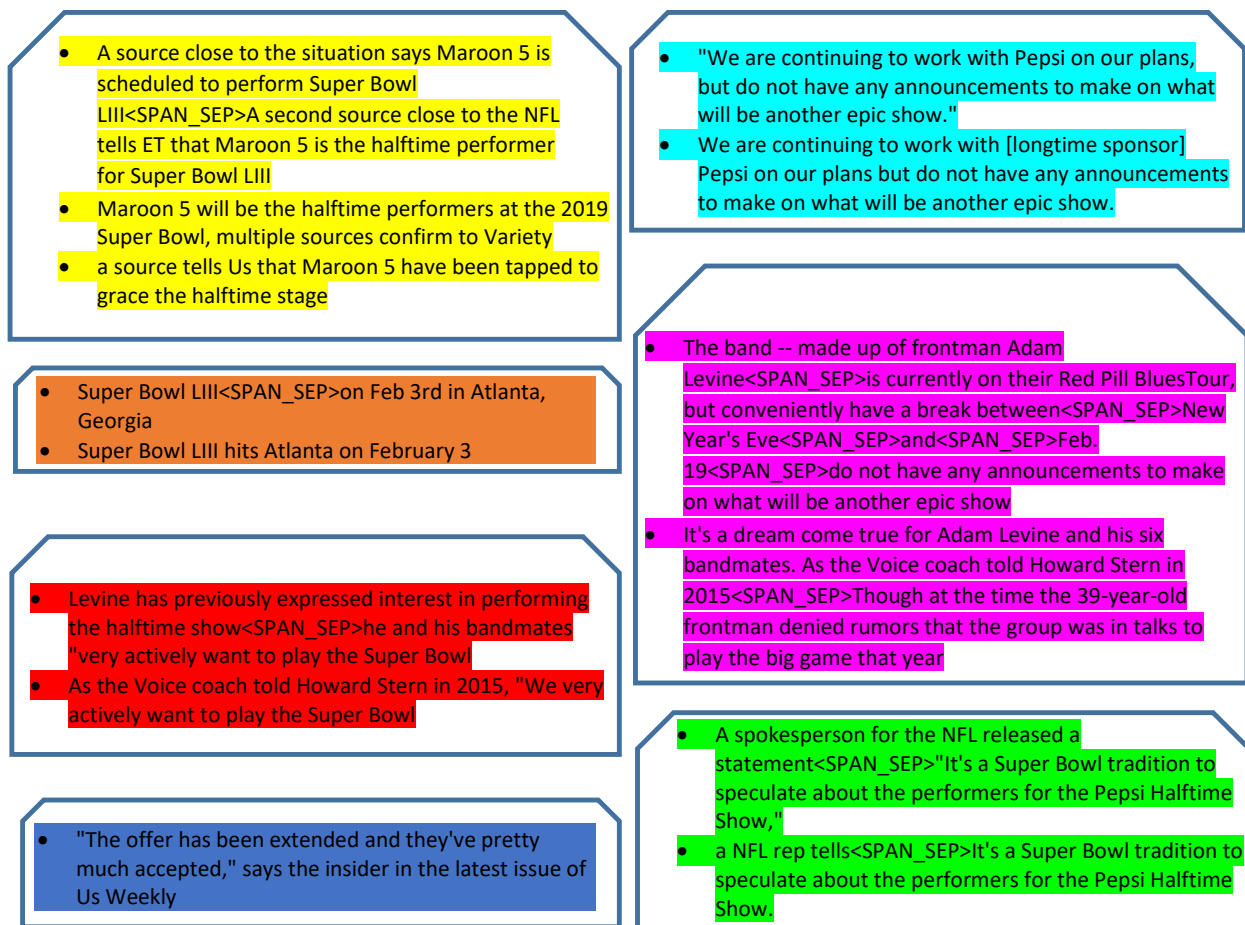


Figure 6: An example of a **Proposition Clustering** instance derived from the alignments in Figure 4. Clusters contain document propositions that are aligned to the same summary proposition.

