

Cross-Document Event-Keyed Summarization

William Walden¹ Pavlo Kuchmiichuk² Alexander Martin¹ Chihsheng Jin²
Angela Cao² Claire Sun² Curisia Allen² Aaron Steven White²

¹Johns Hopkins University ²University of Rochester

{wwalden1}@jhu.edu

Abstract

Event-keyed summarization (EKS) requires summarizing a specific event described in a document given the document text and an event representation extracted from it. In this work, we extend EKS to the cross-document setting (CDEKS), in which summaries must synthesize information from accounts of the same event as given by *multiple* sources. We introduce SEAMUS (Summaries of Events Across Multiple Sources), a high-quality dataset for CDEKS based on an expert reannotation of the FAMUS dataset for cross-document argument extraction. We present a suite of baselines on SEAMUS—covering both smaller, fine-tuned models, as well as zero- and few-shot prompted LLMs—along with detailed ablations and a human evaluation study, showing SEAMUS to be a valuable benchmark for this new task.

1 Introduction

Providing useful information about events requires the ability not only to extract relevant, user-specified information from documents, but also to present that information in a readable form. Drawing on this observation, Gantt et al. (2024) recently proposed *event-keyed summarization* (EKS), a task that entails summarizing a *particular* event, given a document and an event representation extracted from it. EKS thus seeks to satisfy both requirements—reconciling the *specific* information needs of IE end users with the more generic outputs of traditional summarization models—in order to communicate *precise* information about a single event in a *contextualized* and *readable* form. EKS can thus be viewed as event-centric controllable summarization (Fan et al., 2018), where the controlled attributes are the event and roles of interest.

However, adequately understanding a particular event often requires synthesizing information across *multiple* sources—evidenced in part by the rapidly growing interest in *retrieval augmented generation* (RAG; Lewis et al., 2020b). Accordingly,




Figure 1: Schematic illustration of the SEAMUS report and cross-document event-keyed summarization tasks. Letters represent event arguments.

this work extends EKS to the cross-document setting (CDEKS), drawing on—and enhancing—the FAMUS dataset for cross-document argument extraction (CDAE) to do so (Vashishtha et al., 2024). We summarize our contributions as follows:

1. We collect and release an expert reannotation of the FAMUS CDAE dataset, correcting the existing crowdsourced annotations.
2. Based on (1), we collect and release SEAMUS, an expert-annotated dataset of single- and cross-document event-keyed summaries—the first ever dataset for CDEKS.¹
3. We present a suite of baselines on SEAMUS using both smaller, fine-tuned models and prompted LLMs, showing CDEKS to be challenging relative to single-document EKS.
4. We conduct fine-grained ablations and a human evaluation, detailing CDEKS demands as a task as well as models’ current capabilities.

2 Background

FAMUS (Vashishtha et al., 2024) is a dataset of

¹<https://github.com/wgantt/SEAMUS>

short English Wikipedia passages (*reports*) paired with much longer, genre-diverse English *source* documents cited by those reports.² FAMUS supports two tasks: (1) *Source Validation* (SV), where the goal is to determine whether a candidate source document is *valid for*—i.e. describes the same event as—an event identified in a provided report; and (2) *Cross-Document Argument Extraction* (CDAE), which entails extracting arguments for an identified event from *both* the report and a valid source document. SEAMUS builds on the FAMUS CDAE data, which contains 1,265 report-source document pairs (split 3:1:1 across train, dev, and test), and annotates arguments of the same target event for each document in a pair using a subset of the FrameNet ontology restricted to frames denoting events, states, or processes (Baker et al., 1998). A single, maximally “informative” mention is annotated for each argument, where proper names > nominal expressions > pronouns (see Li et al., 2021b). In both report and source texts, arguments may be distributed across sentences.

Event-Centric Summarization In introducing EKS, Gantt et al. (2024) released MUCSUM, an EKS dataset based on the classic MUC-4 template filling dataset (Sundheim, 1992). MUCSUM contains abstractive event-keyed summaries for each event template in MUC-4, written so as to faithfully express the role of each template argument, plus any minimal additional context required for the summary to act as a standalone account of the event. Gantt et al. present baselines on MUCSUM, and also conduct a human evaluation of model outputs, which inspires our own (§5).

Other event-centric summarization research has focused on *timeline summarization* (TLS), which constructs chronological lists of events, often with timestamps and usually based on multiple documents (Allan et al., 2001; Chieu and Lee, 2004; Li et al., 2021a; Rajaby Faghihi et al., 2022, *i.a.*). Beyond TLS, S Hussain et al. (2022) use extracted event-related keywords to condition single-document summarization, and integrate an event-oriented attention mechanism into BART to encourage models to cover *all* events discussed. Additionally, Vallurupalli et al. (2022) introduce the POQue dataset, which has annotations that characterize the subevent structure of complex events in stories and the changes undergone by their participants. Among these annotations are *process summaries*,

²All documents are from MegaWika (Barham et al., 2023).

which give high-level descriptions of a complex event, and *change summaries*, which describe the changes experienced by a participant as a result.

Multi-Document Summarization CDEKS is an event-centric multi-document summarization (MDS) task. Work on MDS has pursued a variety of goals, including synthesizing reviews (Ganesan et al., 2010; Chu and Liu, 2019, *i.a.*), summarizing dialogues (Kraaij et al., 2005; Chen et al., 2021, *i.a.*), distilling news articles (notably, via DUC³ and TAC⁴), and generating reports (Mayfield et al., 2024). *Event-centric* MDS datasets include MultiNews (Fabbri et al., 2019) and DiverseSumm (Huang et al., 2024), which focus on new stories, but SEAMUS is most similar to AutohMDS (Zopf, 2018) and WCEP (Gholipour Ghandari et al., 2020) in being built on Wikipedia articles and their sources.

CDEKS departs from all of these, however, in *responding to an explicit information need*. It is thus an event-centric form of *query-oriented* MDS (Ma et al., 2020), where a query expressing the kind of information to be summarized is provided as additional input. But whereas queries from prior work are given in natural language—e.g. article titles (Liu and Lapata, 2019) or web searches (Pasunuru et al., 2021)—ours are structured event representations, drawing on the IE tradition of leveraging event ontologies to encode information needs, and enabling extraction-to-summarization pipelines.

Our Work We summarize three key differences between prior work and our own. We focus on:

1. Synthesizing information about a *single* event across *multiple* sources. Both multi-event (e.g. TLS) and single-source (e.g. EKS) summarization have their place, but many practical information needs depend on the rich understanding of an *individual* event that is attainable only via *cross-source* synthesis.
2. Responding to a *specific* event-centric information need, not *generically* summarizing event-related content (*contra* S Hussain et al., 2022; Vallurupalli et al., 2022).
3. Leveraging detailed, structured event representations to achieve (1) and (2)—not short, unstructured queries like web searches (Pasunuru et al., 2021) or topics (Allan et al., 2001; Rajaby Faghihi et al., 2022, *i.a.*).

³<https://duc.nist.gov/>

⁴<https://tac.nist.gov/publications/index.html>

3 Annotation

Annotation of SEAMUS was divided into two phases. In the first phase, abstractive **report summaries** were written for each event in FAMUS (see §2) based only on its *report* document, and were then annotated for event arguments (§3.1). In the second phase, abstractive **cross-document summaries** were written for each event based *jointly* on its report and source documents, and were then annotated for event arguments as in the first phase (§3.2). In both phases, annotators were instructed to amend spurious, missing, or otherwise incorrect argument annotations in the report or source document before writing their summary. Thus, both phases involve (1) correcting existing FAMUS argument annotations; (2) writing a summary based on the corrected annotations; and (3) annotating arguments in the summary. The phases differ only in the documents on which the summaries are based (report only vs. report and source). All annotations were performed by authors of this work.⁵

3.1 Phase 1: Report Summaries

Similar to the summaries in MUCSUM (§2), the report summaries in SEAMUS are concise summaries of a single event as recounted in a single document (a FAMUS report) that aim to faithfully represent the role of each participant and to provide the minimum additional context needed to serve as an accurate, standalone account of the event. Although the FAMUS report documents are already relatively short (typically, 2-3 sentences), they often discuss multiple events.⁶ Thus, the report summaries are further distilled descriptions focused on just *one* event from the report.

Three authors completed the Phase 1 annotation, with each summary and its arguments singly annotated. Items from the train split were randomly and evenly divided among these three authors; items from the dev and test splits were similarly divided between two of them. All items were provided in JSON files containing the following information for each example: (1) a unique example ID, (2) the FAMUS report text; (3) the FAMUS-annotated frame, trigger, and arguments of the target event from the report; and (4) definitions of the annotated frame and roles as given in FrameNet. Annotators

⁵Appendix E has additional details and agreement results.

⁶E.g. for reports in the SEAMUS train split, the MegaWika dataset (Barham et al., 2023), from which the reports are taken, has an average of 21.4 FrameNet frames annotated.

	Report		Cross-Doc	
	Train	Dev	Train	Dev
Examples	759	253	759	253
Avg. Words	21.8	24.6	30.5	34.5
Avg. Sentences	1.0	1.0	1.2	1.2
Avg. Arguments	3.1	3.5	4.1	4.6

Table 1: Summary statistics for the SEAMUS report and cross-document summaries. See Table 7 for more.

were provided with detailed instructions written by the first author and completed a 10-example practice task before beginning the main annotation. Consistent with FAMUS, both the corrected report arguments and the report summary arguments were annotated as single, maximally informative mentions (see §2). Annotators were encouraged to use the same mentions in their summaries as were annotated in the (corrected) report arguments, but were permitted to alter them in the summary in order to preserve clarity or naturalness. Annotations were validated to ensure that (1) they were shorter than the report they summarized and (2) the number of arguments for a given role matched between each report and its summary. All initially invalid annotations were then corrected.

3.2 Phase 2: Cross-Document Summaries

The *cross-document summaries* are intended as enriched versions of the report summaries, synthesizing details about the target event from both the report and the target event’s source document.

Five of the authors completed the Phase 2 annotation, with all summaries and arguments singly annotated as in Phase 1. Items from all three splits were randomly and evenly distributed to the five annotators. Given the complexity of the Phase 2 task, annotation was performed in two parts, using adapted versions of Vashishtha et al.’s (2024) interface for FAMUS CDAE annotation.⁷

In Part A, annotators corrected FAMUS argument annotations in the source documents and then wrote the cross-document summary based *jointly* on the report and source texts and their corrected arguments. Annotators were encouraged to use the most informative mention of an argument across *both* the report and source documents, but again were allowed to make alterations for clarity.

In Part B, annotators annotated arguments in the

⁷Interface source code was obtained from Vashishtha et al. Screenshots are shown in Appendix E.

