

# A Unifying Scheme for Extractive Content Selection Tasks

Shmuel Amar<sup>◇</sup> Ori Shapira<sup>♣</sup> Aviv Slobodkin<sup>◇♣</sup> Ido Dagan<sup>◇♣</sup>

<sup>◇</sup>Bar-Ilan University, Israel <sup>♣</sup>Google Research, Israel <sup>♣</sup>OriginAI, Israel  
shmulikamar@gmail.com

## Abstract

A broad range of NLP tasks involve selecting relevant text spans from given source texts. Despite this shared objective, such *content selection* tasks have traditionally been studied in isolation, each with its own modeling approaches, datasets, and evaluation metrics. In this work, we propose *instruction-guided content selection (IGCS)* as a beneficial unified framework for such settings, where the task definition and any instance-specific request are encapsulated as instructions to a language model. To promote this framework, we introduce IGCS-BENCH, the first unified benchmark covering diverse content selection tasks. Further, we create a large generic synthetic dataset that can be leveraged for diverse content selection tasks, and show that transfer learning with these datasets often boosts performance, whether dedicated training for the targeted task is available or not. Finally, we address generic inference time issues that arise in LLM-based modeling of content selection, assess a generic evaluation metric, and overall propose the utility of our resources and methods for future content selection models.<sup>1</sup>

## 1 Introduction

Various NLP tasks essentially perform extractive *content selection*, where, given a single or multiple source texts as input, the system has to select targeted spans within them as the output. In some tasks, particularly extractive text summarization (Barzilay and Elhadad, 1997; Carbonell and Goldstein, 1998), the selection criterion is generic, where only the source texts are given as input while the output selection criteria are determined by the task itself. In other tasks, such as query-focused (Xu and Lapata, 2020) or aspect-based (Ahuja et al., 2022) summarization, and evidence detection (Wadden et al., 2020; Ernst et al., 2024), a specific instance-level input is provided by the user, which specifies the

requested information for that instance (e.g., a query, an aspect label, or a given claim for which evidence is sought, respectively). Additionally, content selection often becomes a natural sub-task in broader applications. For example, in attributable text generation, identifying attributions for a generated sentence in source (Saha et al., 2023; Slobodkin et al., 2024) or reference (Gao et al., 2023) texts is in essence a content selection task, where the instance-specific input is the model-generated sentence. Table 1 illustrates six existing content selection tasks (those included in our benchmark, described below). In sync with the tasks illustrated in the table, we focus in this paper on the setting where the expected selected source content conveys propositional information (complete facts), typically of a substantial length.

Traditionally, such content selection tasks were considered and studied in isolation, with tailored models, datasets, and evaluation methods for each. While earlier models relied on task-specific classifications over spans in the source texts (Devlin et al., 2019; Liu and Lapata, 2019), recent approaches are typically based on large language models (LLMs), where task-specific information is provided in the LLM prompt. In this paper, we suggest that the success of the latter approaches opens up the opportunity for a unified framework that would address effectively a broad range of content selection tasks.

Specifically, we propose such a unified framework, termed *instruction-guided content selection (IGCS)*; §3.1), where the task definition and instance-specific input (the “query”, when relevant) are given to the model as an instruction in the prompt. Following this scheme, we provide a first unified benchmark, IGCS-BENCH (§3.2), which we created by converting six existing datasets for different content selection tasks into a unified structure, while providing suitable instructions for each. This benchmark facilitates the development and evaluation of general-purpose content selection models that can address multiple tasks. In this

<sup>1</sup>Models and datasets available at <https://github.com/shmuelamar/igcs>.

Task	Instruction
EVIDSENT (Wadden et al., 2020)	Given the following <u>abstract document(s)</u> of medical papers, <u>select the sentences</u> that provide either <u>supporting or refuting evidence for the claim</u> : “ <u>① Dexamethasone decreases risk of postoperative bleeding</u> ”.
EVIDPROP (Ernst et al., 2024)	Given the following <u>documents</u> on the same topic, <u>extract short and concise text phrases</u> that provide <u>references to the following statement</u> : “ <u>the Cubs put him back in the lineup</u> ”.
SALIENCE (Ernst et al., 2024)	Given the following <u>documents</u> on the same topic, <u>extract short and concise salient text phrases</u> .
ASPSEL (Amar et al., 2023)	Given the following <u>news articles</u> on the topic “ <u>① Hurricane Andrew</u> ”, <u>extract all sentences related to</u> “ <u>② Government response</u> ”.
ASPSUM (Ahuja et al., 2022)	Given the following <u>news article</u> , <u>select at least 1 and at most 3 sentences that are the most relevant to</u> the <u>① Fraud’s nature of the fraud: the amount of money taken, benefits for the fraudster, and how the fraud worked</u> .
ARGMINE (Roush and Balaji, 2020)	Given the following <u>document</u> , <u>select short and concise text phrases that summarize all the evidence for the argument</u> : “ <u>① School choice is politically popular even among dems</u> ”.

Table 1: Illustration of the natural language content selection instructions for the six IGCS-BENCH tasks (§3.2). For each instruction, we highlight the type of single- or multi-document input, the requested information, the instance-specific query term(s) (when relevant), and the output requirements of the task.

context, we also propose to employ a particular existing evaluation metric as a generic metric for content selection settings (§3.3), while showing that this metric correlates well with task-specific metrics employed in prior works.

Notably, in the context of developing fine-tuned (small) language models over training data, our unified approach facilitates investigating *transfer learning* across datasets that were originally designed for different content selection settings. In particular, we show that the performance of a generic content selection model on a specific task often improves when fine-tuned on data created for other tasks. To further explore transfer learning across content selection settings, we leveraged top-performing LLMs to develop a larger *synthetic dataset*, which comprises a broad range of content selection scenarios (§4). Our results show that fine-tuning a generic model on this novel synthetic dataset improves performance across several tasks. These improvements are observed both when the synthetic data is used in a pure transfer setting, where targeted training data for the specific task at hand is not available, as well as when the synthetic data is used in combination with available task-specific training data (§6.1).

Finally, leveraging our unified framework, we investigate general inference-time issues that arise when utilizing LLMs for content selection, and propose strategies to reduce their impact—namely, document-level inference and alignment of the output with the source documents (§5.3). Overall, we suggest that future research on either existing or novel content selection tasks, with

or without targeted task-specific training data, would obtain significant gains by harnessing our provided datasets and methods.

In summary, our contributions include: (1) providing a unified scheme, benchmark suite, and evaluation measure for diverse content selection tasks (§3); (2) developing a synthetic training dataset that captures a diverse range of content selection scenarios (§4); (3) suggesting inference-time design choices when utilizing LLMs for content selection (§5); and (4) investigating and assessing transfer learning benefits across diverse content selection datasets, while showing the benefits of our novel synthetic dataset (§6).

## 2 Background

In this section we first provide an overview of the rich space of content selection settings (§2.1), which motivates our work, and a short review of content selection modeling approaches (§2.2).

### 2.1 Content Selection Tasks

Many end-tasks, as well as intermediate sub-tasks, can be considered as instances of a generic content selection setting, where spans within given source texts are extracted to satisfy an information need. A notable end-task example is extractive (generic) text summarization, where the output is typically a concatenation of selected salient source sentences (Barzilay and Elhadad, 1997; Carbonell and Goldstein, 1998; Wong et al., 2008; Nallapati et al., 2017; Zhang et al., 2023). Closely related, highlight summarization (Chen and Bansal, 2018;

Arumae et al., 2019; Cho et al., 2020; Ernst et al., 2024) selects salient spans of arbitrary length, which are then highlighted for the user within their original context. Similarly to these extractive end tasks, certain approaches for abstractive summarization employ salient-content selection as a first step, where the selected spans are then fed into an abstractive generation step (Chen and Bansal, 2018; Mao et al., 2020; Pilault et al., 2020; Li et al., 2021; Ernst et al., 2022; Adams et al., 2023; Wu et al., 2023; Saha et al., 2023; Slobodkin et al., 2024). Such approaches provide the advantage of easier traceability and attribution from the generated output texts back to the corresponding inputs from which they were generated.

In some extractive summarization variants, an instance-level input is provided to specify the requested output, such as an aspect label in aspect-based summarization (Ahuja et al., 2022; Gunel et al., 2023; Wang et al., 2024), a query in query-focused summarization (Xu and Lapata, 2020; Kulkarni et al., 2020; Hofmann-Coyle et al., 2022), or a question in long-form extractive question answering (Zhu et al., 2020; Potluri et al., 2023). Here again, explicit content selection from the sources either provides the eventual task output, in extractive settings, or supplies intermediate information that is further passed to an abstractive generation step.

Another content selection setting involves detecting supporting evidence for (pre-) given information. For example, in post-hoc attribution, evidence is sought for abstractive information that was (previously) generated by a model, such as in the RARR architecture (Gao et al., 2023) or LAQuer (Hirsch et al., 2025). In fact-verification and evidence extraction, supporting (or refuting) source spans are sought for externally given facts or claims (Thorne et al., 2018; Wadden et al., 2020; Krishna et al., 2023). Relatedly, argument mining extracts spans, from source documents, that function as claims or evidence pertaining to a particular stance (Stab and Gurevych, 2014; Roush and Balaji, 2020; Guo et al., 2023).

Finally, we note that the scope of our content selection setting should be distinguished from two related settings, which fall outside our scope. The first setting is passage retrieval (Callan, 1994; Karpukhin et al., 2020; Thakur et al., 2021), which is typically employed as an intermediate task that retrieves potentially relevant passage *candidates* from a large text corpus. These are

then passed to a more precise selection module (e.g., a ‘reader’) for making the correct selections, which are the target of our task setting (Chen et al., 2017; Karpukhin et al., 2020; Arivazhagan et al., 2023). The second setting involves the extraction of short phrasal spans, such as in extractive factoid question-answering (Rajpurkar et al., 2016; Kwiatkowski et al., 2019; Lewis et al., 2020) and information extraction (Cardie, 1997; Xu et al., 2024). In contrast, our setting focuses on tasks that require extracting text spans that jointly convey ‘open’ propositional content.