Task	Prompt
Saliency Detection	Below are documents on the same topic in different user messages. Please copy exactly salient sub-sentential spans. Do not change the copied text.
Proposition Clustering	Below are text spans with indexes. Please cluster them into groups. Each group should contain spans that share the same information. Return a dict in the following format : <CLUSTER IDX>. Do not add anything beside the dict.
Evidence Detection	Below are documents on the same topic and a query. Please extract exactly short text spans from the documents that match the information in the query. Separate the spans with a new line and start each span with -.
Sentence and Passage Planning	Your task is to structure a set of information units, each pertinent to a central topic, into a cohesive paragraph. Begin by analyzing and logically arranging these units to ensure a seamless progression of ideas. Once you've established a coherent sequence, segment the units into subgroups that represent distinct conceptual sentences. Your final output should adhere to a Python list of lists format. Each internal list must encompass the indices of information units that belong to a particular conceptual sentence within the paragraph. Output Examples: "[[3, 4, 1], [0, 2], [5]]" Your output format MUST be a simple Python list of lists only, with no comments.
Sentence Fusion	Merge the following text clusters into a single coherent sentence: Your response MUST contain only one sentence.
In-Context Fusion	Below are documents on the same topic. Please summarize the spans marked in <> while using the set of documents as context.

Table 4: Prompts used as input to gpt-3.5-turbo to solve each of the six tasks in zero-shot mode. Similar prompts were used for the in-context learning mode, with the addition of an example.

Task	Prompt
Sentence and Passage Planning	<p>Task: Paragraph planning</p> <p>Your task is to structure a set of information units, each pertinent to a central topic, into a cohesive paragraph.</p> <p>Begin by analyzing and logically arranging these units to ensure a seamless progression of ideas. Once you've established a coherent sequence, segment the units into subgroups that represent distinct conceptual sentences.</p> <p>Your final output should adhere to a Python list of lists format. Each internal list must encompass the indices of information units that belong to a particular conceptual sentence within the paragraph.</p> <p>Output Examples: <code>"[[3, 4, 1], [0, 2], [5]]"</code></p>
Sentence Fusion	Merge the following text clusters into a single coherent sentence:

Table 5: Prompts used as input to `flan-t5-xxl` for finetuning the model for the respective tasks.

Query

Two sources tell Variety the Super Bowl LIII halftime act has been chosen, and it's Maroon 5.

Document 1

Football fans, get ready to rock. While there's no telling which two teams will face off when Super Bowl LIII hits Atlanta on February 3, a source tells Us that Maroon 5 have been tapped to grace the halftime stage.

"The offer has been extended and they've pretty much accepted," says the insider in the latest issue of Us Weekly. It's a dream come true for Adam Levine and his six bandmates. As the Voice coach told Howard Stern in 2015, "We very actively want to play the Super Bowl."

Though at the time the 39-year-old frontman denied rumors that the group was in talks to play the big game that year, he also told Stern they had been on a "short list" of groups for a while.

-
-

Document 2

Looks like it's not a "Secret" anymore!

A source close to the situation says Maroon 5 is scheduled to perform Super Bowl LIII halftime on Feb 3rd in Atlanta, Georgia.

A second source close to the NFL tells ET that Maroon 5 is the halftime performer for Super Bowl LIII.

The band -- made up of frontman Adam Levine with Jesse Carmichael, James Valentine, PJ Morton, Matt Flynn, Sam Farrar and Mickey Madden -- is currently on their Red Pill BluesTour, but conveniently have a break between a New Year's Eve show in Las Vegas and the Australian leg of their tour, which kicks off on Feb. 19.

-
-

Document 3

Maroon 5 will be the halftime performers at the 2019 Super Bowl, multiple sources confirm to Variety.

While the group is stylistically more similar to recent performers like Justin Timberlake and Bruno Mars than more controversial ones such as Beyonce and Lady Gaga, Maroon 5's most recent hit was a tag-team with Cardi B - "Girls Like You" - and it recently collaborated with Kendrick Lamar on the song "Don't Wanna Know," so some left-field guest appearances on Super Bowl Sunday are certainly a possibility.

-
-

Figure 7: An example of a **Evidence Detection** instance derived from the alignments in Figure 4. The evidences are the document propositions aligned to the query from the summary. The documents have been shortened for presentation purposes.