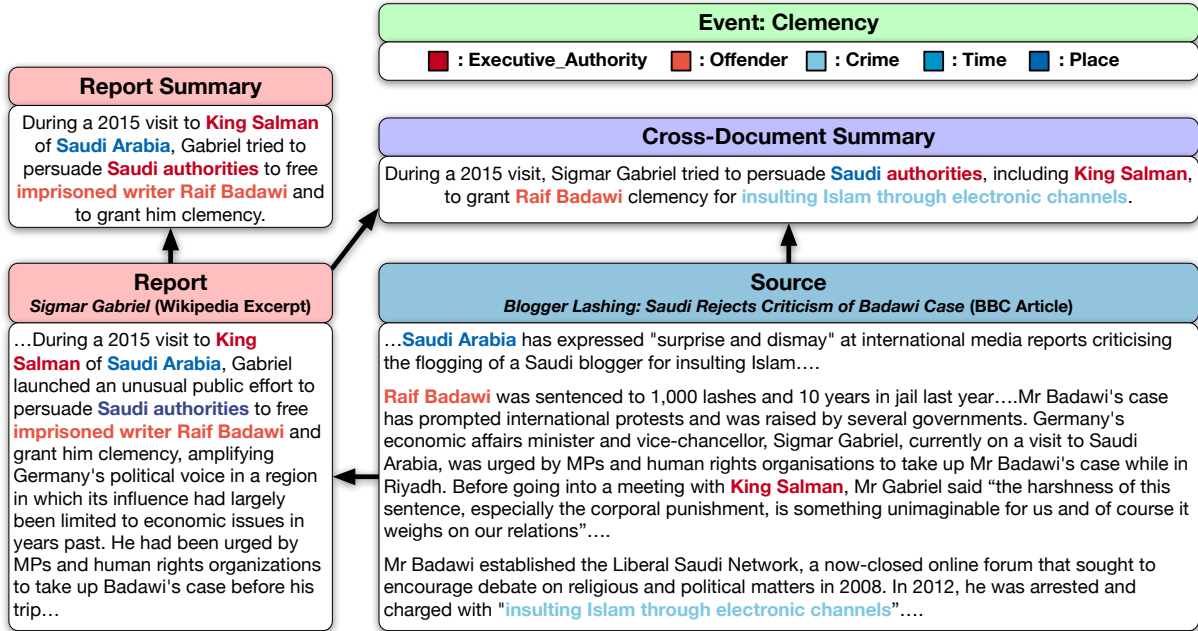


Figure 2: An example from our SEAMUS dataset. **Report** documents (bottom left) are Wikipedia passages that describe some event (top right) and that cite a longer (non-Wikipedia) **source** article (bottom right) as evidence, with event arguments annotated in both documents. SEAMUS features simple summaries of these events based on *only* the report (top left) as well as enriched, cross-document summaries based on *both* the report and its source, which typically contain additional information about the event (here, the CRIME). [Appendix A](#) has further examples.

summaries from Part A. As in Phase 1, all annotators were provided with detailed instructions and completed a 10-example practice annotation before doing the main task. Summary argument annotations were again validated for length and to ensure that they featured as many arguments for a given role as the maximum number annotated for that role between the report and source texts.

Summary statistics for both the report and cross-document summaries can be found in [Table 1](#) and an example is shown in [Figure 2](#). Both types of summary average roughly a sentence in length, though cross-document summaries tend to be longer and to have more arguments—consistent with the richer information they provide.

4 Experiments

4.1 Overview

Tasks We present experiments on both the report (§4.2) and cross-document (§4.3) summarization tasks. In the report task (single-document EKS), both the report and its annotated event are provided as input. The cross-document task (CDEKS) is analogous, but also includes the corresponding source document and its event annotation as input. Next, in §4.4, we briefly discuss some ablations on the input inspired by similar ones from [Gantt et al.](#)

(2024), with full results in [Appendix F](#). Finally, §4.5 evaluates the impact of degraded argument extractions on summary quality.

Models We benchmark SEAMUS using models of two types. First, we consider several classic pre-trained encoder-decoder models widely used for summarization: BART ([Lewis et al., 2020a](#)), PE-GASUS ([Zhang et al., 2020](#)), and T5 ([Raffel et al., 2020](#)), fine-tuning the large versions of all three on the SEAMUS training data. Second, we consider some of the latest proprietary LLMs, evaluated in both the zero- and few-shot settings: GPT-4o⁸, GPT-4o Mini (GPT-4o M in [Table 2](#))⁹, Claude 3 Haiku (CLAUDE H)¹⁰, and Claude 3.5 Sonnet (CLAUDE S)¹¹. For the few-shot examples, we use the three examples from the train split whose frame matches that of the target example. Finally, we also give results for a *report baseline* (RB) that treats the report text itself as the predicted summary.

Metrics We report several standard summarization metrics, including ROUGE-1 (\mathbf{R}_1), ROUGE-2 (\mathbf{R}_2), and ROUGE-LCS F_1 scores (\mathbf{R}_L ; [Lin, 2004](#)),

⁸<https://openai.com/index/hello-gpt-4o/>

⁹<https://openai.com/index/>

[gpt-4o-mini-advancing-cost-efficient-intelligence/](https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/)

¹⁰<https://www.anthropic.com/news/claude-3-haiku>

¹¹<https://www.anthropic.com/news/claude-3-5-sonnet>

Model	S	Report							Cross-Document						
		R ₁	R ₂	R _L	BS	CR	A	F	R ₁	R ₂	R _L	BS	CR	A	F
RB	-	56.2	46.1	48.4	91.6	52.6	99.1	98.7	48.5	33.3	39.3	89.6	31.0	99.3	93.1
GPT-4o M	ZS	62.2	42.3	51.3	93.2	58.5	86.0	75.8	51.8	29.9	39.0	91.3	39.0	81.5	88.9
	FS	72.0	55.4	61.0	94.3	66.8	94.1	83.3	57.5	36.9	45.7	92.1	39.8	88.5	89.8
GPT-4o	ZS	64.0	45.2	53.0	93.2	61.4*	83.9	74.8	58.0*	36.4	45.8	92.2*	41.3*	86.6	88.4
	FS	72.5 [†]	56.6 [†]	62.3 [†]	94.4	69.6 [†]	94.7	81.6	61.2 [†]	40.7 [†]	49.4 [†]	92.7[†]	42.7 [†]	90.6	88.5
CLAUDE H	ZS	64.8	46.2	54.7	93.4	58.8	84.9	77.6	57.7	36.9*	46.5	92.1	36.2	90.4	91.4
	FS	71.7	55.9	61.1	94.3	63.2	94.8	82.5	59.4	39.5	48.6	92.1	37.2	91.0	90.5 [†]
CLAUDE S	ZS	67.4*	48.1*	56.5*	93.8*	61.1	93.0*	80.6*	56.7	34.8	45.3	91.9	35.2	93.4*	91.7*
	FS	72.2	54.6	61.3	94.5 [†]	65.7	95.9 [†]	83.9 [†]	57.9	38.1	47.4	92.1	37.3	95.1[†]	90.4
BART	FT	74.5	61.7	66.4	94.6	69.9	91.6	79.3	63.8	45.5	53.0	92.6	45.0	85.6	85.3
PEGASUS	FT	75.2	62.5	67.0	94.7	70.0	96.1	82.2	63.7	46.2	53.2	92.5	43.7	93.9	90.5
T5	FT	76.6	64.4	68.9	95.0	74.2	98.2	85.0	64.1	46.4	52.8	92.6	44.7	92.5	90.2

Table 2: **Report** and **Cross-Document** summarization results on SEAMUS. Best overall results are **bolded**; * and [†] denote best zero- and few-shot results, respectively. S=setting; RB=report baseline; ZS=zero-shot; FS=few-shot; FT=fine-tuned. See §4.1 for an explanation of metrics; higher is better for all. See Tables 10 and 11 for 95% CIs. Best **A** and **F** results exclude RB, for reasons explained in Appendix F.

as well as BERTScore F₁ (**BS**; Zhang et al., 2019).

Given EKS’s focus on producing summaries that recover *specific* pieces of information—as represented by an event’s roles—we report several other metrics that evaluate this. First, we report CEAF-REE F₁ (**CR**; Du et al., 2021a), a form of argument F₁ that allows us to compare arguments extracted from a *predicted* summary against those in a *reference* summary, aligning arguments based on exact match.¹² Following Gantt et al. (2024), we train the event extraction model of Xia et al. (2021)¹³ on SEAMUS and use it to extract arguments from the predicted summaries, constraining extraction to arguments that fill roles of the target event only.

The summaries in SEAMUS also make *claims* about these arguments that reflect their role in the target event. To evaluate these claims’ fidelity to the text, we report AlignScore (**A**; Zha et al., 2023), a learned metric that provides a score in [0, 1] that indicates how well a claim (here, a summary) is supported by a given context (the report for the report task, and the concatenated report and source for the cross-document task). We also report FACTSCORE (**F**; Min et al., 2023), which uses LMs to (1) decompose a generation into a set of *atomic* facts, and (2) determine the % of these facts supported by a given knowledge source, where **F** is the average % supported over all examples. We use as knowledge sources the contexts used for **A**.

4.2 Report Summarization

Setup As input for BART, PEGASUS, and T5, we provide the full report text concatenated with a linearized representation of the annotated report

event that contains the frame name, the event trigger, and the role names, each followed by a list of the arguments annotated for that role. We train each model against a standard conditional language modeling objective w.r.t. the gold report summaries for a maximum of 30 epochs, using a patience of 5 epochs, with dev **R**₁ as the stopping criterion.¹⁴ For inference, we use beam search decoding with a beam size of 5 and a max of 256 new tokens.

For the Claude and GPT models, our system prompt asks the model to analyze and summarize a specific event. The user prompt provides more detailed task instructions, followed by the full report text, and a description of the target event that includes (1) the frame name and definition from FrameNet; (2) the trigger; and (3) a bulleted list, where each item includes a role name, its definition, and the arguments annotated for that role. In the few-shot setting, we format the three few-shot examples (see §4.1) the same way, but with the target summary shown at the end of each. We set temperature to 0.7 and the max new tokens to 256, leaving other API defaults unchanged.¹⁵

Results are shown in the left half of Table 2. First, we find that T5 obtains the best performance across all metrics, followed by PEGASUS and BART, with T5 exhibiting particularly strong results for **CR**, indicating its ability to accurately recover event arguments in its summaries. Second, the LLMs almost universally outperform the report baseline (RB)—even in the zero-shot setting (ZS), where Claude Sonnet generally obtains the best results. Third, adding just three few-shot examples

¹²Appendix F reports a soft match variant of this metric.

¹³<https://hub.docker.com/r/hltcoe/lome>

¹⁴Details on training and input formats are in Appendix B.

¹⁵Appendix C has further details on models and prompts.