## 2.2 Modeling Content Selection

Several approaches were proposed in prior works for modeling content selection using instruction-tuned LLMs. In cases where the output selection consisted of complete sentences, a typical approach is to split the text into sentences and then ask an LLM to generate the indices of the selected sentences (Parmar et al., 2024; Sahu et al., 2025). This approach, however, is unsuitable when extracting spans of arbitrary length. Other methods either augment the input with special labels or prompt the model to repeat the entire input with such labels (Mallick et al., 2023; Sundar et al., 2024), but this becomes prohibitively expensive as the input size grows.

We focus on another modeling approach, where an LLM is instructed to select the requested input spans and copy them verbatim when generating the output. Yet, as observed by Ernst et al. (2024), models sometimes fail to verbatim copy the input as instructed, and instead generate outputs that deviate from the copied source. To address this issue, some works extended the decoder with constrained decoding techniques that ensure verbatim copying of input spans (Castel et al., 2022; Slobodkin et al., 2024). However, incorporating such extensions is not accessible when using API-based models. We take a simpler approach that recovers the selected source spans via a fuzzy match with the model’s output (§5.3).

## 3 Unified Content Selection Benchmark

In this section, we introduce our unified content selection benchmark, IGCS-BENCH, which was created by converting six existing datasets for different content selection tasks (§3.2) into our proposed general scheme (§3.1). Further, we propose adopting an existing and simple generic

evaluation metric for content selection (§3.3), which we later show to strongly correlate with four other metrics that were proposed for specific tasks (§6.3).

### 3.1 The IGCS Task Definition

A content selection task extracts a set of (typically disjoint) spans from input source texts, which jointly satisfy that specific task’s information need. As mentioned earlier, our scope focuses on settings where the sought information expresses propositional information (complete facts), rather than just phrases (e.g., in response to factoid questions). To generalize and unify the various conceivable settings of content selection, we suggest capturing the task-specific information need via a natural language *instruction*, given as an additional input to the model, as illustrated in Table 1.

For each task, the instruction structure is defined by a pre-specified template, which may include slots that are filled with instance-specific input, termed *query*. For example, in the aspect-based sentence selection task (ASPSEL) (Amar et al., 2023), the instruction template includes slots for the topic name associated with the input documents and the requested aspect label, pertaining to the given task instance. The selected text spans composing the output are expected to jointly convey the overall information requested by the instruction, and to follow its output format specification, as exemplified in Table 1.

### 3.2 IGCS-BENCH Tasks and Datasets

To compose IGCS-BENCH, we identified six existing content-selection tasks with human-annotated datasets from prior works, listed below. To ensure our benchmark data quality, we chose content selection tasks with high-quality datasets, which were created via reliable human annotation. Further, to make our benchmark useful for current research, we chose datasets on which current performance leaves room for improvement. Table 6 in Appendix A presents inter-annotator agreement figures and reported performances for the selected datasets.

In the original tasks, the input for each instance consists of a document set (possibly a single document), and in most cases also an instance-specific query that specifies the requested output for that instance (e.g., a specific aspect label or claim). To convert these tasks’ instances into IGCS instances,

we formulate an instruction template for each task (Table 1). Technical details for reproducing IGCS-BENCH are in Appendix A.

**Evidence Retrieval (EVIDSENT).** SCIFACT (Wadden et al., 2020) defines a task in which, given a set of medical abstracts and a scientific claim, the goal is to select sentences that either refute or support the claim.

**Proposition-level Evidence Detection (EVIDPROP).** Ernst et al. (2024) define a task in which, given multiple news documents and a proposition-level text span (representing a fact), the goal is to identify all proposition-level spans within the input documents that provide evidence for the given proposition. Compared to EVIDSENT, this task only detects supporting evidence and targets sub-sentence spans as output.

**Saliency Detection (SALIENCY).** Ernst et al. (2024) define a saliency detection task in which, given a set of input documents, the goal is to select the most salient proposition-level spans, capturing the information that could be incorporated in a generic summary. Notice that this task generically defines the requested output, without any instance-specific input.

**Aspect-based Sentence Selection (ASPSEL).** OPENASP (Amar et al., 2023) defines a task in which, given a set of documents on the same topic and an aspect label, the goal is to identify all sentences related to the specified aspect label.

**Extractive Aspect-based Summarization (ASPSUM).** ASPECTNEWS (Ahuja et al., 2022) is a dataset for extractive aspect-based summarization, where the task input is a single document and an aspect label. The task requires selecting between 1 and 3 sentences from the document that are most relevant to the aspect. In this dataset, multiple reference selections, produced by different annotators, are provided in each instance, capturing the higher level of subjectivity in this task.

**Argument Mining (ARGMINE).** DebateSum (Roush and Balaji, 2020) is a large argument mining dataset in which, given an article and an argument, the task is to extract all evidence spans supporting the argument. The dataset was originally annotated for evidence to be read aloud in debate competitions.

Task	# Instances	Query	Output Granularity	MD	∅	Source Token Len. (21–19389)	Selection Token Len. (0–5310)	Span Token Len. (0–718)
EVIDSENT	1109	Claim	Sentence	✓	✓	304	46.8	23.2
EVIDPROP	1332	Proposition	Proposition	✓	✗	2145	32.1	14.2
SALIENCE	98	–	Proposition	✓	✗	1909	436	12.8
ASPSSEL	51	Aspect	Sentence	✓	✗	8088	955	28.3
ASPSUM	400	Aspect	Sentence	✗	✗	265	83.1	30.4
ARGMINE	3000	Argument	Span	✗	✗	665	244	25.5
GENCS <sub>Union</sub>	12490	Instruction	Span	✓	✓	1181	86.0	30.1
GENCS <sub>Majority</sub>	12490	Instruction	Span	✓	✓	1181	75.7	28.1

Table 2: Content selection task properties, as detailed in §3.4. The six IGCS-BENCH tasks (§3.2) are at the top, and our two GENCS variants (§4) are at the bottom. The middle section describes task setup: **Query** — the instance-specific query type; **Output Granularity** — the type of selections for the output; **MD** — whether the source is multi-document or not; **∅** — the possibility for empty selections. The right section describes quantitative metrics, measured in token length. The histograms are shown on a  $\log_2$  scale on both axes, with the minimum and maximum value across all tasks displayed in the header, and the dark line marking the mean value, also written to the left of each histogram.

### 3.3 Evaluation Method

In gold instances, the reference (ground-truth) output specifies the set of source spans that should be selected for that instance. Prior works utilized various evaluation metrics to measure the degree of overlap between the gold and predicted source spans, often using variants of an  $F_1$  measure (Wadden et al., 2020; Ahuja et al., 2022; Amar et al., 2023; Ernst et al., 2024). As a generic evaluation metric, we suggest adopting (the existing) token-level source index comparison between the gold and predicted source spans, measuring token-level recall, precision, and  $F_1$  (Tjong Kim Sang and Buchholz, 2000; Ernst et al., 2024). As shown in §6.3, this metric highly correlates with the other four evaluation metrics that were originally used for the tasks in our benchmark.

Formally, given a reference selection  $S_r$  and a predicted selection  $S_p$ , we define  $T_r$  and  $T_p$  as the sets of token indices corresponding to the spans in  $S_r$  and  $S_p$ , respectively. We then compute precision, recall, and  $F_1$  scores between  $T_p$  and  $T_r$ . In cases where multiple alternative reference selections are provided in the dataset, we evaluate against each reference separately, and report the scores for the reference that yields the highest  $F_1$  score.<sup>2</sup> Finally, the system level scores are reported with average  $F_1$ , recall, and precision across all test instances.

<sup>2</sup>When there are multiple references, the predicted selection is expected to match just one reference, hence the maximum score among the references is used.

**Overall Scores.** As customary among other popular benchmark suites (e.g., Hendrycks et al., 2021; Suzgun et al., 2023), we compute a combined score for the IGCS-BENCH benchmark as the average of the individual task scores. A combined score enables convenient comparison between models tested on the benchmark. Specifically, each of the overall precision, recall, and  $F_1$  scores is a macro-average of the corresponding token-level metrics across the six IGCS-BENCH tasks. The confidence intervals for the overall  $F_1$  score is calculated with bootstrap resampling (Efron and Tibshirani, 1994; see Appendix B for details). Aside from the combined token-level scores, we report each individual task’s score using its original metric, as a reference point when examining a task independently. An overall “original” score is also computed, as the average original score across the six tasks (see Appendix A for details).<sup>3</sup>

### 3.4 Dataset and Task Properties

In the upper part of Table 2, we compare notable data properties that vary across tasks in IGCS-BENCH, collectively covering diverse content selection settings in our unified benchmark.

The leftmost section of the table specifies the number of instances included for each task. Next (middle section), we compare qualitative content selection properties across the tasks. Five of the

<sup>3</sup>As shown in §6.3, the two overall score variants correlate almost perfectly.

six tasks have an instance-specific *query* input, with the exception of SALIENCE. Three tasks define the *output granularity* at the sentence level, while the other three operate at the sub-sentence level; for example, in ARGMine, a single-word span may be part of the larger (non-consecutive) reference selection. Four tasks receive a multi-document set in the input (MD), which may be more challenging for a model to handle since overlapping and related content is scattered across documents. Finally, EVIDSENT is the only task where an empty output selection is possible ( $\emptyset$ ), i.e., an empty set of tokens; 37% of EVIDSENT instances fall into this category.