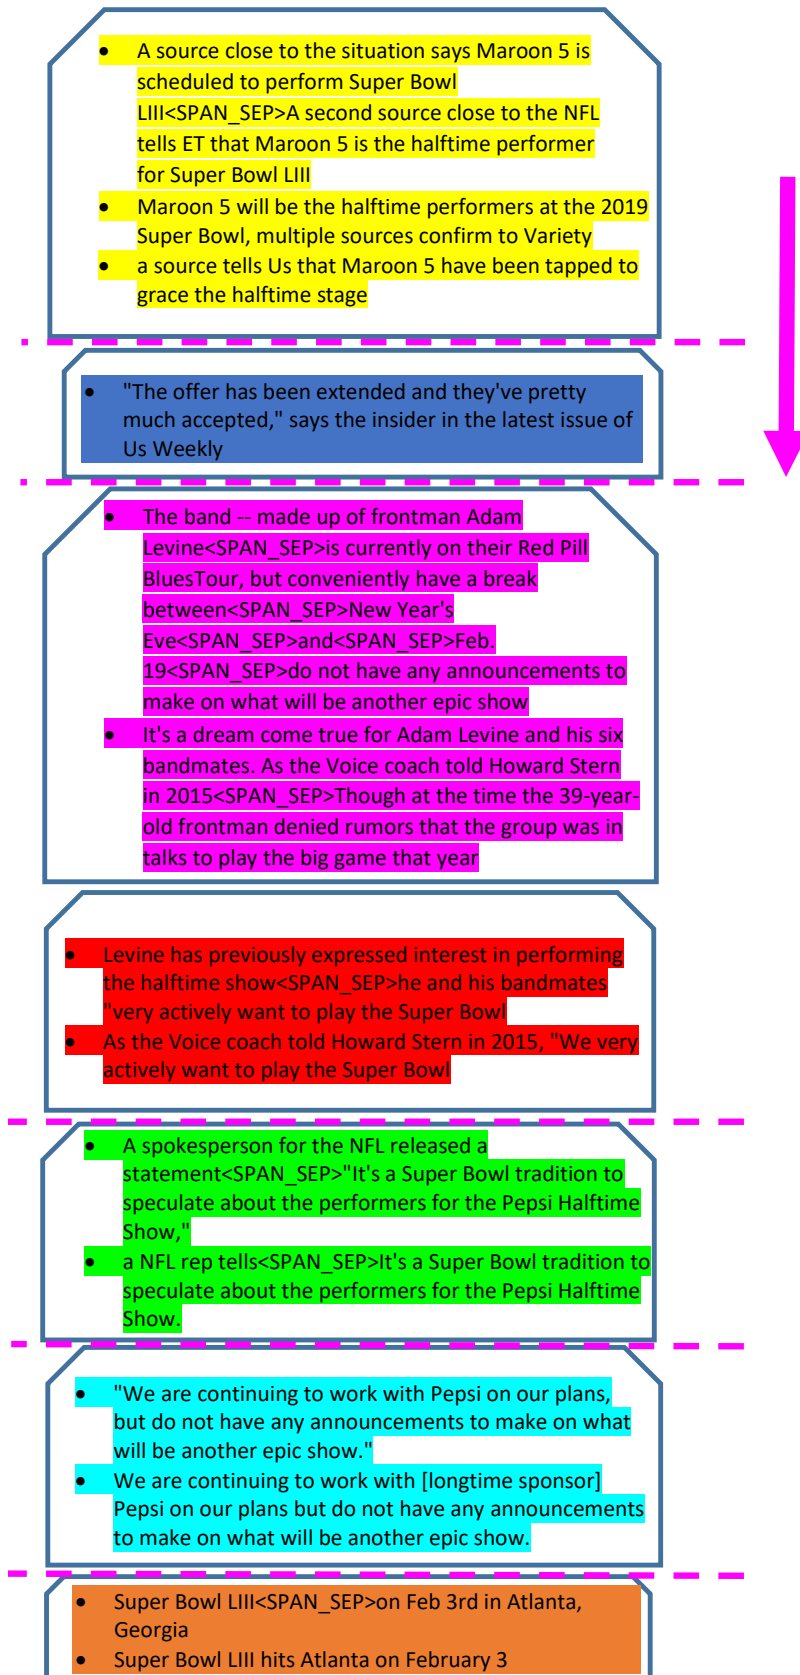


Figure 8: An example of a **Sentence & Paragraph Planning** instance derived from the alignments in Figure 4. The clusters are ordered by the order of their aligned summary propositions, and grouped with clusters with the same aligned summary sentence.