Model	p	Report							Cross-Document						
		R_1	R_2	R_L	BS	CR	A	F	R_1	R_2	R_L	BS	CR	A	F
T5	0.0	76.6	64.4	68.9	95.0	74.2	98.2	85.0	64.1	46.4	52.8	92.6	46.3	92.5	90.2
	0.1	75.6	62.8	67.8	93.9	71.4	97.6	84.7	62.8	45.3	51.8	91.5	47.2	92.0	89.9
	0.2	74.0	61.7	66.2	93.6	69.6	98.0	84.6	62.0	44.3	50.7	91.4	43.5	89.3	88.0
	0.3	72.1	60.0	64.7	93.3	67.5	98.2	83.0	60.0	42.8	49.3	91.0	43.3	87.3	89.0
	0.4	70.3	57.5	62.1	92.9	66.4	95.8	83.2	58.4	40.9	47.8	90.8	44.4	87.4	86.8
	0.5	68.3	55.2	60.6	92.6	63.2	96.3	83.5	56.6	39.1	46.3	90.4	43.1	87.3	88.0
CLAUDE H (FS)	0.0	71.7	55.9	61.0	94.3	63.2	94.8	82.9	57.7	36.9	45.7	92.1	36.2	91.0	90.5
	0.1	67.5	51.5	56.7	93.7	59.4	94.8	83.6	57.2	37.3	45.4	91.8	37.1	82.5	88.9
	0.2	65.6	48.8	55.1	93.5	55.1	94.7	83.2	56.2	37.0	45.1	91.7	37.8	79.4	88.6
	0.3	64.7	47.8	54.1	93.3	52.8	94.6	84.1	56.0	36.2	44.9	91.5	32.7	82.2	89.2
	0.4	64.1	47.2	54.1	93.3	52.2	95.0	83.1	54.5	34.3	43.1	91.3	31.4	85.0	89.0
	0.5	63.1	46.8	54.0	93.1	52.3	94.7	83.8	54.3	34.6	43.3	91.3	33.1	86.4	89.2
GPT-4o M (FS)	0.0	72.0	55.4	61.0	94.3	66.8	94.1	83.3	57.5	36.9	45.7	92.1	39.8	88.5	89.8
	0.1	69.2	52.8	59.5	94.0	64.0	94.5	81.8	58.8	38.2	46.2	92.1	42.2	74.5	90.6
	0.2	67.6	50.8	57.0	93.7	59.8	94.2	84.3	56.6	36.2	45.1	91.2	39.4	75.2	89.8
	0.3	66.9	50.1	57.0	93.7	59.3	94.9	81.8	56.4	36.2	44.5	91.8	37.8	77.2	90.2
	0.4	65.2	48.1	54.9	93.4	56.7	93.8	84.2	54.8	34.0	42.8	91.6	36.6	77.8	90.6
	0.5	65.1	47.5	54.8	93.4	55.2	95.4	82.7	54.2	33.2	42.6	91.4	34.4	80.9	90.8

Table 3: Performance of three models from Table 2 when the argument annotations for each role in the report event (**Report**) or additionally in the source event (**Cross-Document**) are corrupted with probability p (see §4.5).

(FS) yields major gains over the zero-shot setting for all LLMs on all metrics. Even here, however, few-shot results still trail the best fine-tuned results (T5) by sizable margins on most metrics.

4.3 Cross-Document Summarization

Setup The setup for the cross-document task is similar to that of the report task, but adds the source text and its annotated event to the input alongside the report text and its event. As the source texts are full web articles, most are long (e.g. dev texts average almost 62 sentences and over 1,500 words). While this is no obstacle for the LLMs, the smaller models do not support contexts of this size. Thus, to enable a fair comparison across models, we apply a sentence retriever to the source, using the report text as a query to select the top k most relevant sentences to use as context.¹⁶ We consider $k \in \{3, \dots, 10\}$ and selected the maximum value such that $\geq 95\%$ of the resulting dev set contexts would fit untruncated in the input, yielding $k = 7$. We experimented with the dense retrievers `all-mpnet-base-v2` (based on MPNet; Song et al., 2020) and `e5-large-v2` (Wang et al., 2022), but obtained our best results with BM25 (Robertson et al., 2009), which we use in all experiments.¹⁷

We use the same training and inference settings from §4.2; see Appendices B, C for further details.

¹⁶This approach can also be justified by the fact that typically only a small portion of the source concerns the event.

¹⁷Models were evaluated on recall of annotated arguments in the retrieved contexts for the dev set for fixed k . At $k = 7$, BM25 recovered $\sim 76\%$ of annotated source arguments.

Results are shown in the right half of Table 2 and are qualitatively similar to those for the report task, with the fine-tuned models generally showing the best overall numbers ($R_{1,2,L}$, CR) or nearly so (BS), although GPT-4o obtains the highest scores on BS and Claude Sonnet on A and F. Once again, nearly all models outperform the report baseline across the board (ZS and GPT-4o Mini excepted). Finally, we note that results on most metrics are much lower in absolute terms compared to the corresponding results from §4.2, testifying to the greater difficulty of the cross-document task.

4.4 Input Ablations

Following Gantt et al. (2024), Appendix F considers ablations on the input for both tasks, in which we omit the annotated events (TEXT ONLY) or the texts (EVENT ONLY), and condition summary generation on the resulting ablated inputs. We also present a novel third ablation that omits the *arguments*, but leaves in information about the frame and roles (TEXT+SCHEMA). Consistent with Gantt et al., we find that *both the text and the full event annotations are needed to obtain the best results* (Tables 8 and 9), indicating that the SEAMUS tasks are *not* reducible to standard summarization (TEXT ONLY), structure-to-text (EVENT ONLY), or even a hybrid objective (TEXT+SCHEMA). Moreover, if some results in Table 2 (e.g. $R_{\{1,2,L\}}$) appear high, these ablations show that this is due in large part to access to gold event structures. We next turn to the case of imperfect event extractions.

4.5 Impact of Extraction Quality

Setup §4.2 and §4.3 use gold event structures in the input to facilitate fair cross-model comparisons. But in real-world scenarios, one rarely has access to gold arguments, and certainly not in the extraction-to-summarization pipelines CDEKS aims to support. It is thus essential to understand models’ tolerance to noise in the extracted events.

To probe robustness to extraction errors in a controlled manner, we apply variable amounts of noise to the gold event annotations and evaluate model performance on the resulting inputs. Concretely, for each role R of each event, we edit R ’s arguments with probability p . If a role is selected for editing, we then make *one* of the following edits with equal probability:

1. INSERT: A new (incorrect) span from the text is *added* to the argument list for R .
2. DELETE: An argument span is *removed* at random from the argument list for R .
3. REPLACE: An argument span is *replaced* at random with an incorrect span from the text.

For the cross-document task, we apply these edits to the event annotations for both the report and the source. We sample the edits to be made uniformly and then prompt an LLM (GPT-4o) to apply them by supplying in a prompt: (1) the text (report or source), (2) the (JSON-formatted) report or source event annotations, and (3) instructions for the edits to be made, generated automatically by populating templatic statements based on the edits sampled. The LLM is free to select an appropriate *new* span to be used for the INSERT and REPLACE operations. We consider $p \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$, using the same sampled edits for all models for a given p .

We evaluate one fine-tuned model (T5) and one model each from the Claude (Claude Haiku) and GPT (GPT-4o Mini) families. For T5, we use the same checkpoint as presented in Table 2. For Claude Haiku and GPT-4o Mini, we use few-shot prompts similar to those used in the FS setting in Table 2, but with two key changes. First, we alter the task instructions to say that the event annotations for the target example *may contain errors*, and that the model must correct these errors when generating its summary by consulting the text(s). Second, we show the model how to do this by substituting noised versions of the event annotations in the few-shot examples while leaving their associated texts and summaries unchanged.

p	Summary
0.0	The gradual accumulation of partially decayed plant material in a bog functions as a carbon sink.
0.1	The gradual accumulation of decayed plant material in a bog functions as a carbon sink.
0.2	The gradual accumulation of decayed plant material in a bog acts as a carbon sink.
0.3	The gradual accumulation of decayed plant material in a bog functions as a carbon sink.
0.4	The gradual accumulation of decayed plant material, including peat, in bogs functions as a carbon sink.
0.5	The gradual accumulation of decayed plant material in a bog functions as a carbon sink.

Table 4: Example outputs from GPT-4o Mini on the cross-document task as role annotations are corrupted with probability p . In many cases (as here), we find minimal degradation in quality from $p = 0$ to $p = 0.5$.

Results for both tasks are in Table 3. For all models, we observe (near-)monotonic drops in performance for most metrics as p increases. While performance drops are sizable in some cases, they are arguably less radical than we might expect, given the destructiveness of the changes at $p = 0.5$, where roughly half of all roles contain extraction errors. This is especially evident in the results for Claude Haiku and GPT-4o Mini on the cross-document task, where (e.g.) $\mathbf{R}_{1,2,L}$ scores decrease by only about 3 points from $p = 0$ to $p = 0.5$, \mathbf{BS} by less than 1, and \mathbf{F} showing no drop at all. Further, losses on \mathbf{CR} (the most explicit measure of extraction ability) are only ~ 5 points for GPT-4o Mini and ~ 3 points for Claude Haiku.

These findings are confirmed by manual inspection of model outputs, where we often see relatively little degradation in summary quality (Table 4). This suggests an intriguing strength of this task relative to traditional event extraction: the ability to *counteract extraction errors post-hoc* by using imperfect event extractions as a query to locate relevant passages in the input and then leveraging those passages to avoid analogous errors in the summary.

5 Human Evaluation

Setup Lastly, we conduct a human evaluation of the reference and model-generated summaries. We focus our evaluation on the cross-document task, comparing the summaries generated by models presented in Table 2 (excluding RB). For the GPT and Claude models, we use the FS (few-shot) summaries only, owing to their superiority over the ZS results. We randomly sampled 30 test set examples and presented the 7 model-generated summaries for these examples, along with the references, to 3 human raters—all English-speaking NLP researchers who did not participate in other




Figure 3: Histograms of summary quality scores (1-5, higher is better) from our human evaluation (§5). The bottom right plot (red) aggregates scores across all three raters; each of the other plots (blue) shows a single rater’s scores.

parts of this work. Each rater provided a single quality score for each summary based on the ordered list of attributes used by Gantt et al. (2024): *factuality*, *adequacy*, *coherence*, *relevancy*, and *fluency*. Scores were given from 1 (low) to 5 (high), with half points allowed. Each rater thus provided $30 \times (7 + 1) \times 1 = 240$ judgments. Raters were not shown which model produced which summary, and summary presentation order was randomized.¹⁸

Results Four sets of histograms of scores for each model (and the reference) are shown in Figure 3. The bottom right set (red) shows scores aggregated across annotators, while the other three (blue) each show scores of a single rater. For all raters, scores are consistently high across models and the reference, with modes of ≥ 4 for each. Comparing preferences across raters, however, we see significant variability: GPT-4o achieved the highest average score for one rater (top right, 4.28); GPT-4o Mini for the second (bottom left, 4.57); and PEGASUS for the third (top left, 4.23).

Looking at intra-rater distributions, however, it’s unclear how robust these preferences are. Using Wilcoxon rank-sum tests to evaluate pairwise differences in each rater’s scores for a given pair of models, we find that some of these preferences are reliable at $\alpha = .05$ (e.g. GPT-4o $>$ T5 with $p = .016$ for the first rater), but none holds up when applying the Bonferroni correction for multiple comparisons. We take these results to indicate that our baselines are fairly effective at producing

good summaries, and that while they may somewhat differentiate themselves on individual metrics¹⁹, the best models on a more holistic picture may come down to user preference, and there may not be *definitive* bests even at this scope. This plurality of solid modeling options is encouraging, and suggests flexibility in the application of CDEKS to a range of use cases.