Next, we compute four quantitative measurements across the six tasks (rightmost section of Table 2). The first three columns display the averages and distributions of token-level lengths of the input source text, output selection, and individual spans, for each dataset.<sup>4</sup> Overall, ASPSEL has the largest input and output size on average, with the largest instance containing 19,389 input tokens and the largest output containing 5,310 tokens. In §6.2, we observe that a task’s average selection size may influence modeling performance.

## 4 Synthetic Dataset for IGCS

IGCS-BENCH (§3.2) represents a set of prior tasks that adhere to the IGCS scheme. Each dataset is structured differently in terms of the IGCS properties, and has a limited amount of training data. Our objective is to facilitate large-scale fine-tuning of IGCS models over diverse sets of properties and instructions, and not only over corpora derived from particular existing tasks. The typical approach to curate such datasets involves manual annotation which is labor-intensive. To address this, we build upon two methods from previous works that leverage existing large-scale corpora with human-written documents. Specifically, we employ targeted distillation, as proposed in Zhou et al. (2024), to transfer knowledge from general-purpose LLMs into smaller models tailored to the specific task of content selection, by generating synthetic instructions for these documents, as done by Köksal et al. (2024).

In this section we describe our three-step automated annotation pipeline which utilizes three top performing LLMs to generate two versions

<sup>4</sup>Throughout this paper, we use the spaCy (Honnibal et al., 2020) tokenizer, `en_core_web_sm`.

Corpus	# Inst.	# Docs	Source Len.	Agr.
PubMed	2252	1.0	424	81.8
Wikipedia	2025	1.0	1428	76.2
Email Threads	1991	1.0	735	69.8
Books	1704	1.0	1917	59.2
Multi-News	1654	2.6	1428	63.8
Hotel Reviews	1487	11.0	1567	54.4
GitHub	1377	1.9	1075	47.2
Overall	12490	2.5	1181	65.2

Table 3: Statistics of the source corpora forming GENCS. **# Inst.**—total instances; **# Docs**—average number of documents per document set; **Source Len.**—average tokens per document set; **Agr.**—inter-annotator agreement among the three models on selection annotation.

of a synthetic dataset for IGCS, called **Generic Content Selection (GENCS)**— $\text{GENCS}_{\text{Union}}$  and  $\text{GENCS}_{\text{Majority}}$ . In §6.2 we compare the effectiveness of several pipeline configurations for training a generic IGCS model.

### 4.1 Synthetic Dataset Generation

An IGCS dataset includes instructions for selecting content and the corresponding selected spans from the input text sources. The process for synthesizing the dataset is as follows.

**Source Corpora.** Inspired by the creation of The Pile corpus (Gao et al., 2021), we collected single- and multi-document sets from seven corpora spanning different domains, as detailed in Table 3. Specifically, we leveraged news article clusters from Multi-News (Fabbri et al., 2019), email threads (Klimt and Yang, 2004), English Wikipedia articles,<sup>5</sup> PubMed medical abstracts,<sup>6</sup> hotel reviews (Wang et al., 2010), books (Rae et al., 2020), and GitHub code<sup>7</sup> (technical details in Appendix C). For annotation, we sampled 500 document sets from each corpus, where each document set has an average of 1181 tokens (between 350 and 3500).<sup>8</sup>

**Step 1: Synthesizing Instructions.** In the first annotation phase we employed GPT-4<sup>9</sup> (OpenAI et al., 2023) to write five content selection instructions  $I_i^j$  for every sampled document set

<sup>5</sup><https://en.wikipedia.org>.

<sup>6</sup><https://huggingface.co/datasets/ncki/pubmed>.

<sup>7</sup><https://huggingface.co/datasets/codeparrot/github-code-clean>.

<sup>8</sup>Based on `nlk.tokenize.word.tokenize`.

<sup>9</sup>Snapshot `gpt-4-turbo-2024-04-09`.

$D_i$ , encouraging generation of diverse instructions (see details and prompts in Appendix C.2). In a real-world scenario, an instruction may yield an empty selection. We thus asked the LLM to generate challenging instructions with no relevant content in the source text, for 5% of the document sets in each corpus. Overall, we gathered 17,500 instructions, i.e., 5 instructions for 500 document sets in 7 corpora.

**Step 2: Synthesizing Candidate Content Selections.** In the second annotation phase, we prompted GPT-4, Claude3-Opus,<sup>10</sup> and Gemini-1.5-Pro,<sup>11</sup> to follow each of the instructions generated in the first phase and select content from the respective document set. Since the outputs from the LLMs may deviate from the exact wording in the source documents, we aligned the outputs with the source via a grounding method, described in §5.3, and rejected spans that could not be grounded to the source text. Additionally, we discarded any instruction instance where one of the models produced a response in an invalid format. Out of the 17,500 potential IGCS instructions, we gathered 12,490 that had three valid model selections (per corpus statistics are shown in Table 3).

**Step 3: Merging Possible Selections.** To produce the final reference selection for an instance, we explored two natural merging strategies for the three annotated selections: (1) the reference selection is set as the union of all selected tokens from the 3 selections, producing the recall-oriented  $\text{GENCS}_{\text{Union}}$  dataset; (2) the reference selection is the set of tokens selected by at least 2 models, producing the precision-oriented  $\text{GENCS}_{\text{Majority}}$  dataset. As shown in §6, different tasks might benefit more from either recall-oriented or precision-oriented data, hence we release both versions for future research. To conform to our definition of selection spans, a span is formed by concatenating consecutive selected tokens.

To conclude, this step results in the creation of the  $\text{GENCS}_{\text{Union}}$  and  $\text{GENCS}_{\text{Majority}}$  synthetic datasets. Each contains 12,490 instances of  $(D_i, I_i^j, S_i^j)$ , such that the selections differ in the two datasets according to the merging strategy

<sup>10</sup><https://www.anthropic.com/news/claude-3-family>, Snapshot claude-3-opus-20240229.

<sup>11</sup>gemini-1.5-pro-latest.

(in §6.2 we compare different dataset generation variants). The annotation cost is approximately \$550, with each dataset being twice the size of the entire IGCS-BENCH and can be further expanded by annotating additional samples from the source corpora.

## 4.2 GENCS Quality and Diversity

The GENCS dataset is extrinsically evaluated in §6 by demonstrating its utility for transfer learning in the content selection setting. In addition, we wish to directly assess its quality, and whether it meets our design goal of diversity in both the generated instructions and selections. To that end, we randomly sampled from the dataset three document sets, one of which has instructions for empty selections, from each of the seven source corpora, resulting in a total of 105 dataset instances. We instructed two annotators (NLP students) to rate each instruction and perform content selection as detailed below.

**Quality of Instructions.** A diverse dataset is expected to comprise instructions for various content selection use cases that require varying levels of informational *specificity*. The instructions should also be *natural* in the context of the given document set. Accordingly, the annotators rated each instruction on a Likert scale of 1 to 5 for *naturalness* and *specificity* (articulated in Figure 7 in Appendix E). The high average score of 4.0 ( $\sigma = 1.3$ ) for *naturalness* indicates that, overall, the instructions are plausible and relevant to their document sets. The average *specificity* score of 3.1 ( $\sigma = 0.8$ ) indicates that the values are scattered around 3, which reflects a broad range of scenarios, where the requested information varies from generic to anecdotal in relation to the topic of the document set.

**Quality of Selections.** A high-quality selection must accurately adhere to the given instruction by including only the relevant text spans from the input sources. To assess the selections in the GENCS dataset we measured their agreement with human-annotated selections.

We instructed human annotators to manually perform the content selection task for the sampled instances (see the annotation interface in Figure 6 in Appendix E), and computed the inter-rater agreement between annotations, following Hripacsak and Rothschild (2005) (see Appendix E for more details). Overall, we

measured a Cohen’s  $\kappa$  score of 0.7 among the three models, 0.59 among the two human annotators, and 0.61 human-LLM agreement, which indicate moderate to high agreement. In addition to preventing the significant effort from human annotators (as reported by our annotators), the LLMs were evidently capable of reliably producing selections.

Finally, selection diversity is measured through our content selection properties, as presented in the lower part of Table 2. The histograms in the table demonstrate diverse source, selection, and span sizes. This diversity emulates a wide range of content selection scenarios, including single large-span selections, empty selections, multiple short-span selections, and selections spanning multiple text sources.

## 5 Modeling

Following our proposed unified scheme for content selection (§3.1), our main objective is to assess whether, when modeling a particular content selection task, we can leverage generic training data such as GENCS. Accordingly, the focus of our modeling is to apply such transfer learning in different configurations, by fine-tuning feasibly-sized small language models, as described in §5.1.<sup>12</sup> In addition, we address, at inference time, two issues that arise when applying LLMs for content selection, by fragmenting the inference to apply over one document at a time, and by post-hoc matching between the generated output and their corresponding source spans (§5.3).

### 5.1 Transfer Learning Configurations

To model an Instruction-guided Content Selection (IGCS) task, we prompt an LLM, for each task instance, with the instance-specific instruction along with the source texts (or text) for that instance. In order to evaluate the effects of transfer learning between different content selection tasks, we fine-tune a popular LLM, **Llama-3-8B**, with various mixtures of training data, while utilizing the training datasets in IGCS-BENCH (§3) and our synthetic GENCS dataset (§4). In our fine-tuned models, we address two transfer learning scenarios: (1) a *transfer-only* setting, when no training data for the targeted task is available, thus

<sup>12</sup>We fine-tune small models of up to 8B parameters, which fits typical research computation budgets.

fine-tuning the model only over data for other tasks; (2) *supervision+transfer*, where training data for the targeted task is available, we use the same training data in the transfer-only setting but additionally include training data for the targeted task. To corroborate the robustness of the observed trends and behaviors in the Llama model, we also conduct analyses using fine-tuned models from other families, namely, **Qwen2.5** (Yang et al., 2024) and **SmolLM2** (Allal et al., 2025). These families offer multiple small-scale models that can be fine-tuned with modest computational resources.