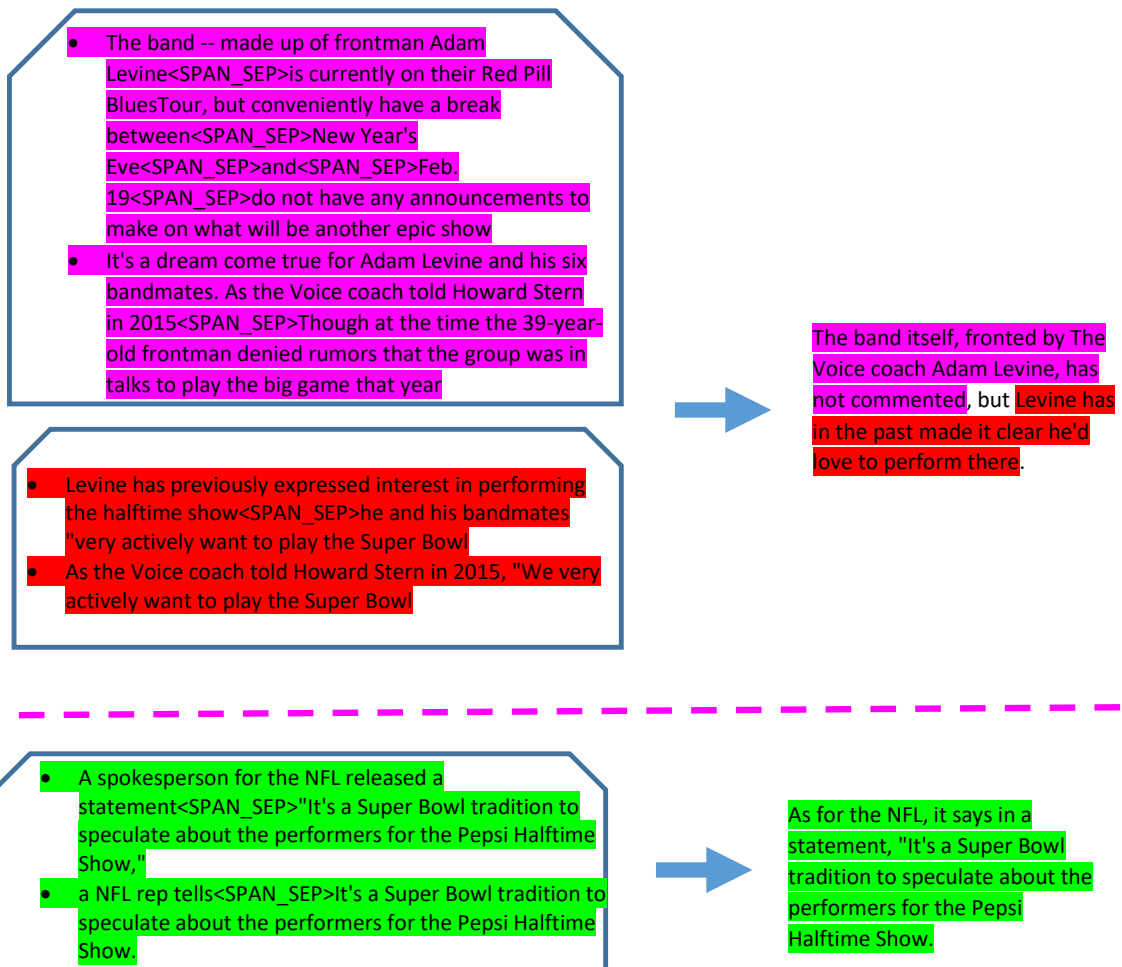


Figure 9: An example of some of the **Sentence Fusion** instances derived from the alignments in Figure 4. The clusters that are aligned to the same summary sentence should be fused to generate this sentence.