6 Conclusion

This work has extended the task of *event-keyed summarization* (EKS) to the cross-document setting (CDEKS). To enable this, we provided an expert reannotation of the FAMUS CDAE dataset, yielding high-quality event argument annotations on all 1,265 examples. We then leveraged these improved annotations to construct SEAMUS—a collection of single- (*report*) and cross-document summaries on top of FAMUS, further annotating the summaries themselves for event arguments (§3). We benchmarked SEAMUS on a diverse set of baselines, including smaller fine-tuned models, as well as zero- and few-shot prompted LLMs (§4.2, §4.3). We then presented more detailed analysis, conducting a comprehensive set of input ablations (§4.4), assessing the impact of degraded event extraction on summary quality (§4.5), and finally concluding with a human evaluation of summary quality (§5). We release SEAMUS, along with our baseline results, to facilitate further work on EKS in both the single- and cross-document settings.

¹⁸See Appendix D for further details.

¹⁹See our discussion of argument recovery in Appendix F.

Limitations

One limitation of this work is SEAMUS’s size: 1,265 examples is sufficient for fine-tuning smaller models and for conducting prompting experiments with larger ones, but is likely insufficient for substantive fine-tuning of very large models.

A second limitation is that the cross-document setting considers only two documents per example. This constraint was imposed by the choice of the FAMuS dataset as the basis for SEAMUS, as cross-document argument annotations in the former were provided only for *pairs* of report and source texts. Future work expanding the set of source texts would be valuable, and would allow both for richer summaries and for more robust evaluation of models’ ability to accurately synthesize information across possibly differing accounts of events (cf. [Huang et al. \(2024\)](#)), as information conflicts are more common outside of Wikipedia citations.

We note, however, that addressing either limitation may require relaxing data quality standards—relying on crowdsourcing or LLM-powered annotation techniques—as scaling our annotation procedure to many more examples or source documents would demand considerable resources. It was only thanks to the above restrictions that we were able to provide expert annotations for SEAMUS.

Finally, while the experiments in §4.5 offer a helpful picture of the impact of event extraction quality on cross-document event-keyed summaries, further experiments on the outputs of actual event extraction systems would likely provide a better one. Much of the difficulty of deploying CDEKS in practical settings undoubtedly lies in the development of an effective document-level event extractor, as evidenced by the ongoing challenges documented by much prior work in this domain ([Du et al., 2021b](#); [Chen et al., 2023b](#); [Gantt et al., 2023](#); [Vashishtha et al., 2024, i.a.](#))

Ethics

As the report and source texts in SEAMUS are the same as those in the FAMuS dataset, and as the summaries in SEAMUS are simply distillations of (parts of) these texts, we do not believe our dataset introduces any novel risks as a resource. Nonetheless, these texts do discuss real people, places, and institutions, and models trained on this data may thus be liable to make untrue claims about them or otherwise misrepresent them. We intend SEAMUS for academic use only, as a benchmark

to evaluate systems for single- and cross-document event-keyed summarization.

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 information retrieval*, pages 10–18.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. [The Berkeley FrameNet project](#). In *COLING 1998 Volume 1: The 17th International Conference on Computational Linguistics*.
- Samuel Barham, Orion Weller, Michelle Yuan, Kenton Murray, Mahsa Yarmohammadi, Zhengping Jiang, Siddharth Vashishtha, Alexander Martin, Anqi Liu, Aaron Steven White, et al. 2023. Megawika: Millions of reports and their sources across 50 diverse languages. *arXiv preprint arXiv:2307.07049*.
- Yanran Chen and Steffen Eger. 2023. [MENLI: Robust evaluation metrics from natural language inference](#). *Transactions of the Association for Computational Linguistics*, 11:804–825.
- Yulong Chen, Yang Liu, Liang Chen, and Yue Zhang. 2021. [DialogSum: A real-life scenario dialogue summarization dataset](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 5062–5074. Association for Computational Linguistics.
- Yunmo Chen, William Gantt, Tongfei Chen, Aaron White, and Benjamin Van Durme. 2023a. [A unified view of evaluation metrics for structured prediction](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12868–12882. Association for Computational Linguistics.
- Yunmo Chen, William Gantt, Weiwei Gu, Tongfei Chen, Aaron White, and Benjamin Van Durme. 2023b. [Iterative document-level information extraction via imitation learning](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1858–1874. Association for Computational Linguistics.
- 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*, pages 425–432.
- Eric Chu and Peter Liu. 2019. Meansum: a neural model for unsupervised multi-document abstractive summarization. In *International conference on machine learning*, pages 1223–1232. PMLR.
- Xinya Du, Alexander Rush, and Claire Cardie. 2021a. [GRIT: Generative role-filler transformers for document-level event entity extraction](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 634–644. Association for Computational Linguistics.
- Xinya Du, Alexander Rush, and Claire Cardie. 2021b. [Template filling with generative transformers](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 909–914. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Al-lonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stéphane Collot, Suchin Gu-

rurangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baeovski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal Chawla, Kushal Lakhota, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Maria Tsim-

poukelli, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratan-chandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Ruty Rinott, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vitor Albiero, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihalescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaofang Wang, Xiaojian Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. 2024. [The llama 3 herd of models](#).

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. Association for Computational Linguistics.

Angela Fan, David Grangier, and Michael Auli. 2018. [Controllable abstractive summarization](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 45–54. Association for Computational Linguistics.

Kavita Ganesan, ChengXiang Zhai, and Jiawei Han. 2010. [Opinosis: A graph based approach to abstractive summarization of highly redundant opinions](#). In

- Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 340–348. Coling 2010 Organizing Committee.
- William Gantt, Reno Kriz, Yunmo Chen, Siddharth Vashishtha, and Aaron White. 2023. [On event individuation for document-level information extraction](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12938–12958. Association for Computational Linguistics.
- William Gantt, Alexander Martin, Pavlo Kuchmiichuk, and Aaron Steven White. 2024. Event-keyed summarization. *arXiv preprint arXiv:2402.06973*.
- Demian Gholipour Ghalandari, Chris Hokamp, Nghia The Pham, John Glover, and Georgiana Ifrim. 2020. [A large-scale multi-document summarization dataset from the Wikipedia current events portal](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1302–1308. Association for Computational Linguistics.
- Kung-Hsiang Huang, Philippe Laban, Alexander Fabri, Prafulla Kumar Choubey, Shafiq Joty, Caiming Xiong, and Chien-Sheng Wu. 2024. [Embrace divergence for richer insights: A multi-document summarization benchmark and a case study on summarizing diverse information from news articles](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 570–593. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.
- Wessel Kraaij, Thomas Hain, Mike Lincoln, and Wilfried Post. 2005. The ami meeting corpus. In *Proc. International Conference on Methods and Techniques in Behavioral Research*, pages 1–4.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020a. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020b. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Manling Li, Tengfei Ma, Mo Yu, Lingfei Wu, Tian Gao, Heng Ji, and Kathleen McKeown. 2021a. [Timeline summarization based on event graph compression via time-aware optimal transport](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6443–6456. Association for Computational Linguistics.
- Sha Li, Heng Ji, and Jiawei Han. 2021b. [Document-level event argument extraction by conditional generation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 894–908. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81. Association for Computational Linguistics.
- Yang Liu and Mirella Lapata. 2019. [Hierarchical transformers for multi-document summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5070–5081. Association for Computational Linguistics.
- Weicheng Ma, RuiBo Liu, Lili Wang, and Soroush Vosoughi. 2020. [Multi-resolution annotations for emoji prediction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6684–6694. Association for Computational Linguistics.
- James Mayfield, Eugene Yang, Dawn Lawrie, Sean MacAvaney, Paul McNamee, Douglas W Oard, Luca Soldaini, Ian Soboroff, Orion Weller, Efsun Kayi, et al. 2024. On the evaluation of machine-generated reports. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1904–1915.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FActScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100. Association for Computational Linguistics.
- Ramakanth Pasunuru, Asli Celikyilmaz, Michel Galley, Chenyan Xiong, Yizhe Zhang, Mohit Bansal, and Jianfeng Gao. 2021. Data augmentation for abstractive query-focused multi-document summarization. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13666–13674.
- Nima Pourdamghani, Kevin Knight, and Ulf Hermjakob. 2016. [Generating English from Abstract Meaning Representations](#). In *Proceedings of the 9th International Natural Language Generation conference*, pages 21–25. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text

- transformer. *Journal of machine learning research*, 21(140):1–67.
- Hossein Rajaby Faghihi, Bashar Alhafni, Ke Zhang, Shihao Ran, Joel Tetreault, and Alejandro Jaimes. 2022. [CrisisLTLSum: A benchmark for local crisis event timeline extraction and summarization](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5455–5477. Association for Computational Linguistics.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Aafiya S Hussain, Talha Z Chafekar, Grishma Sharma, and Deepak H Sharma. 2022. [Event oriented abstractive summarization](#). In *Proceedings of the 19th International Conference on Natural Language Processing (ICON)*, pages 99–108. Association for Computational Linguistics.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. MpNet: Masked and permuted pre-training for language understanding. *Advances in neural information processing systems*, 33:16857–16867.
- Beth M. Sundheim. 1992. [Overview of the fourth Message Understanding Evaluation and Conference](#). In *Fourth Message Understanding Conference (MUC-4): Proceedings of a Conference Held in McLean, Virginia, June 16-18, 1992*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#).
- Sai Vallurupalli, Sayontan Ghosh, Katrin Erk, Niranjan Balasubramanian, and Francis Ferraro. 2022. [POQue: Asking participant-specific outcome questions for a deeper understanding of complex events](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8674–8697. Association for Computational Linguistics.
- Siddharth Vashishtha, Alexander Martin, William Gantt, Benjamin Van Durme, and Aaron White. 2024. [FAMuS: Frames across multiple sources](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8250–8273. Association for Computational Linguistics.
- Jiarui Wang, Richong Zhang, Junfan Chen, Jaein Kim, and Yongyi Mao. 2022. [Text style transferring via adversarial masking and styled filling](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7654–7663. Association for Computational Linguistics.
- Miriam Wanner, Seth Ebner, Zhengping Jiang, Mark Dredze, and Benjamin Van Durme. 2024. [A closer look at claim decomposition](#). In *Proceedings of the 13th Joint Conference on Lexical and Computational Semantics (*SEM 2024)*, pages 153–175. Association for Computational Linguistics.
- Patrick Xia, Guanghui Qin, Siddharth Vashishtha, Yunmo Chen, Tongfei Chen, Chandler May, Craig Harman, Kyle Rawlins, Aaron Steven White, and Benjamin Van Durme. 2021. [LOME: Large ontology multilingual extraction](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 149–159. Association for Computational Linguistics.
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. [AlignScore: Evaluating factual consistency with a unified alignment function](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11328–11348. Association for Computational Linguistics.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International conference on machine learning*, pages 11328–11339. PMLR.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2019. [Bertscore: Evaluating text generation with bert](#). *ArXiv*, abs/1904.09675.
- Markus Zopf. 2018. [Auto-hMDS: Automatic construction of a large heterogeneous multilingual multi-document summarization corpus](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. European Language Resources Association (ELRA).