We test our fine-tuned models over each of the six tasks in IGCS-BENCH. For transfer-only training, we fine-tune with two different compositions of data: (1) Leave-one-out (**LOO**)—mixing all available training sets in IGCS-BENCH (available for EVIDSENT, ASPSUM, and ARG MINE), except for the set of the targeted task being tested (simulating “out-of-domain” testing); (2) synthetic dataset (**GENCS**)—fine-tuning over one of the automatically generated dataset variants, **GENCS<sub>Union</sub>** or **GENCS<sub>Majority</sub>**. Analogously, in the supervision+transfer setting, we mix the training data of the targeted task with the same two compositions above of transfer data.<sup>13</sup> See Appendix F for technical details.

### 5.2 Prompt-based Models

As reference points for the results of our fine-tuned transfer models, we also report results for larger LLMs, obtained via zero- and few-shot prompting. Specifically, for zero-shot prompting, we use the proprietary **GPT-4**,<sup>14</sup> and **Claude3-Opus**<sup>15</sup> models, as well as the open-source **Llama-3** family of models (Dubey et al., 2024)—of 8B,<sup>16</sup> 70B, and 405B<sup>17</sup> parameters, respectively. For the few-shot in-context learning setting, we experimented with GPT-4 and Llama-3-8B (Dong et al., 2024), denoted **GPT-4<sub>ICL</sub>** and **Llama-3-8B<sub>ICL</sub>**, respectively. After experimenting with the number of in-context examples, we found 2-shot to perform best.

<sup>13</sup>We note that combining the two types of transfer data did not yield notable improvements in our experiments.

<sup>14</sup>[gpt-4-turbo-2024-04-09](https://openai.com/index/gpt-4-turbo-2024-04-09/).

<sup>15</sup>[claude-3-opus-20240229](https://claude.ai/claude-3-opus-20240229).

<sup>16</sup>[meta-llama/Meta-Llama-3-8B-Instruct](https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct).

<sup>17</sup>70B and 405B with <https://www.together.ai/blog/meta-llama-3-1>.

### 5.3 Inference-time Configurations

When employing LLMs for selecting text spans from their input, two issues arise. First, for multi-text inputs, the input context length, as well as the output length, may become challenging (Liu et al., 2024; Levy et al., 2024). Second, as mentioned in §2.2, while instructed to copy the selected source spans verbatim, LLMs sometimes deviate from the source, e.g., omitting words, generating paraphrases, or hallucinating. We next address these two issues.

**Document-level Inference.** Given multiple documents as input, the typical approach would be to concatenate all of them in a single prompt. However, we observed that when both the input and the expected output are relatively long, models tend to produce shorter outputs than required, decreasing performance (§6.2). Yet, in a content selection task, the output for a multi-document input could simplistically be viewed as the concatenation of selections from all documents. Hence, we experimented with prompting the model separately with each input document, then concatenating all output selections. While in this approach selection decisions for each document cannot consider information from other documents, we found it to perform better overall thanks to the shorter inputs and outputs, and hence adopted it in our modeling.

**Grounding the Output to the Source Text.** Addressing the second issue above, where the output selections generated by the model might deviate from the source text, we ground the model output back to the source. In case an exact match is not found, we exhaustively search for the closest source span in terms of token-level Levenshtein distance (technical details in Appendix G).

## 6 Results and Analysis

Our results and analysis assess the utility of transfer learning based on our unified content selection scheme (§6.1), the effectiveness of document level inference (§6.2), and the generality of our proposed generic evaluation metric (§6.3).

### 6.1 Transfer-learning of Fine-tuned Models

Table 4 presents the results of applying our various transfer learning configurations, including *transfer-only* and *supervised+transfer* (§5.1), as

well as results for larger prompt-based models (§5.2), tested over the six IGCS-BENCH datasets. The results of each task are measured with the corresponding evaluation metric of the original task dataset, while overall scores, averaged over all datasets (as explained in §3.3), are presented in the right-hand side of the table. Table 9 in Appendix I analogously presents the scores measured by token-level  $F_1$ , which we advocate as a generic content selection measure, yielding very similar trends.

#### 6.1.1 Transfer-only Configurations

The top section of Table 4 presents the results obtained in the absence of training data for the targeted tested task. The first two rows provide the baselines in this setting, namely, zero-shot and few-shot prompting, without any fine-tuning on transfer data. The 3rd row corresponds to leave-one-out fine-tuning over the IGCS-BENCH training data (i.e., testing on a task when the model is fine-tuned with the *other* tasks’ train sets). Rows 4-5 correspond to fine-tuning with GENCS.

As can be seen in the +LOO row, performance degrades substantially, relative to the prompt-based baselines, when transferring training data from a few specific content selection tasks to another task. In such a setting, the model seems to be steered toward distinct use-cases, which fails to generalize to different tasks. On the other hand, when fine-tuning with our *generic* GENCS dataset, the fine-tuned models outperform the prompt-based baselines with an overall  $F_1$  score of 45.7 compared to the highest scoring baseline with 41.8.

Additionally, the overall recall score of  $\text{GENCS}_{\text{Union}}$  and precision score of  $\text{GENCS}_{\text{Majority}}$  are at least 5.4 and 6.0 points higher than both baselines, respectively. This outcome is expected, as the union merging strategy encourages the fine-tuned model to select more tokens than the majority strategy, which is more conservative and therefore induces higher precision. Thus, the two GENCS fine-tuned variants offer complementary trade-offs, allowing one to be chosen when a given task prioritizes precision over recall or vice versa. For example, as shown in the task-specific results (left part of the table), ARG MINE, for which selection size is relatively long, benefits more from  $\text{GENCS}_{\text{Union}}$  than from  $\text{GENCS}_{\text{Majority}}$ . On the other hand,  $\text{GENCS}_{\text{Majority}}$  provides greater benefit for

		ASP	ARG	EVID	ASP	SALI-	EVID	Avg	Token-level		
		SUM	MINE	SENT	SEL	ENCE	PROP		P	R	$F_1 \pm CI$
transfer-only	Llama-3-8B <sub>ICL</sub>	34.6	<b>46.9</b>	44.8	28.6	<b>42.9</b>	13.5	35.2	42.3	51.0	41.2 $\pm$ 1.5
	Llama-3-8B	29.4	42.4	47.5	41.9	36.6	27.3	37.4	44.0	50.5	41.9 $\pm$ 1.6
	+ LOO	34.0	29.1	25.8	30.4	41.9	10.2	28.6	24.0	<b>61.9</b>	29.8 $\pm$ 1.5
	+ GENCS <sub>Union</sub>	<b>37.0</b>	36.7	42.6	<b>49.3</b>	37.5	<u>33.6</u>	<b>39.5</b>	45.7	56.4	<b>45.7</b> $\pm$ 1.7
	+ GENCS <sub>Majority</sub>	35.6	25.2	<b>48.1</b>	47.1	32.4	<b>35.6</b>	37.3	<b>50.0</b>	46.3	43.2 $\pm$ 1.7
supervision+transfer	Llama-3-8B <sub>Sup</sub>	40.6	63.5	66.0	–	–	–	–	–	–	–
	+ LOO	<u>42.3</u>	<b>64.1</b>	<u>70.0</u>	–	–	–	–	–	–	–
	+ GENCS <sub>Union</sub>	<u>42.7</u>	63.7	<b>72.1</b>	–	–	–	–	–	–	–
	+ GENCS <sub>Majority</sub>	<b>43.2</b>	63.6	68.8	–	–	–	–	–	–	–
	prompt-based models	Claude3-Opus	31.3	49.6	54.5	52.2	43.7	28.2	43.2	49.0	58.4
Llama-3-70B		29.3	40.7	58.5	56.8	33.2	44.9	43.9	55.7	49.4	47.8 $\pm$ 1.7
Llama-3-405B		30.0	45.4	56.2	59.8	35.1	42.1	44.7	51.6	57.1	49.3 $\pm$ 1.8
GPT-4		32.8	39.0	58.6	57.4	39.1	50.1	46.2	60.0	53.9	51.4 $\pm$ 1.9
GPT-4 <sub>ICL</sub>		33.9	45.5	57.3	55.0	39.9	47.5	46.5	59.7	55.5	52.2 $\pm$ 1.7

Table 4: Performance on IGCS-BENCH tasks (§6.1) using their original evaluation metrics (left) and overall **Token-level** metrics (right), comparing *transfer-only* (top), *supervision+transfer* (middle), to prompt-based methods (bottom). In the two topmost sections, the baseline(s) appear above the dashed line, and the transfer configurations of fine-tuned Llama-3-8B variants (§5.1) below it. In the *transfer-only* section, only GENCS is used for fine-tuning, whereas in the *supervision+transfer* section, the task-specific training set is also included in the fine-tuning mix, when the task has such a set. For each task in each section, **bold** indicates the highest score, and underline indicates statistical significance ( $p < 0.05$ ) with respect to the baselines in each section. The four rightmost columns report the overall average score across the six tasks (**Avg**) and token-level **Precision**, **Recall**, and  $F_1 \pm$  confidence interval ( $\alpha = 0.05$ ).