A Additional Examples

Below, we show a few examples of the report summaries and their corresponding cross-document summaries to illustrate how the latter typically provide greater detail about an event of interest relative to the former. We note, however, that this is not always the case: sometimes the source document offers no additional information about the event beyond what is contained in the report.

Example 1

- **Frame:** CAUSE TO RESUME
- **Report:** *Areva **renewed** a uranium deal with Niger in January 2008.*
- **Cross-Doc:** *On January 13, 2008, French state-controlled nuclear reactor maker Areva CEFi said it had **renewed** a uranium mining deal with the state of Niger and would invest over 1 billion euros.*

Example 2

- **Frame:** SMUGGLING
- **Report:** *A woman pled guilty to possession and attempting to **smuggle** 89 grams of heroin out of Thailand.*
- **Cross-Doc:** *Scot Sandra Gregory pled guilty to possession and attempting to **smuggle** 89 grams of heroin out of Thailand in 1993 and did her time in Thai jails.*

Example 3

- **Frame:** HOSTILE ENCOUNTER
- **Report:** *The plot of *Reign of Shadows* involves players returning to the dark side of the moon of Luclin to **face** the snake-like Shissar race led by Emperor Ssraeshza.*
- **Cross-Doc:** *The plot of *Reign of Shadows* involves players returning to the heart of the dark side of the moon of Luclin to **face** the snake-like Emperor Ssraeshza and his unyielding throngs of insidious zealots and enslaved minions to take back the ancient citadel of Vex Thal and end their march.*

B Training and Evaluation

Models and Hardware The BART, T5, and PEGASUS models were all trained on a single NVIDIA Quadro RTX 6000 GPU using CUDA version 11.7. Results reported with these models are based on single runs with a fixed random seed. We fine-tune the following pretrained checkpoints available from HuggingFace:

- t5-large
- facebook/bart-large
- google/pegasus-large

Libraries Models were developed using Python 3.11.9. We used the following libraries for model training, inference, and evaluation:

- accelerate (0.34.2)
- bert-score (0.3.13)
- bm25s (0.2.1)
- datasets (3.0.1)
- deepspeed (0.15.1)
- editdistance (0.8.1)
- evaluate (0.4.3)
- metametric (0.1.2)
- numpy (1.26.4)
- rouge-score (0.1.2)
- sentence-transformers (3.1.1)
- spacy (3.7.5)
- torch (2.0.1+cu117)
- transformers (4.45.1)
- tokenizers (0.20.0)

Metrics We use the implementations of ROUGE ($\mathbf{R}_{1,2,L}$) and BERTScore (\mathbf{BS}) provided by the HuggingFace evaluate library. We implement CEAF-REE (\mathbf{CR}) and its soft-match variant (see Tables 8, 9) using the metametric package (Chen et al., 2023a). We use the implementation of AlignScore released by the metric’s authors (Zha et al., 2023).²⁰ Lastly, for FActScore, we use the few-shot examples from Wanner et al. (2024) for decomposition and use Llama3.1-8B Instruct (Touvron et al., 2023; Dubey et al., 2024) for both atomic fact decomposition and verification.

Hyperparameters BART, T5, and PEGASUS were all trained for a maximum of 30 epochs with a patience of 5 epochs, using ROUGE-1 (\mathbf{R}_1) F_1 score on the dev set as the evaluation criterion. We use the Adam optimizer (Kingma and Ba, 2014) with default hyperparameters ($\beta_1 = 0.9$, $\beta_2 = 0.990$, $\epsilon = 1e^{-8}$, $\eta = 0.001$) for all models. For inference, we use beam search decoding with a beam size of 5 and set the maximum tokens to 256.

²⁰<https://github.com/yuh-zha/AlignScore>

Input Formats Below, we show in greater detail the input format for BART, PEGASUS, and T5 for the report and cross-document results reported in Table 2 and Table 3. (Note: these input formats were also used to obtain the BART, PEGASUS, and T5 results in the TEXT+EVENT rows in Table 8 and Table 9.) Here, $\langle B \rangle$ and $\langle E \rangle$ denote the model’s start-of-sequence and end-of-sequence tokens, respectively (if applicable), and $\langle S \rangle$ denotes a special token used to delineate information pertaining to a particular event role. Other text set between angle brackets ($\langle \dots \rangle$) denotes a variable placeholder. We add spaces between separators and adjacent text to improve readability below; they are not present in the actual input.

The input format for the **report** task is:

```

 $\langle B \rangle$  Report:  $\langle$ Report Text $\rangle$   $\langle E \rangle$   $\langle B \rangle$ 
Frame  $\langle S \rangle$   $\langle$ Frame Name $\rangle$   $\langle S \rangle$  Trigger  $\langle S \rangle$ 
 $\langle$ Trigger $\rangle$   $\langle S \rangle$   $\langle$ Role 1 Name $\rangle$   $\langle S \rangle$   $\langle$ Arg
1 $\rangle$ ;  $\langle$ Arg 2 $\rangle$ ; ...  $\langle S \rangle$   $\langle$ Role N Name $\rangle$   $\langle S \rangle$ 
 $\langle$ Arg 1 $\rangle$ ;  $\langle$ Arg 2 $\rangle$ ; ...  $\langle S \rangle$   $\langle E \rangle$ 

```

The input format for the **cross-document** task is:

```

 $\langle B \rangle$  Report:  $\langle$ Report Text $\rangle$   $\langle S \rangle$  Source:
 $\langle$ Source Text $\rangle$   $\langle E \rangle$   $\langle B \rangle$  Report Event:
Frame  $\langle S \rangle$   $\langle$ Frame Name $\rangle$   $\langle S \rangle$  Trigger  $\langle S \rangle$ 
 $\langle$ Trigger $\rangle$   $\langle S \rangle$   $\langle$ Role 1 Name $\rangle$   $\langle S \rangle$   $\langle$ Arg
1 $\rangle$ ;  $\langle$ Arg 2 $\rangle$ ; ...  $\langle S \rangle$   $\langle$ Role N Name $\rangle$   $\langle S \rangle$ 
 $\langle$ Arg 1 $\rangle$ ;  $\langle$ Arg 2 $\rangle$ ; ...  $\langle S \rangle$  Source Event:
Frame  $\langle S \rangle$   $\langle$ Frame Name $\rangle$   $\langle S \rangle$   $\langle$ Role 1
Name $\rangle$   $\langle S \rangle$   $\langle$ Arg 1 $\rangle$ ;  $\langle$ Arg 2 $\rangle$ ; ...  $\langle S \rangle$ 
 $\langle$ Role N Name $\rangle$   $\langle S \rangle$   $\langle$ Arg 1 $\rangle$ ;  $\langle$ Arg 2 $\rangle$ ;
...  $\langle S \rangle$   $\langle E \rangle$ 

```

The ablation settings presented in Table 8 and Table 9 (TEXT ONLY, EVENT ONLY, TEXT+SCHEMA) do not fundamentally change this overall structure, but merely omit parts of it (e.g. TEXT+SCHEMA omits all \langle Arg N \rangle).

C LLMs

GPT All GPT models were accessed through the OpenAI Chat API²¹, via the OpenAI Python SDK (openai 1.50.2). As noted in §4, we set temperature to 0.7 and set the maximum output tokens to 256 (consistent with the fine-tuned models) for all experiments reported in this paper and leave the

²¹<https://platform.openai.com/docs/api-reference/chat>

other API defaults unchanged ($n = 1$, top_p is not set, and we use no frequency penalty, presence penalty, or logit bias). For GPT-4o, we used model version gpt-4o-2024-08-06. For GPT-4o Mini, we used model version gpt-4o-mini-2024-07-18. Results reported throughout the paper are based on a single generation per prompt.

Claude All Claude models were accessed through the Anthropic Messages API²² via the Anthropic Python SDK (anthropic 0.34.2). As with the GPT models, we set temperature to 0.7 for all experiments in this paper and leave the other defaults unchanged (we do not set top_p or top_k, as recommended, and we do not set any stop sequences). For Claude 3.5 Sonnet, we used model version claude-3-5-sonnet-20240620. For Claude 3 Haiku, we used model version claude-3-haiku-20240307. Results reported throughout the paper are based on a single generation per prompt.

Prompts We use the same prompts for all LLMs. Complete prompts will be available in the public GitHub repository for this work. Here, we provide prompt templates used to obtain the results in Table 2 and Table 3, for both tasks (report or cross-document) and for both the zero- (ZS) and few-shot (FS) settings. Text set between angle brackets ($\langle \dots \rangle$) denote placeholders.

We use the same system prompt for both tasks:

```

You are an expert intelligence briefer.
Your task is to analyze a specific, important event based ONLY on certain information, and to compile a concise summary of that event to be presented to a high-ranked decision maker.

```

For the **report** task in the **zero-shot (ZS)** setting, the user prompt has the following structure:

```

The Report text below describes a situation. The Report Template provides specific details about the same situation. Focus ONLY on information relevant to the Situation Type.

```

```

Please write a short, accurate summary that is one sentence long and that is based ONLY on the provided information. DO NOT include any extraneous details. DO NOT use more than one sentence.

```

²²<https://docs.anthropic.com/en/api/messages>

Situation Type: ⟨Frame Name⟩ (⟨Frame Def⟩)

Report: ⟨Report Text⟩

Report Template:

- ⟨Role 1⟩ (⟨Role 1 Def⟩): ⟨Arg 1⟩;
⟨Arg 2⟩;...

- ...

- ⟨Role N⟩ (⟨Role N Def⟩): ⟨Arg 1⟩;
⟨Arg 2⟩;...

Summary:

The **few-shot (FS)** user prompt for the **report** task had the following structure:

The Report text below describes a situation. The Report Template provides specific details about the same situation. Focus ONLY on information relevant to the Situation Type.

Please write a short, accurate summary that is one sentence long and that is based ONLY on the provided information. DO NOT include any extraneous details. DO NOT use more than one sentence.

Here are a few examples to show you how to complete the task:

Example 1

Situation Type: ⟨Frame Name⟩ (⟨Frame Def⟩)

Report: ⟨Report Text⟩

Report Template:

- ⟨Role 1⟩ (⟨Role 1 Def⟩): ⟨Arg 1⟩;
⟨Arg 2⟩;...

- ...

- ⟨Role N⟩ (⟨Role N Def⟩): ⟨Arg 1⟩;
⟨Arg 2⟩;...

Summary: ⟨summary text⟩

Example 2

⟨ same format as above ⟩

Example 3

⟨ same format as above ⟩

Now here is the target example for you to complete:

Target

⟨same format, but with summary text omitted⟩

The **zero-shot** user prompt for the **cross-document** task had the following structure:

The Report text below describes a situation, and the Report Template provides specific details about the same situation. The Source text provides additional context about this situation, and the Source Template provides additional details. Focus ONLY on information relevant to the Situation Type.

Please write a short, accurate summary that is preferably one sentence long (and no more than two sentences long) based ONLY on the provided information. DO NOT include any extraneous details. TRY to use one sentence and DO NOT use more than two.

Situation Type: ⟨Frame Name⟩ (⟨Frame Def⟩)

Report: ⟨Report Text⟩

Report Template:

- ⟨Role 1⟩ (⟨Role 1 Def⟩): ⟨Arg 1⟩;
⟨Arg 2⟩;...

- ...

- ⟨Role N⟩ (⟨Role N Def⟩): ⟨Arg 1⟩;
⟨Arg 2⟩;...

Situation Type: ⟨Frame Name⟩ (⟨Frame Def⟩)

Source: ⟨Source Text⟩

Source Template:

- ⟨Role 1⟩ (⟨Role 1 Def⟩): ⟨Arg 1⟩;
⟨Arg 2⟩;...

- ...

- ⟨Role N⟩ (⟨Role N Def⟩): ⟨Arg 1⟩;
⟨Arg 2⟩;...

Summary:

The **few-shot** user prompt for the **cross-document** task (not explicitly shown) follows exactly the same structure as the few-shot prompt for the report task, but naturally uses the cross-document example format in lieu of the report format.

D Human Evaluation

Full instructions for the human evaluation, along with a JSON file containing the items that were rated, are provided in our GitHub repo (<https://github.com/wganttt/SEAMuS>).

E Data & Annotation

E.1 License

We release SEAMUS and our code under a CC-BY-SA-4.0 license. As noted in the **Ethics** section, we intend SEAMUS for research use only, not for commercial purposes.

E.2 Additional Summary Statistics

Additional summary statistics—about the report and source texts are shown in [Table 7](#)

E.3 Inter-Annotator Agreement

As we note in §3, there was no redundancy in the SEAMUS annotation process: corrections to the FAMuS arguments, writing of summaries, and annotation of summary arguments were performed by a single annotator for each example. However, annotators did conduct a 10-example practice annotation for both the report and cross-document tasks. Thus, to give some (limited) sense of the inter-annotator agreement, [Table 5](#) and [Table 6](#) present pairwise comparisons of annotators’ annotations on these 10 items for the report and cross-document tasks (respectively) using the *reference-based* metrics from [Table 2](#) (plus the edit distance version of **CR**, **CR_{soft}**; see [Appendix F](#)). We treat annotations produced by annotators in the *P* column as “predictions” to be evaluated against the “reference” annotations produced by annotators in the *R* column. Two important notes:

1. Because all of these metrics are F_1 scores, the distinction between *P* and *R* is moot and reversing *P* and *R* for any given pair would yield the same results. In both tables, we report results for all unordered annotator pairs, as well as the average across all pairs.
2. Because these were practice annotations, none of them were included in the final SEAMUS dataset. We would thus expect the numbers reported here to be an *underestimate* of the level of agreement on the main task, had we had redundancy.

<i>P</i>	<i>R</i>	R₁	R₂	R_L	BS	CR	CR_{soft}
A1	A2	67.1	44.5	52.6	93.4	72.7	86.9
A1	A3	72.4	53.9	63.1	94.4	76.4	86.5
A2	A3	77.2	60.5	62.5	94.7	75.9	89.6
Avg.		72.2	53.0	59.2	94.2	75.0	87.7

Table 5: Inter-annotator agreement on the 10 practice examples from the SEAMUS **report summary** annotation, as given by the reference-based metrics we report in §4, treating annotator *P*’s responses as predictions and *R*’s responses as references (the reverse is equivalent, since these metrics are symmetric).

<i>P</i>	<i>R</i>	R₁	R₂	R_L	BS	CR	CR_{soft}
A1	A2	64.8	42.4	53.3	94.0	40.5	56.6
A1	A3	64.7	43.8	51.6	93.0	50.6	65.0
A1	A4	45.8	22.7	33.0	90.7	49.8	65.2
A1	A5	69.0	48.1	58.8	94.6	46.8	62.8
A2	A3	77.7	66.8	72.4	94.9	50.6	65.6
A2	A4	55.4	33.9	40.8	91.2	50.7	66.8
A2	A5	72.0	56.1	64.9	94.1	50.6	66.9
A3	A4	55.4	33.2	41.9	90.7	51.8	69.0
A3	A5	71.0	57.2	61.8	93.6	52.6	69.8
A4	A5	48.2	27.2	37.3	90.7	52.8	69.6
Avg.		62.4	43.1	51.6	92.8	49.7	65.7

Table 6: Inter-annotator agreement on the 10 practice examples from the SEAMUS **cross-document summary** annotation, as given by the reference-based metrics we report in §4, treating annotator *P*’s responses as predictions and *R*’s responses as references (the reverse is equivalent, since these metrics are symmetric).

	Report		Source	
	Train	Dev	Train	Dev
Examples	759	253	759	253
Avg. Words	59	60	1,084	1,511
Avg. Sentences	2.0	2.0	44.7	61.5
Avg. Arguments	3.1	3.5	3.8	4.2

Table 7: Summary statistics for the SEAMUS report (left) and source documents, which are the same as those in the FAMUS dataset, albeit with slightly different arguments due to our corrections of the original FAMUS argument annotations.

E.4 Annotation Interface

Here, we include screenshots of the annotation interface used to complete the Phase 2 annotation.²³ As noted in §3, the interface was adapted from Vashishtha et al.’s (2024) annotation interface for the FAMUS cross-document argument extraction task (cf. Figures 5 and 6 in Appendix A of their paper). Tasks were run via Turkle, an open-source tool with similar functionality to Amazon Mechanical Turk.²⁴

In the first part of the Phase 2 annotation, the existing (crowdsourced) FAMUS argument annotations for the source text were reviewed and corrected, and the cross-document summaries were written jointly on the basis of these corrected annotations and the corrected report text argument annotations from Phase 1 (see Figure 4). The interface was pre-populated with (a) the corrected report text arguments from Phase 1 (in the “Report Text” tab, highlighted); the report summary from Phase 1 (in the “Report Summary” field); and (c) the uncorrected source text arguments (in the “Source Text” tab). The source text arguments were reviewed and corrected by toggling to the “Source Text” tab and making any necessary edits to the existing selections. The cross-document summaries were then written in the “Combined Summary” field. The UI for selecting, adding, and removing arguments was unchanged relative to Vashishtha et al.’s implementation. The major differences here are the addition of the “Report Summary” and “Combined Summary” fields, and the inability to alter the selected FrameNet frame for annotation.

In the second part, arguments were annotated on the summaries written in the first part (Figure 5). The interface is similar to the interface for the first part of the Phase 2 annotation, except that the “Report Summary” and “Combined Summary” fields have been removed, and a new tab (“Summary Text”) containing the cross-document summary to be annotated was added. Summary arguments were annotated by toggling to this tab and making argument selections in the same way as before. Here, the corrected argument annotations for *both* the report text *and* for the source text were pre-populated for each task under their respective tabs, allowing annotators to toggle between these for reference in

The screenshot shows a web-based annotation interface. At the top, there is an "Instructions" button. Below it, the "Select Event Type" is set to "Emotions_success_or_failure" and the "Active Event Type" is also "Emotions_success_or_failure". The "Event Definition" states: "A Topic or Stimulus falls or succeeds at causing an emotion in an Experiencer, which is conveyed by an Expressor or manifested in a State or description of an Event." An "Event Example" shows "They were dissatisfied with the arrangements." A link provides more details about the event type. The "Report Summary" field contains the text: "On October 26, an immigration tribunal stated that it was 'satisfied that the appellant has engaged in the unacceptable behaviour of fostering hatred, ' which it said could lead to intercommunity violence in the UK'". Below this is the "Combined Summary" field with a placeholder: "Enter your combined source-report summary here!". A "Unique ID" is displayed as "EN-4992-441-frame-Emotions_success_or_failure". There are two tabs: "Report Text" (active) and "Source Text". The "Report Text" tab shows a snippet of text: "On October 26, an immigration tribunal concluded that May had been justified in her position. The tribunal stated that it is 'satisfied that the appellant has engaged in the unacceptable behaviour of fostering hatred which might lead to intercommunity violence in the UK...'". Below the text are "Active Role" and "Answer" fields, both set to "Time" and "October 26" respectively, with an "Add Role" button. A "Roles" section contains buttons for "Experiencer", "Topic", "Stimulus", "Expressor", "State", "Event", and "Time", with "Time" selected. There are also "previous role" and "next role" buttons. A "Role Definition" explains: "The Time when the Experiencer, Expressor, Event, or State can be described as having said emotion." A "Role Example" is provided. At the bottom, there is a "Submit" button.

Figure 4: Interface for source text argument correction and cross-document summary writing (the first part of the Phase 2 annotation).

annotating the summary arguments.

As can be seen in both Figure 4 and Figure 5, details about the frame for the target event, including the frame name, its definition, as well as role names and their definitions, were provided as in the original FAMUS interface. Instructions were also accessible at any time via the dropdown shown at the top of the screen.

E.5 Annotation Instructions

Annotation instructions for both phases are available on our GitHub repo (<https://github.com/wgantt/SEAMuS>).

E.6 Annotator Demographics

The full set of annotators consists of six students (five graduate and one undergraduate) pursuing degrees in Computer Science (3), Linguistics (2), and Cognitive Science (1), all of whom are fluent English speakers. Only one was financially compensated for the annotations (at a rate of \$15 per hour), as this person initially became involved with the project through a university job board posting for the task, whereas the others were members of the lab from which the project originated. The project, and the intended use of their annotations, was clearly explained to all participants in meetings before they began any annotation.

²³Recall that the Phase 1 annotation, which involved correcting the FAMUS report text argument annotations and writing the report summaries, was done in JSON files.

²⁴<https://github.com/hltcoe/turkle-client>

Figure 5: Interface for annotation of arguments on the cross-document summaries (the second part of the Phase 2 annotation).

F Additional Results

F.1 Main Results

Table 10 and Table 11 contain 95% confidence intervals of the results in Table 2 based on non-parametric bootstraps ($n = 1,000$).

F.2 Input Ablations

Here, we include the full results of the ablations on the inputs introduced briefly in §4.4, which were inspired by similar ones conducted by Gantt et al. (2024). In the TEXT ONLY setting, we omit information about the target event entirely and include only the text in the input—either the report for the report task, or both the report and source for the cross-document task—effectively reducing the problem to standard summarization. In the EVENT ONLY setting, we omit the text(s) and include only information about the target event—either the report event annotations for the report task, or both the report and source event annotations for the cross-document task—making this ablation similar to structure-to-text tasks, such as AMR-to-text (Pourdamghani et al., 2016)). In the TEXT+SCHEMA setting, we omit the argument annotations, but leave in information about the frame and its roles. For the fine-tuned models, we include just the names of the frame and its roles. For the LLMs, we additionally include the definitions of the frame and roles as given in FrameNet. Finally, TEXT+EVENT is the name we assign to the *unablated* setting, used to obtain the results in Table 2 and Table 3, where both the text(s) and the full event annotations are present in the input. For all

ablation settings, BART, PEGASUS, and T5 are fine-tuned on the ablated inputs using the same settings for training and inference as are described in §4. For the GPT and Claude models, the examples provided in the few-shot setting are also ablated in the way called for by each ablation.