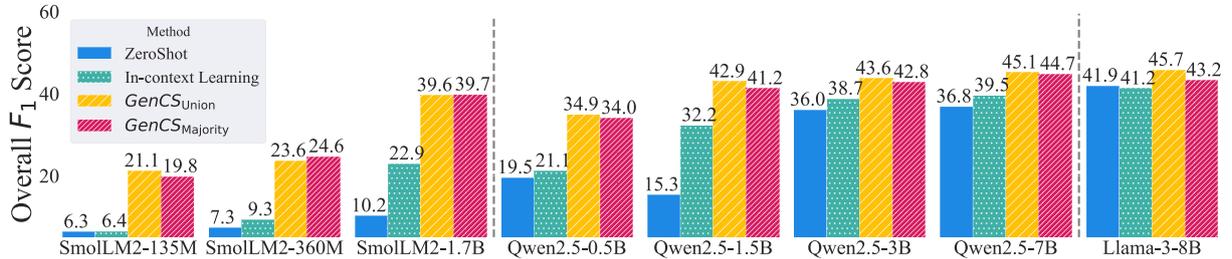


Figure 1: Overall  $F_1$  scores on IGCS-BENCH for four methods across eight small language models of up to 8B parameters from Qwen-2.5, SmolLM2 and Llama-3 families. All tested models benefit from fine-tuning with GENCS, especially smaller ones. The largest confidence interval across models is 2.0 ( $\alpha = 0.05$ ).

EVIDSENT, for which the expected selection size if relatively short (favoring precision). Overall across the benchmark, the union variant is more advantageous than the majority variant. Further, when examining the per-task results, we find that two out of the six tasks do not benefit from transfer-only fine-tuning. This suggests that while our results indicate that such fine-tuning is beneficial overall, when developing a generic content selection model, the utility of such fine-tuning should be verified when targeting a specific task. As a reference point, Table 6 in Appendix A

compares the results of transfer learning from our GENCS datasets to the analogous previously reported result for each dataset.

To more broadly investigate the advantages of fine-tuning with GENCS compared to prompt-based methods, we report overall  $F_1$  scores for seven additional models in Figure 1. Across all models evaluated, the fine-tuned variants consistently outperform their prompt-based counterparts. While, as may be expected, the smallest models exhibit the greatest gains from GENCS (transfer) fine-tuning, the larger models

we tested also exhibit significant gains of several  $F_1$  points.

To further assess the reliability of these results, we repeated each model run two more times, with two additional prompt variants that were derived automatically from the original human-written instruction, following a common-practice LLM-based prompt tuning, as elaborated in Appendix L. The results across these runs, presented in Figure 8 in the appendix, consistently show similar trends exhibited in Figure 1, while exhibiting some variations in absolute performances, demonstrating the typical sensitivity of LLMs to prompt variants (Sclar et al., 2024; Mizrahi et al., 2024).

To summarize, the results thus far suggest that fine-tuning with a generic dataset of diverse content selection scenarios allows the model to better generalize to different tasks, showing the potential value of our synthetic dataset when addressing tasks for which targeted training data is absent.

### 6.1.2 Supervision+transfer Configurations

The middle section of Table 4 presents results for the setting where training data *is* available for the tested task. The first row in this section provides the baseline for this setting, namely the zero-shot model (2nd row in the top section) fine-tuned with the training data of the tested task (available for three tasks). Naturally, such task-specific fine-tuning improves results for all three tasks (comparing the baseline rows of the two sections).

Notably, results improve by enriching the fine-tuning data with transfer training sets (second row of the section), and statistically significantly so (see Appendix J for details) for two of the three tasks in this setting. The magnitude of improvement is somewhat smaller than in the transfer-only setting, which is seemingly anticipated, since the transfer-training data supplements the existing task-specific training data. Interestingly, we observe that while transfer learning from different tasks was not helpful in the transfer-only setting (§6.1.1), it does improve performance for all three tasks when it is combined with the available task-specific training data. Finally, our GENCS datasets prove beneficial in this setting as well.

### 6.1.3 Larger Prompt-based Models

The bottom section of Table 4 provides the results for larger models using prompt-based approaches,

	SmolLM2-1.7B	Qwen2.5-7B	Llama-3-8B
Prompt-based	22.9 ± 1.4	39.5 ± 2.0	41.9 ± 1.6
GENCS <sub>1-step</sub>	32.9 ± 1.5	47.0 ± 1.6	44.7 ± 1.6
GENCS <sub>1-inst</sub>	34.2 ± 1.4	46.5 ± 1.7	44.1 ± 1.8
GENCS <sub>1-model</sub>	38.2 ± 1.5	44.7 ± 1.6	43.1 ± 1.8
GENCS <sub>Union</sub>	39.6 ± 1.6	45.1 ± 1.7	45.7 ± 1.7

Table 5: Overall  $F_1$  scores  $\pm$  confidence intervals ( $\alpha = 0.05$ ) on IGCS-BENCH of three fine-tuned models with training sets obtained using different synthetic pipeline configurations (§4). All fine-tuned models outperform the best prompt-based setting (first row), while different pipeline configurations are more effective for different models.

as reference points. These models mostly outperform the smaller 8B models in the transfer-only version, as can be expected when no task-specific training data is available for the targeted task. Still, transfer learning offers value when access to large models is limited or cost-prohibitive. When task-specific training data *is* available for fine-tuning, smaller models notably outperform larger models, and this performance gap widens further when transfer learning using GENCS is applied.

Taken together, our findings suggest that the proposed generic IGCS scheme, combined with transfer learning on our datasets, offers an effective and appealing approach for a variety of content selection tasks and use cases.

## 6.2 Ablation Analysis

### Synthetic Dataset Generation Configurations.

To test the impact of various design choices in our automatic dataset generation process (see §4), we generated several ablated variants of this process and tested the performance of 3 models when fine-tuned with the different dataset variants. Specifically, we created four variants of the GENCS dataset (details in Appendix D): (1) GENCS<sub>1-step</sub>—combining steps 1 and 2 such that an instruction and its selection are generated with a single prompt; (2) GENCS<sub>1-inst</sub>—generating only a single instruction in step 1; (3) GENCS<sub>1-model</sub>—using only a single model in step 2; and (4) GENCS<sub>Union</sub>—the full process used to generate the GENCS<sub>Union</sub> dataset, as described in §4. Then, we fine-tuned three models on each dataset variant. Table 5 presents the results of each fine-tuning variant, in comparison to the

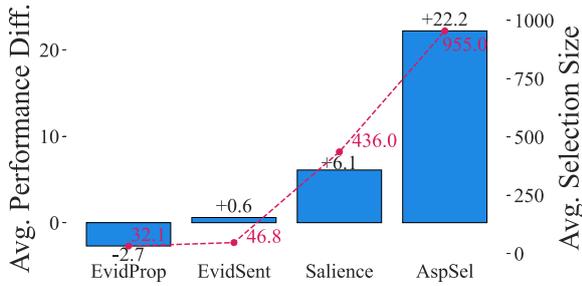


Figure 2: Difference in original metric scores between applying document-level inference (§5.3) and processing document-sets as a whole, averaged across 9 models. The red line shows the average *selection size* (§3.4) per task. It seems that document-level inference is more beneficial when a task has a higher selection size.

best prompt-based (zero-shot or in-context) baseline for the respective model (row 1). As shown, all fine-tuned models achieved higher overall performance on IGCS-BENCH than the baseline, where two models performing best with the full configuration.

**Document-level Inference.** We next analyze the impact of document-level inference (§5.3) over the four IGCS-BENCH tasks that have multi-document inputs, shown in Figure 2. To that end, we measure model performance on these tasks when feeding the model the full input versus feeding it document by document and then concatenating all document-level selections (§5.3). For each task, we measured the average performance of 9 models (details in Appendix K) on the task for each of the two settings (single vs. multi-document input), and plotted the difference between these averages in Figure 2. Notably, performance for ASPSEL and SALIENCE improves substantially, across all 9 models. Meanwhile, performance is negligibly affected for EVIDSENT, and slightly degrades for EVIDPROP. The differentiating factor seems to be the output *selection size*, where the model struggles to generate sufficient selections when processing all input documents at once. Processing one document at a time, on the other hand, triggers the model to generate more complete selections for each document, and hence for all documents as a whole. Further analysis appears in Appendix K.

### 6.3 Assessing a Generic Evaluation Metric

Next, we assess our proposal to use token-level  $F_1$  as a generic evaluation metric for content selection

tasks (§3.3). As mentioned earlier, each of the six IGCS-BENCH tasks has been originally evaluated using its own metric: While SALIENCE and EVIDPROP already employed token-level  $F_1$ , the other four tasks employed different metrics (see Appendix A for details). To assess the generality of token-level  $F_1$ , we measured its system level correlation with the other four metrics (see Table 7 in Appendix H). We find that the token-level  $F_1$  measure exhibits strong or very strong correlation with the all other metrics, suggesting its suitability as a generic evaluation metric for content selection tasks.

**Overall Score.** To evaluate the effectiveness of the overall score based on token-level  $F_1$ , we computed its correlation with the overall score derived from the original task-specific metrics (described in §3.3). The correlation considers the scores of all models and methods presented in Figure 1. We find that the two variants exhibit an almost perfect correlation, with Pearson’s  $r > 0.99$ . This suggests that the overall score based on the generic token-level metric is a reliable indicator of model performance on the IGCS-BENCH benchmark.

## 7 Conclusion

We introduced *instruction-guided content selection (IGCS)*, a unified scheme that generalizes a broad range of content selection tasks by encoding the task objective and input as a natural language instruction. To support this framework, we developed the first unified benchmark for this setting, IGCS-BENCH, and an extensive generic synthetic dataset for training, GENCS, which covers diverse content selection requests. Notably, we showed that leveraging these datasets for transfer learning is often effective, whether training data for the specific targeted task is available or not. Additionally, we propose document-level inference to circumvent the shortcomings of large language models when addressing content selection for long contexts. Finally, we proposed using token-level  $F_1$  evaluation as a standard generic metric for content selection, showing that it strongly correlates with prior task-specific metrics. Overall, we suggest the utility of our framework for future modeling of diverse content selection tasks, while paving the way for future research to model ad-hoc user-generated content selection instructions.

## 8 Limitations

IGCS-BENCH is built upon six particular content selection tasks. While these tasks are shown to be diverse, an alternative set of tasks may behave differently in terms of transfer learning, and lead to slightly different findings.