Report Results for the report task are in Table 8. Here and in the cross-document results to follow (Table 9), we include a variant of CEAF-REE (CR) that we dub CR_{soft} , which aligns and scores predicted arguments against reference arguments using normalized levenshtein distance rather than exact match—enabling a more nuanced comparison of different models’ ability to recover event arguments in the summaries they produce.

Across all models and most metrics, we see significant drops in performance when ablating any component of the input. Notably, a number of models, especially the LLMs, fall to numbers near or below those of the report baseline (RB) on a variety of metrics.

There are, however, some unsurprising exceptions here. First, in many cases, results on CR and CR_{soft} in the EVENT ONLY ablation are markedly stronger than the report baseline, and are even competitive with the results in the unablated setting (TEXT+EVENT) for most of the zero-shot-evaluated LLMs. This echoes a similar finding by Gantt et al. (2024), who note that “the document [is not] needed to generate *some* string that contains all the [event] template’s arguments.” If this is correct, we would *expect* to see strong CR scores in the EVENT ONLY setting, even though the summaries may be poorer overall (as reflected in other metrics).

An intriguing, related observation is that whereas the fine-tuned models look dominant against the LLMs on CR in the unablated setting, this advantage sharply diminishes when we turn to CR_{soft} . This is likely explained by the fact that the fine-tuned models are able to learn the conventions adopted by annotators in selecting argument spans, whereas the (prompted) LLMs do not—even though they may still be generating outputs with approximately correct spans that are nonetheless harshly penalized by an exact match.

A second exception is the results on AlignScore (A) and FActScore (F) in the TEXT ONLY setting, which are competitive with—and in some cases superior to—the results in the unablated setting across models. Recall that both A and F here eval-

uate how well the report summary is supported by the report text. It is thus intuitively possible, and evidently quite feasible, to generate a summary that is adequately supported by the text without relying at all on the event annotations—which is exactly what is demanded by the TEXT ONLY setting. This is once again consistent with findings from Gantt et al. (2024) on the NLI-based family of metrics MENLI (Chen and Eger, 2023), which are broadly similar to AlignScore and FActScore: “[event] templates are not needed to generate *some* summary that is entailed by the document.”

We also note that, for the fine-tuned models, we obtain **A** scores in the TEXT+SCHEMA ablation that are comparable (T5) or higher than (BART, PEGASUS) those of the unablated setting. This makes sense, inasmuch as the TEXT+SCHEMA setting contains a superset of the inputs of the TEXT ONLY setting, though it is unclear why we do not find a similar pattern with the LLMs.

Finally, note that the report baseline, which treats the report text itself as the summary, should in theory achieve perfect **A** and **F** scores, and thus does not really represent a fair comparison with the other models (note: this is also true for the cross-document setting). That it does not is surely a reflection of the fact that both metrics rely on outputs from imperfect models. Such flaws of LM-based metrics must not be overlooked.

Cross-Document results on the cross-document task are shown in Table 9 and follow a pattern that is qualitatively very similar to that of the report results above. We consistently find that the best results are obtained in the unablated setting (TEXT+EVENT) for most metrics, with the same exception regarding **CR/CR_{soft}** in the EVENT ONLY setting as we found for the report task. Curiously, however, the findings on **A** are more complicated here: whereas we continue to see the strongest results on this metric in the TEXT ONLY and TEXT+SCHEMA ablations for the fine-tuned models, with the LLMs, we instead see our best results in the unablated setting—following the trend of other metrics.

F.3 Argument Recovery by Role

Table 12 and Table 13 show **CR** and **CR_{soft}** results (respectively) on the cross-document task broken down by role for the 20 roles with highest support (number of annotated arguments) in the SEAMUS training split.

Comparing the tables reveals an interesting dichotomy. For **CR**, no model is consistently dominant across all roles, with fine-tuned models collectively obtaining the best results on 12 of the 20 and few-shot prompted models obtaining the best results on the remaining 8. The **CR_{soft}** results, by contrast, heavily favor GPT-4O, which achieves the best scores on 13 roles. Here, the fine-tuned models are top-performing on only 4 roles.

We believe the same factor discussed in subsection F.2 explains this dichotomy: whereas **CR** requires exact span match—and thus will tend to favor models able to learn span boundary conventions through fine-tuning—**CR_{soft}** does not, and rewards spans proportional to their edit distance from the reference. Thus, **CR_{soft}** reveals the LLMs (and GPT-4O above all) to be effective in producing summaries that recover the correct arguments, albeit with more lexical modifications relative to the reference.

G Use of AI Assistants

GitHub Copilot was used as a coding assistant for parts of model development and data analysis, though its suggestions were carefully reviewed by the authors. AI assistants were **not** used for other parts of this work (writing, brainstorming, etc.).

Model	Ablation	Setting	R ₁	R ₂	R _L	BS	CR	CR _{soft}	A	F
Report Baseline	-	-	56.15	46.05	48.37	91.57	52.58	62.56	99.11	98.73
GPT-4o M	TEXT ONLY	ZS	49.96	28.18	39.23	91.31	34.59	53.13	95.74*	83.11
	EVENT ONLY	ZS	53.11	34.04	43.67	91.51	52.13	77.37	60.98	53.42
	TEXT+SCHEMA	ZS	53.29	31.60	42.91	91.28	38.24	56.92	79.07	76.38
	TEXT+EVENT	ZS	62.18	42.32	51.26	93.17	58.48	78.71	86.04	75.80
	TEXT+EVENT	FS	71.98	55.35	61.03	94.34	66.80	83.66	94.06	83.32
GPT-4o	TEXT ONLY	ZS	51.52	29.90	40.90	91.50	33.75	52.06	94.49	84.00*
	EVENT ONLY	ZS	56.39	38.34	46.34	91.93	59.35	83.35	70.66	57.14
	TEXT+SCHEMA	ZS	56.57	37.19	47.08	92.00	42.37	61.50	81.66	73.05
	TEXT+EVENT	ZS	63.95	45.21	52.95	93.18	61.39*	82.60*	83.87	74.78
	TEXT+EVENT	FS	72.54 [†]	56.59 [†]	62.34 [†]	94.40	69.61 [†]	87.27[†]	94.72	81.58
CLAUDE H	TEXT ONLY	ZS	50.41	30.39	40.53	91.11	32.35	51.46	93.10	83.77
	EVENT ONLY	ZS	55.03	36.37	45.71	91.79	54.36	78.25	72.15	56.29
	TEXT+SCHEMA	ZS	57.67	38.51	47.68	92.08	41.36	59.10	83.24	77.05
	TEXT+EVENT	ZS	64.75	46.19	54.67	93.44	58.75	78.92	84.87	77.57
	TEXT+EVENT	FS	71.73	55.86	61.05	94.29	63.21	80.95	94.82	82.54
CLAUDE S	TEXT ONLY	ZS	46.98	22.83	36.24	90.78	25.68	45.88	91.31	82.41
	EVENT ONLY	ZS	55.66	36.89	46.21	92.13	56.38	78.54	72.15	60.37
	TEXT+SCHEMA	ZS	57.33	36.18	46.98	92.30	41.71	61.46	88.93	77.85
	TEXT+EVENT	ZS	67.38*	48.11*	56.52*	93.84*	61.07	81.35	92.96	80.59
	TEXT+EVENT	FS	72.16	54.64	61.29	94.54 [†]	65.66	83.68	95.89 [†]	83.86 [†]
BART	TEXT ONLY	FT	57.13	43.53	50.46	91.77	46.27	58.59	97.42	84.64
	EVENT ONLY	FT	58.34	40.96	48.51	91.83	59.82	75.34	51.17	52.41
	TEXT+SCHEMA	FT	62.23	49.43	55.55	92.59	52.92	65.83	95.01	83.34
	TEXT+EVENT	FT	74.46	61.68	66.42	94.57	69.88	82.72	91.59	79.25
PEGASUS	TEXT ONLY	FT	60.33	46.19	52.44	92.13	45.95	60.40	97.45	85.20
	EVENT ONLY	FT	59.69	41.97	49.46	91.90	57.14	74.34	53.93	53.43
	TEXT+SCHEMA	FT	63.28	49.79	55.91	92.71	53.69	66.28	96.94	84.33
	TEXT+EVENT	FT	75.18	62.53	66.96	94.70	70.00	82.68	96.08	82.23
T5	TEXT ONLY	FT	58.38	45.25	51.81	91.96	49.70	60.75	98.88	87.85
	EVENT ONLY	FT	63.14	45.62	52.47	92.67	64.00	80.08	68.42	62.63
	TEXT+SCHEMA	FT	65.82	51.90	58.46	93.11	56.18	68.42	97.92	82.93
	TEXT+EVENT	FT	76.64	64.44	68.90	95.02	74.20	85.22	98.15	85.02

Table 8: Input ablation results for the **report** summarization task. Best overall results are in **bolded**. * and [†] denote best zero- and few-shot results, respectively. See §4.1 for an explanation of metrics. See Appendix F for an explanation of the settings.

Model	Ablation	Setting	R ₁	R ₂	R _L	BS	CR	CR _{soft}	A	F
Report Baseline	-	-	48.52	33.28	39.31	89.58	31.00	42.04	99.29	93.12
GPT-4o M	TEXT ONLY	ZS	37.56	16.93	26.97	88.98	21.86	40.48	73.58	91.60
	EVENT ONLY	ZS	52.45	31.15	40.04	91.17	37.48	66.51	69.97	75.00
	TEXT+SCHEMA	ZS	41.88	20.40	30.32	89.72	24.04	44.76	76.64	89.12
	TEXT+EVENT	ZS	51.87	29.90	39.10	91.31	38.99	64.13	81.46	88.89
	TEXT+EVENT	FS	57.48	36.99	45.74	92.08	39.78	62.93	88.48	89.79
GPT-4o	TEXT ONLY	ZS	41.59	19.28	30.70	89.48	21.60	42.04	69.09	92.06
	EVENT ONLY	ZS	54.03	33.98	42.13	91.51	41.75*	69.63*	81.02	80.55
	TEXT+SCHEMA	ZS	49.87	27.04	37.76	90.86	25.80	48.53	85.44	89.75
	TEXT+EVENT	ZS	57.97	36.42	45.89	92.22*	41.34	68.04	86.61	88.41
	TEXT+EVENT	FS	61.17 [†]	40.62 [†]	49.38 [†]	92.67[†]	42.72 [†]	69.27 [†]	90.62	88.45
CLAUDE H	TEXT ONLY	ZS	47.27	25.48	36.49	90.23	22.64	43.20	84.29	92.59
	EVENT ONLY	ZS	53.35	33.01	42.94	91.39	38.64	66.08	77.70	76.83
	TEXT+SCHEMA	ZS	51.79	30.45	41.04	90.87	26.38	48.02	87.10	90.87
	TEXT+EVENT	ZS	57.72*	36.88*	46.35*	92.05	36.22	60.03	90.37	91.36
	TEXT+EVENT	FS	59.42	39.40	48.56	92.13	37.20	59.70	90.99	90.50 [†]
CLAUDE S	TEXT ONLY	ZS	44.13	20.08	32.73	89.88	19.93	40.24	87.26	92.30
	EVENT ONLY	ZS	53.51	33.51	42.73	91.53	39.78	66.17	84.12	81.91
	TEXT+SCHEMA	ZS	51.37	29.33	40.06	90.94	28.05	49.07	88.64	89.33
	TEXT+EVENT	ZS	56.77	34.75	45.27	91.91	35.24	59.47	93.41*	91.71*
	TEXT+EVENT	FS	57.95	38.05	47.53	92.09	37.32	59.31	95.09 [†]	90.39
BART	TEXT ONLY	FT	48.57	30.30	39.70	89.99	27.12	44.43	90.06	86.87
	EVENT ONLY	FT	56.37	37.04	45.14	91.21	39.12	62.90	56.01	68.10
	TEXT+SCHEMA	FT	51.67	35.12	44.15	90.42	32.31	49.47	94.45	90.52
	TEXT+EVENT	FT	63.77	45.50	52.98	92.59	44.97	66.36	85.55	85.27
PEGASUS	TEXT ONLY	FT	50.85	33.44	42.51	90.29	30.22	47.46	97.63	91.80
	EVENT ONLY	FT	58.52	38.41	46.46	91.42	39.98	64.06	67.05	75.80
	TEXT+SCHEMA	FT	51.21	34.18	43.11	90.28	30.15	47.04	97.99	92.72
	TEXT+EVENT	FT	63.66	46.24	53.18	92.51	43.73	64.51	93.85	90.48
T5	TEXT ONLY	FT	49.18	33.15	41.39	89.94	30.98	46.58	98.75	91.60
	EVENT ONLY	FT	59.96	40.55	47.51	91.84	45.30	68.85	73.73	78.98
	TEXT+SCHEMA	FT	53.06	35.64	44.93	90.64	31.87	50.14	94.11	91.30
	TEXT+EVENT	FT	64.14	46.36	52.79	92.56	44.67	65.66	92.48	90.19

Table 9: Input ablations on the **cross-document** summarization task. Best overall results are in **bolded**. * and [†] denote best zero- and few-shot results, respectively. See §4.1 for an explanation of metrics. See Appendix F for an explanation of the settings.

		Report							
Model	S	R ₁	R ₂	R _L	BS	CR	A	F	
GPT-4O M	ZS	[60.0, 64.5]	[39.3, 45.2]	[48.7, 54.0]	[92.8, 93.6]	[52.8, 60.5]	[80.5, 86.9]	[72.8, 78.6]	
	FS	[69.8, 74.0]	[52.5, 58.1]	[58.5, 63.5]	[94.0, 94.7]	[60.7, 68.8]	[93.4, 95.8]	[80.8, 85.8]	
GPT-4O	ZS	[61.6, 66.3]	[42.4, 48.2]	[50.5, 55.5]	[92.8, 93.6]	[54.2, 62.3]	[83.0, 88.8]	[71.2, 78.0]	
	FS	[70.3, 74.9]	[53.6, 59.8]	[59.6, 65.0]	[94.0, 94.8]	[62.7, 70.5]	[92.6, 95.4]	[78.4, 84.5]	
CLAUDE H	ZS	[62.7, 67.2]	[43.5, 49.1]	[52.4, 57.3]	[93.1, 93.9]	[52.5, 60.4]	[81.5, 87.7]	[74.1, 80.1]	
	FS	[69.4, 73.9]	[52.7, 58.7]	[58.5, 63.6]	[93.9, 94.7]	[58.1, 66.5]	[93.3, 96.1]	[79.6, 85.2]	
CLAUDE S	ZS	[65.1, 69.6]	[45.3, 50.8]	[54.1, 59.0]	[93.5, 94.2]	[55.0, 63.3]	[90.8, 94.8]	[77.6, 83.5]	
	FS	[69.8, 74.4]	[51.7, 57.5]	[58.8, 53.8]	[94.2, 94.9]	[60.7, 68.6]	[94.8, 96.7]	[80.8, 86.5]	
BART	FT	[71.9, 76.6]	[58.7, 64.6]	[63.7, 69.1]	[93.3, 94.1]	[64.3, 72.1]	[89.2, 93.9]	[76.1, 82.2]	
PEGASUS	FT	[72.9, 77.5]	[59.5, 65.4]	[64.2, 69.5]	[93.3, 94.1]	[65.4, 72.4]	[94.4, 97.5]	[79.4, 85.0]	
T5	FT	[74.3, 78.9]	[61.4, 67.3]	[66.1, 71.5]	[93.6, 94.4]	[69.7, 76.9]	[97.4, 98.8]	[82.4, 87.5]	

Table 10: 95% confidence intervals [low, high] from a non-parametric bootstrap ($n = 1000$) of the **report** results given in Table 2.

		Cross-Document							
Model	S	R ₁	R ₂	R _L	BS	CR	A	F	
GPT-4O M	ZS	[49.9, 53.7]	[27.8, 32.0]	[37.0, 41.0]	[91.0, 91.6]	[35.9, 42.2]	[78.5, 84.1]	[86.7, 90.7]	
	FS	[55.2, 59.7]	[34.3, 39.5]	[43.4, 47.8]	[91.7, 92.4]	[35.6, 42.9]	[86.0, 90.6]	[87.7, 91.6]	
GPT-4O	ZS	[55.7, 60.0]	[33.8, 38.9]	[43.5, 48.1]	[91.8, 92.6]	[36.4, 43.5]	[83.8, 89.0]	[86.0, 90.6]	
	FS	[59.0, 63.3]	[38.1, 43.1]	[47.1, 51.5]	[92.3, 93.0]	[38.0, 45.3]	[88.4, 92.8]	[86.3, 90.5]	
CLAUDE H	ZS	[55.6, 59.7]	[34.3, 39.2]	[44.0, 48.7]	[91.7, 92.4]	[32.8, 39.9]	[88.1, 92.3]	[89.7, 92.9]	
	FS	[57.0, 61.5]	[36.7, 42.1]	[46.0, 50.9]	[91.8, 92.5]	[33.4, 40.4]	[89.0, 92.9]	[88.8, 92.2]	
CLAUDE S	ZS	[54.7, 58.9]	[32.4, 37.2]	[43.1, 47.7]	[91.6, 92.3]	[31.1, 37.9]	[91.7, 95.0]	[89.6, 93.4]	
	FS	[55.6, 60.4]	[35.4, 40.8]	[45.2, 49.9]	[91.7, 92.5]	[33.9, 41.5]	[94.1, 96.0]	[88.5, 92.2]	
BART	FT	[61.5, 66.1]	[42.7, 48.4]	[50.5, 55.6]	[91.5, 92.2]	[41.3, 49.1]	[82.3, 88.6]	[82.6, 87.7]	
PEGASUS	FT	[61.2, 66.0]	[43.1, 49.0]	[50.4, 55.8]	[91.3, 92.1]	[40.9, 48.4]	[91.7, 95.6]	[88.7, 92.3]	
T5	FT	[61.5, 66.4]	[43.6, 49.2]	[50.2, 55.3]	[91.3, 92.2]	[40.3, 48.4]	[90.1, 94.4]	[88.2, 91.9]	

Table 11: 95% confidence intervals [low, high] from a non-parametric bootstrap ($n = 1000$) of the **cross-document** results given in Table 2.

Role	Support	GPT-4O M	GPT-4O	CLAUDE H	CLAUDE S	BART	PEGASUS	T5
TIME	523	39.13	40.69	34.17	37.25	42.07	44.32	47.09
PLACE	499	33.33	38.49	25.00	26.12	27.42	33.11	38.56
AGENT	240	34.67	32.89	27.40	23.13	42.38	32.43	39.74
THEME	94	49.12	44.07	43.33	35.09	40.00	39.44	39.34
ENTITY	65	29.27	35.90	35.00	35.90	30.00	35.90	41.03
PATIENT	53	41.18	30.30	50.00	34.29	43.75	51.61	48.48
GOAL	49	36.84	50.00	37.84	27.03	30.00	40.91	45.00
EVENT	43	14.81	20.69	20.69	7.14	18.75	25.00	6.45
CAUSE	42	6.45	24.24	11.43	12.12	31.25	10.81	26.67
EXPERIENCER	39	38.10	70.00	70.00	54.55	52.17	38.46	54.55
VICTIM	39	41.38	53.33	48.28	34.48	31.25	37.50	32.26
GOODS	38	26.67	75.00	0.00	28.57	50.00	50.00	14.29
PROTAGONIST	38	26.67	37.50	37.50	40.00	40.00	37.50	50.00
SOURCE	30	66.67	77.78	55.56	63.16	63.16	50.00	66.67
TOPIC	26	13.33	13.33	0.00	14.29	15.38	13.33	0.00
SPEAKER	25	50.00	50.00	66.67	50.00	50.00	70.59	58.82
ADDRESSEE	22	33.33	60.00	16.67	40.00	60.00	40.00	33.33
STIMULUS	21	15.38	30.77	33.33	33.33	46.15	33.33	50.00

Table 12: **CR** F₁ results on test set **cross-document** summaries for the top 20 roles with highest support (# arguments) in the SEAMUS training split (which has 3,004 total arguments). Results with GPT and Claude models are from the few-shot (FS) setting. Best results for each role are **bolded**.

Role	Support	GPT-4o M	GPT-4o	CLAUDE H	CLAUDE S	BART	PEGASUS	T5
TIME	523	57.37	65.22	49.22	50.52	60.54	61.40	65.43
PLACE	499	44.84	53.39	37.46	37.67	46.93	47.67	52.15
AGENT	240	65.40	67.47	59.78	50.16	66.72	57.99	62.06
THEME	94	73.21	73.96	69.22	72.96	71.75	64.71	68.59
ENTITY	65	66.33	76.13	63.16	60.07	63.34	63.12	69.44
PATIENT	53	76.28	73.52	74.41	68.62	70.88	73.28	73.36
GOAL	49	50.53	66.30	50.19	45.42	47.30	53.63	60.30
EVENT	43	52.75	63.52	52.76	42.95	45.97	52.09	35.38
CAUSE	42	47.66	52.99	39.41	42.03	43.23	43.87	49.82
EXPERIENCER	39	68.83	91.48	82.86	69.44	56.03	60.99	61.06
VICTIM	39	62.13	74.79	71.75	60.13	65.62	60.53	55.89
GOODS	38	42.77	79.33	33.35	50.68	61.94	57.46	24.39
PROTAGONIST	38	60.13	72.52	54.02	66.03	66.57	67.63	61.24
SOURCE	30	78.00	84.13	60.54	65.55	66.80	55.03	69.62
TOPIC	26	19.06	19.06	17.75	20.42	21.70	35.80	12.96
SPEAKER	25	58.21	61.41	71.85	64.91	64.92	73.03	66.28
ADDRESSEE	22	44.60	89.49	40.21	52.63	63.63	73.95	41.94
STIMULUS	21	38.61	70.46	54.55	66.70	53.39	74.75	57.65

Table 13: $\mathbf{CR}_{\text{soft}}$ (distinct from \mathbf{CR} ; see §F.2) F_1 results on test set **cross-document** summaries for the top 20 roles with highest support (# arguments) in the SEAMUS training split (which has 3,004 total arguments). Results with the GPT and Claude models are from the few-shot (FS) setting. Best results for each role are **bolded**.