Throughout our study we experimented with several prompt variants, yet it is still possible that better prompts exist.

Finally, there is a lack of detailed documentation regarding the pre-training data used by the tested models. This makes it challenging to determine whether our test data is included in their training corpora, raising the possibility of data contamination.

## Acknowledgments

We thank the reviewers and the Action Editor for their constructive feedback, which substantially improved this manuscript. This work was supported by Israel Science Foundation grant 2827/21.

## References

- Griffin Adams, Alex Fabbri, Faisal Ladhak, Noémie Elhadad, and Kathleen McKeown. 2023. Generating EDU extracts for plan-guided summary re-ranking. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2680–2697, Toronto, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.acl-long.151>, PubMed: 38770277
- Ojas Ahuja, Jiacheng Xu, Akshay Gupta, Kevin Horecka, and Greg Durrett. 2022. ASPECTNEWS: Aspect-oriented summarization of news documents. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6494–6506, Dublin, Ireland. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.acl-long.449>
- Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Gabriel Martín Blázquez, Guilherme Penedo, Lewis Tunstall, Andrés Marafioti, Hynek Kydlíček, Agustín Piqueres Lajarín, Vaibhav Srivastav, Joshua Lochner, Caleb Fahlgrén, Xuan-Son Nguyen, Clémentine Fourier, Ben Burtenshaw, Hugo Larcher, Haojun Zhao, Cyril Zakka, Mathieu Morlon, Colin Raffel, Leandro von Werra, and Thomas Wolf. 2025. Smollm2: When smol goes big - data-centric training of a small language model. *CoRR*, abs/2502.02737v1. <https://doi.org/10.48550/ARXIV.2502.02737>
- Shmuel Amar, Liat Schiff, Ori Ernst, Asi Shefer, Ori Shapira, and Ido Dagan. 2023. OpenAsp: A benchmark for multi-document open aspect-based summarization. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1967–1991, Singapore. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.emnlp-main.121>
- Manoj Ghuhán Arivazhagan, Lan Liu, Peng Qi, Xinchí Chen, William Yang Wang, and Zhiheng Huang. 2023. Hybrid hierarchical retrieval for open-domain question answering. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10680–10689, Toronto, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-acl.679>
- Kristjan Arumae, Parminder Bhatia, and Fei Liu. 2019. Towards annotating and creating summary highlights at sub-sentence level. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 64–69, Hong Kong, China. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-5408>
- Regina Barzilay and Michael Elhadad. 1997. Using lexical chains for text summarization. In *Intelligent Scalable Text Summarization*.
- James P. Callan. 1994. Passage-level evidence in document retrieval. In *SIGIR '94*, pages 302–310, London. Springer London. [https://doi.org/10.1007/978-1-4471-2099-5\\_31](https://doi.org/10.1007/978-1-4471-2099-5_31)
- Jaime Carbonell and Jade Goldstein. 1998. The use of mmr, diversity-based reranking for re-ordering documents and producing summaries. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '98*,

- page 335–336, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/290941.291025>
- Claire Cardie. 1997. Empirical methods in information extraction. *AI Magazine*, 18(4):65. <https://doi.org/10.1609/aimag.v18i4.1322>
- Or Castel, Ori Ram, Avia Efrat, and Omer Levy. 2022. How optimal is greedy decoding for extractive question answering? In *4th Conference on Automated Knowledge Base Construction, AKBC 2022, London, UK, November 3–5, 2022*.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading Wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P17-1171>
- Yen-Chun Chen and Mohit Bansal. 2018. Fast abstractive summarization with reinforce-selected sentence rewriting. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–686, Melbourne, Australia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-1063>
- Sangwoo Cho, Kaiqiang Song, Chen Li, Dong Yu, Hassan Foroosh, and Fei Liu. 2020. Better highlighting: Creating sub-sentence summary highlights. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6282–6300, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.509>
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, Xu Sun, Lei Li, and Zhifang Sui. 2024. A survey on in-context learning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128, Miami, Florida, USA. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.emnlp-main.64>
- 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, Aurélien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Rozière, 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 Allonsius, 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, Grégoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel M. Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelder 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, et al. 2024. The llama 3 herd of models. *CoRR*, abs/2407.21783v1. <https://doi.org/10.48550/ARXIV.2407.21783>
- Bradley Efron and Robert J. Tibshirani. 1994. *An Introduction to the Bootstrap*, 1st edition.

- Chapman & Hall/CRC, New York, USA. <https://doi.org/10.1201/9780429246593>
- Ori Ernst, Avi Caciularu, Ori Shapira, Ramakanth Pasunuru, Mohit Bansal, Jacob Goldberger, and Ido Dagan. 2022. Proposition-level clustering for multi-document summarization. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1765–1779, Seattle, United States. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.naacl-main.128>
- Ori Ernst, Ori Shapira, Aviv Slobodkin, Sharon Adar, Mohit Bansal, Jacob Goldberger, Ran Levy, and Ido Dagan. 2024. The power of summary-source alignments. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 6527–6548, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.findings-acl.389>
- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1102>
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2021. The pile: An 800gb dataset of diverse text for language modeling. *CoRR*, abs/2101.00027v1.
- Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. 2023. RARR: Researching and revising what language models say, using language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16477–16508, Toronto, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.acl-long.910>
- Beliz Gunel, Sandeep Tata, and Marc Najork. 2023. Strum: Extractive aspect-based contrastive summarization. In *Companion Proceedings of the ACM Web Conference 2023, WWW '23 Companion*, page 28–31, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/3543873.3587304>
- Jia Guo, Liying Cheng, Wenxuan Zhang, Stanley Kok, Xin Li, and Lidong Bing. 2023. AQE: Argument quadruplet extraction via a quad-tagging augmented generative approach. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 932–946, Toronto, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-acl.59>
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. In *International Conference on Learning Representations*.
- Eran Hirsch, Aviv Slobodkin, David Wan, Elias Stengel-Eskin, Mohit Bansal, and Ido Dagan. 2025. Laquer: Localized attribution queries in content-grounded generation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vienna, Austria. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.acl-long.746>
- Ella Hofmann-Coyle, Mayank Kulkarni, Lingjue Xie, Mounica Maddela, and Daniel Preotiuc-Pietro. 2022. Extractive entity-centric summarization as sentence selection using bi-encoders. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 326–333, Online only. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.aacl-short.40>
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. spaCy: Industrial-strength natural language processing in Python. *Zenodo*. <https://doi.org/10.5281/zenodo.1212303>

- George Hripcsak and Adam S. Rothschild. 2005. Agreement, the f-measure, and reliability in information retrieval. *Journal of the American Medical Informatics Association : JAMIA*, 12(3):296–298. <https://doi.org/10.1197/jamia.M1733>, PubMed: 15684123
- Neel Jain, Ping yeh Chiang, Yuxin Wen, John Kirchenbauer, Hong-Min Chu, Gowthami Somepalli, Brian R. Bartoldson, Bhavya Kailkhura, Avi Schwarzschild, Aniruddha Saha, Micah Goldblum, Jonas Geiping, and Tom Goldstein. 2023. Neftune: Noisy embeddings improve instruction fine-tuning.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.550>
- Bryan Klimt and Yiming Yang. 2004. The enron corpus: A new dataset for email classification research. In *Machine Learning: ECML 2004*, pages 217–226, Berlin, Heidelberg. Springer Berlin Heidelberg. [https://doi.org/10.1007/978-3-540-30115-8\\_22](https://doi.org/10.1007/978-3-540-30115-8_22)
- Abdullatif Köksal, Timo Schick, Anna Korhonen, and Hinrich Schuetze. 2024. Long-Form: Effective instruction tuning with reverse instructions. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7056–7078, Miami, Florida, USA. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.findings-emnlp.414>
- Kundan Krishna, Prakhar Gupta, Sanjana Ramprasad, Byron Wallace, Jeffrey Bigham, and Zachary Lipton. 2023. USB: A unified summarization benchmark across tasks and domains. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8826–8845, Singapore. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-emnlp.592>
- Sayali Kulkarni, Sheide Chammas, Wan Zhu, Fei Sha, and Eugene Ie. 2020. Aquamuse: Automatically generating datasets for query-based multi-document summarization. *CoRR*, abs/2010.12694v1.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466. [https://doi.org/10.1162/tacl\\_a\\_00276](https://doi.org/10.1162/tacl_a_00276)
- Mosh Levy, Alon Jacoby, and Yoav Goldberg. 2024. Same task, more tokens: the impact of input length on the reasoning performance of large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15339–15353, Bangkok, Thailand. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.acl-long.818>
- Patrick Lewis, Barlas Oguz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. 2020. MLQA: Evaluating cross-lingual extractive question answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7315–7330, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.653>
- Haoran Li, Arash Einolghozati, Srinivasan Iyer, Bhargavi Paranjape, Yashar Mehdad, Sonal Gupta, and Marjan Ghazvininejad. 2021. EASE: Extractive-abstractive summarization end-to-end using the information bottleneck principle. In *Proceedings of the Third Workshop on New Frontiers in Summarization*, pages 85–95, Online and in Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.newsum-1.10>
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173. [https://doi.org/10.1162/tacl\\_a.00638](https://doi.org/10.1162/tacl_a.00638)

- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3730–3740, Hong Kong, China. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1387>
- Prabir Mallick, Tapas Nayak, and Indrajit Bhattacharya. 2023. Adapting pre-trained generative models for extractive question answering. In *Proceedings of the Third Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 128–137, Singapore. Association for Computational Linguistics.
- Yuning Mao, Yanru Qu, Yiqing Xie, Xiang Ren, and Jiawei Han. 2020. Multi-document summarization with maximal marginal relevance-guided reinforcement learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1737–1751, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.136>
- Moran Mizrahi, Guy Kaplan, Dan Malkin, Rotem Dror, Dafna Shahaf, and Gabriel Stanovsky. 2024. State of what art? A call for multi-prompt LLM evaluation. *Transactions of the Association for Computational Linguistics*, 12:933–949. [https://doi.org/10.1162/tacl\\_a\\_00681](https://doi.org/10.1162/tacl_a_00681)
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2017. SummaRuNNer: A recurrent neural network based sequence model for extractive summarization of documents. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1). <https://doi.org/10.1609/aaai.v31i1.10958>
- E. W. Noreen. 1989. *Computer-Intensive Methods for Testing Hypotheses: An Introduction*. Wiley.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély,

- Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, C. J. Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2023. GPT-4 technical report. *CoRR*, abs/2303.08774v1. <https://doi.org/10.48550/ARXIV.2303.08774>
- Mihir Parmar, Hanieh Deilamsalehy, Franck Dernoncourt, Seunghyun Yoon, Ryan A. Rossi, and Trung Bui. 2024. Towards enhancing coherence in extractive summarization: Dataset and experiments with LLMs. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 19810–19820, Miami, Florida, USA. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.emnlp-main.1106>
- Jonathan Pilault, Raymond Li, Sandeep Subramanian, and Chris Pal. 2020. On extractive and abstractive neural document summarization with transformer language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9308–9319, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.748>
- Abhilash Potluri, Fangyuan Xu, and Eunsol Choi. 2023. Concise answers to complex questions: Summarization of long-form answers. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9709–9728, Toronto, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.acl-long.541>
- Jack W. Rae, Anna Potapenko, Siddhant M. Jayakumar, Chloe Hillier, and Timothy P. Lillicrap. 2020. Compressive transformers for long-range sequence modelling. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26–30, 2020*. OpenReview.net.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D16-1264>
- Allen Roush and Arvind Balaji. 2020. DebateSum: A large-scale argument mining and summarization dataset. In *Proceedings of the 7th Workshop on Argument Mining*, pages 1–7, Online. Association for Computational Linguistics.
- Swarnadeep Saha, Shiyue Zhang, Peter Hase, and Mohit Bansal. 2023. Summarization programs: Interpretable abstractive summarization with neural modular trees. In *The Eleventh*

- International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1–5, 2023*. OpenReview.net.
- Gaurav Sahu, Olga Vechtomova, and Issam H. Laradji. 2025. A guide to effectively leveraging LLMs for low-resource text summarization: Data augmentation and semi-supervised approaches. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 1584–1603, Albuquerque, New Mexico. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2025.findings-naacl.86>
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2024. Quantifying language models’ sensitivity to spurious features in prompt design or: How I learned to start worrying about prompt formatting. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7–11, 2024*. OpenReview.net.
- Aviv Slobodkin, Eran Hirsch, Arie Cattan, Tal Schuster, and Ido Dagan. 2024. Attribute first, then generate: Locally-attributable grounded text generation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3309–3344, Bangkok, Thailand. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.acl-long.182>
- Christian Stab and Iryna Gurevych. 2014. Annotating argument components and relations in persuasive essays. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1501–1510, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Kawshik Manikantan Sundar, Shubham Toshniwal, Makarand Tapaswi, and Vineet Gandhi. 2024. Major entity identification: A generalizable alternative to coreference resolution. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11679–11695, Miami, Florida, USA. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.emnlp-main.652>
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, and Jason Wei. 2023. Challenging BIG-bench tasks and whether chain-of-thought can solve them. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13003–13051, Toronto, Canada. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-acl.824>
- Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A heterogeneous benchmark for zero-shot evaluation of information retrieval models. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: A large-scale dataset for fact extraction and VERification. Marilyn Walker, Heng Ji, and Amanda Stent, editors, In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics. <https://doi.org/10.18653/v1/N18-1074>
- Erik F. Tjong Kim Sang and Sabine Buchholz. 2000. Introduction to the CoNLL-2000 shared task chunking. In *Fourth Conference on Computational Natural Language Learning and the Second Learning Language in Logic Workshop*. <https://doi.org/10.3115/1117601.1117631>
- David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. 2020. Fact or fiction: Verifying scientific claims. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7534–7550, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.609>
- Hongning Wang, Yue Lu, and Chengxiang Zhai. 2010. Latent aspect rating analysis on review

- text data: A rating regression approach. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '10, page 783–792, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/1835804.1835903>
- Yiming Wang, Jindong Zhang, Zhiyao Yang, Bing Wang, Jingyi Jin, and Yitong Liu. 2024. Improving extractive summarization with semantic enhancement through topic-injection based bert model. *Information Processing & Management*, 61(3):103677. <https://doi.org/10.1016/j.ipm.2024.103677>
- Kam-Fai Wong, Mingli Wu, and Wenjie Li. 2008. Extractive summarization using supervised and semi-supervised learning. In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008)*, pages 985–992, Manchester, UK. Coling 2008 Organizing Committee. <https://doi.org/10.3115/1599081.1599205>
- Yuping Wu, Ching-Hsun Tseng, Jiayu Shang, Shengzhong Mao, Goran Nenadic, and Xiao-Jun Zeng. 2023. EDU-level extractive summarization with varying summary lengths. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1655–1667, Dubrovnik, Croatia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-eacl.123>
- Derong Xu, Wei Chen, Wenjun Peng, Chao Zhang, Tong Xu, Xiangyu Zhao, Xian Wu, Yefeng Zheng, Yang Wang, and Enhong Chen. 2024. Large language models for generative information extraction: A survey. *Frontiers of Computer Science*, 18(6):186357. <https://doi.org/10.1007/s11704-024-40555-y>
- Yumo Xu and Mirella Lapata. 2020. Coarse-to-fine query focused multi-document summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3632–3645, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.296>
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2024. Qwen2.5 technical report. *CoRR*, abs/2412.15115v2. <https://doi.org/10.48550/ARXIV.2412.15115>
- Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023. Extractive summarization via ChatGPT for faithful summary generation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3270–3278, Singapore. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2023.findings-emnlp.214>
- Wenxuan Zhou, Sheng Zhang, Yu Gu, Muhao Chen, and Hoifung Poon. 2024. Universalner: Targeted distillation from large language models for open named entity recognition. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7–11, 2024*. OpenReview.net.
- Ming Zhu, Aman Ahuja, Da-Cheng Juan, Wei Wei, and Chandan K. Reddy. 2020. Question answering with long multiple-span answers. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3840–3849, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.342>

## A IGCS-BENCH Details

In this section we provide details on creating IGCS-BENCH, introduced in Section 3, based on the six CS datasets. We split every dataset into train, development and test sets except for our two datasets from Ernst et al. (2024) that only include a test set. The full prompt template that we drafted for the tasks is shown in Figure 3.

Table 6 presents the desiderata for the tasks. Specifically, it displays: (1) model scores for each task, as reported in the respective paper. We find

**IGCS-BENCH Prompt Template**

{① instruction}. Output the exact {② selection\_type} from the given {③ source\_type} as a valid json array of strings. Do not change the copied text.

document #0:  
Document 0 text...  
document #1:  
Document 1 text...  
...

Figure 3: The prompt template for IGCS-BENCH tasks. The ① instruction for each task is detailed in Table 1. The ② selection\_type is “sentences” for ASPSEL, ASPSUM, EVIDSENT, and “text phrases” for the other tasks. The ③ source\_type varies by task: “document” for single-document tasks (ASPSUM and ARGMINE), “abstract(s)” for EVIDSENT, and “documents” for the remaining three multi-document tasks.

that there is substantial **room for improvement** on these datasets. (2) Inter annotator agreement of the annotations when the datasets were curated. The high Cohen’s  $\kappa$  scores are an indication for **the quality of the data**.

**Evidence Retrieval (EVIDSENT).** We randomly split the original train set from SCIFACT (Wadden et al., 2020) into train and a development sets with 687 and 122 instances, respectively. The original development set is used as our test set, since the original SCIFACT test set is gated behind a leaderboard submission system.

For evaluation, we followed the original paper and used sentence-level  $F_1$ . Since the model selects tokens, we first identify the sentences that contain them, which are then used for computing the sentence-level metric.

**Saliency and Proposition-level Evidence Detection (SALIENCE and EVIDPROP).** From the SPARK (Ernst et al., 2024) dataset, we only utilize the respective test sets. Ernst et al. (2024) originally used an automatically derived and lower quality training set. SALIENCE and EVIDPROP include 98 and 1,332 instances in their test sets, respectively.

For evaluation, we followed the original paper and used token-level  $F_1$  for both tasks. This is the same as the generic metric that we propose to apply, in §3.3.

Task	Cohen’s $\kappa$	Best-reported	Best-GENCS
ASPSUM	†–	45.0	37.0
ARGMINE	–	‡38.5	36.7
EVIDSENT	0.71	44.0	48.1
ASPSSEL	0.64	34.4	49.3
SALIENCE	0.72	31.0	37.5
EVIDPROP	0.72	32.0	35.6

Table 6: Reported inter-annotator agreement and performance figures for the original datasets incorporated in our IGCS-BENCH benchmark (whose details appear in §3.2). The Cohen’s  $\kappa$  figures specify the reported inter-annotator agreement levels for the respective dataset, when available. Best-reported are the previously reported results for each dataset, using the original evaluation measure proposed for each dataset. We quote here performance figures in the transfer learning setting (except for ASPSUM, see dagger), as this is the setting on which we focus in this paper, with respect to the utility of our generic GENCS training dataset. Finally, for comparison, Best-GENCS presents the best performance obtained in our experiments when utilizing our GENCS training set in the transfer-only setting (maximum value among rows 4 and 5 in Table 4). As shown in the table, our transfer-only results improve over the prior performances for 4 of the datasets (while direct comparison is not available for ARGMINE, see double-dagger). †: Inter-annotator agreement was not reported for this task due to its subjectivity; instead, the dataset includes five reference selections per instance, where model performances are computed against each of them. ‡: Here we quote the best supervised results, rather than transfer learning results, since the latter setting was not previously attempted for this dataset. Further, the original test split for this dataset was not released, hence this figure is not fully comparable to our result (in the last column), which was computed on our introduced test split.

**Aspect-based Sentence Selection (ASPSSEL).** OPENASP (Amar et al., 2023) defines a sentence selection task. We use the original test set with 27 samples and split the original development set into training and development sets with 13 and 11 samples, respectively.

For evaluation, we follow the originally proposed sentence-level micro- $F_1$  metric. Similar to the case of EVIDSENT, we first identify the sentences of the selected tokens, and then compute the measurement.

**Extractive Aspect-based Summarization (ASP-SUM).** For the extractive aspect-based summarization task (single document), Ahuja et al. (2022) annotated 100 documents from two topics, each with two aspects, yielding a total of 400 document-aspect-summary instances. We use the *Fraud* topic and its two aspects – *penalty* and *nature* – as the test set, and split the remaining 100 documents from the Earthquake topic into 160 training and 40 development instances. ASPECTNEWS has 5 annotations per document; we retain these as 5 separate gold references and evaluate them following the multiple-reference selection approach (§3.3).

For evaluation, we use the metric from the original paper, which computes sentence-level  $F_1$  against references with soft labels. As in the other sentence-level evaluations, the sentences used for the evaluation are those that contain the tokens selected by a model.

**Argument Mining (ARGMINE).** DebateSum (Roush and Balaji, 2020) reported results on a test set of 18,738 instances but did not provide the original dataset splits, as confirmed in both our review and the project’s official repository.<sup>18</sup> To support reproducible research, we propose new training, development, and test splits for DebateSum. The dataset, originally sourced from the debate.cards website, spans 2013–2019 with a total of 187,386 instances. We allocate years 2013–2016 for training, 2017 for development, and 2018–2019 for testing. To avoid contamination, we filter instances by removing any that share identical abstract summaries, extractive summaries, or full document fields across splits. For IGCS-BENCH, we randomly sample 1,000 instances from each split.

For evaluation, we followed the original paper and used ROUGE-2- $F_1$  as our primary metric.

## B Computing Confidence Intervals for Token-level $F_1$

We estimated 95% confidence intervals via bootstrap resampling (Efron and Tibshirani, 1994) with 10,000 iterations. At each iteration, we drew instances with replacement from each of the six IGCS-BENCH datasets and computed the corresponding  $F_1$  score for that bootstrap sample.

<sup>18</sup><https://github.com/Hellisotherpeople/DebateSum/issues/3>.

## C GENCS Details

### C.1 Data Collection

From Multi-News (Fabbri et al., 2019), English Wikipedia, PubMed, books (Rae et al., 2020), and hotel reviews (Wang et al., 2010), we uniformly sampled 500 documents, each with between 350 and 3500 tokens, as counted with `nltk.word_tokenize`. From Enron (Klimt and Yang, 2004), we sampled 250 email threads containing multiple emails and 250 threads with a single email.

From GitHub, we sampled one million source code files in one of the following 15 programming languages: Assembly, C, C#, C++, GO, Java, JavaScript, PHP, Perl, Python, Ruby, Rust, Scala, Shell, TypeScript, each with a permissive license of either APACHE-2.0 or MIT. Next, we grouped files from the same repository and folder into multi-source task

**GENCS Instruction Generation Prompt**

**System:** You are a manager of a {software|publishing} firm

**User:** You are a manager of a {software|publishing} firm and you are required to train the best students on how to perform {code|content} selection from given sources. Write 5 short instructions for selecting {code|content} from the given {file|document}(s) to challenge students and train them on how to select relevant {code|content} based on diverse instructions.

Guidelines:

1. You must keep the 5 instructions short and concise - as a single sentence.
2. Instructions must start with the words "Select {code|content}" as they are always for selecting {code|content} from the {file|document}(s) and never for writing a new {code|text} nor paraphrasing the original content.
3. Instructions should not be too specific that hint on the answer and not too vague that cannot be fulfilled.
4. Write the instructions as a numbered list.

{Source File|Document} #0:  
Document 0 text...

{Source File|Document} #1:  
Document 1 text...

...

Figure 4: The prompt template for annotating IGCS instructions is as follows: From each choice of {①|②}, we use ① for the GitHub code dataset and ② for all other datasets. For empty selection annotation, the second sentence is changed as follows: “Write 5 short instructions that have no matching {code|content} from the given {file|document}(s) to challenge students and train them to avoid selecting {code|content} when none matches the instruction.”

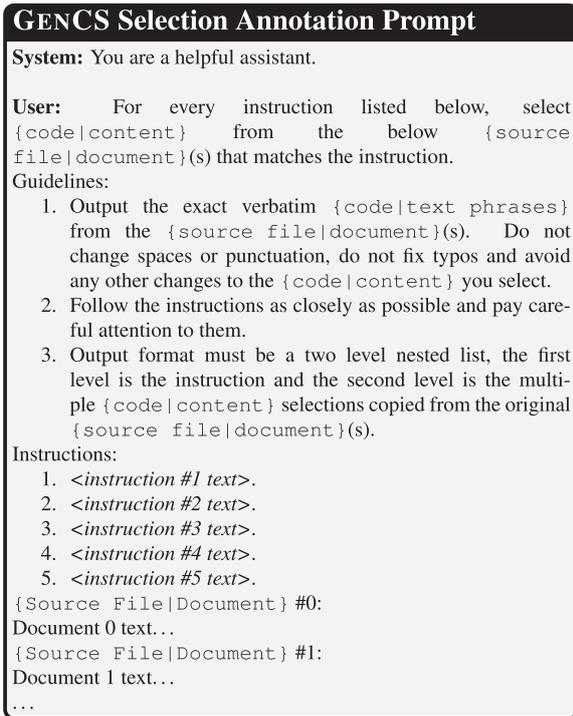


Figure 5: The prompt template for annotating selections. From every choice of {①|②}, we use ① for the GitHub code dataset and ② otherwise.

instances. Finally, we sampled 250 multi-source task instances and 250 single-instance source files, resulting in 500 GitHub samples.

## C.2 Synthesizing Instructions and Annotating Selections

For synthesizing instructions (Step 1 in §4.1), the prompt is presented in Figure 4. For synthesizing possible content selections (Step 2), the prompt is presented in Figure 5.

We intentionally under-specified the definition of an instruction to encourage the generation of diverse instructions.

## D Synthetic Pipeline Configuration Variants Details

In Table 5, we compare models fine-tuned on different pipeline configurations. For GENCS<sub>1-step</sub>, we used GPT-4 as the annotator, issuing a single prompt of combined instruction and guidelines from Figure 4 and Figure 5. We utilized the annotations for GENCS to derive GENCS<sub>1-model</sub> and GENCS<sub>1-inst</sub>. For GENCS<sub>1-model</sub>, we used GPT-4 as the single selection model, while for GENCS<sub>1-inst</sub>, we used only the first generated instruction.

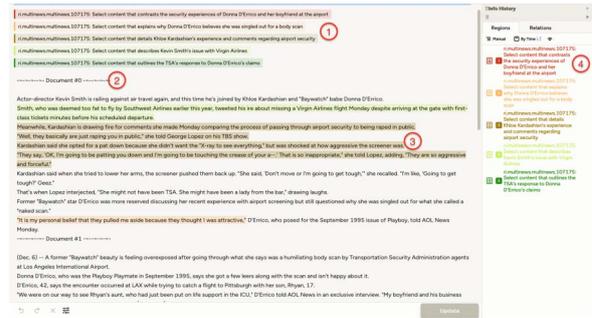


Figure 6: Our annotation interface for manual content selection, based on Label Studio (<https://labelstud.io>). In ①, the five instructions for the document set can be selected in order to enable highlighting of text spans in the source ②, such as the highlighted span shown in ③. The highlights are added to the selected spans panel marked in ④. The text within the figure is for illustrative purposes only and does not need to be read.

## E GENCS Manual Quality Assessment

The guidelines for rating the instructions' *naturalness* and *specificity* are in Figure 7. The annotation interface for selecting spans based on instructions is presented in Figure 6.

**Inter-rater Agreement of Selections.** To quantify the agreement, we cannot directly use Cohen's  $\kappa$  as it requires the number of negative cases, which is ill-defined (or very large) in the content selection setting (e.g., all the spans that are not part of the selection). Instead, we followed Hripcsak and Rothschild (2005), who demonstrated that as the number of possible negative cases increases,  $\kappa$  approaches the average  $F_1$  score, which is the case in our content selection setting. Based on this, to compute  $\kappa$  between two groups of annotators, we computed the average pairwise token-level  $F_1$  score (§3.3) between every possible pair of annotators, where each annotator originates from a different group.

## F Model Configurations

We trained Llama-3-8B meta-llama/Meta-Llama-3-8B-Instruct with various data mixtures, as described in §5.1. Each dataset training mix was shuffled before training.

For training, we first experimented with several hyperparameter options. We also tuned the prompt manually to make it work better for the zero-shot case. For models trained on multiple training sets, we attempted to balance the different sets by

## GENCS Manual Annotation Guidelines for Instructions and Selections

All scores are Likert scores, integers between 1–5.

### Instruction Quality and Diversity

We wish to measure quality and diversity of the generated instructions by manually rating the Naturalness (for quality), and specificity (for diversity).

1. **Naturalness** – The instruction is clear, fluent, plausible, and relevant to the context of the document set and its topic.
  - **Clear** – The instruction is unambiguous and understandable, given the topic.
  - **Fluent** – The instruction is written in natural, human-like language and adheres to proper grammar.
  - **Plausible** – The instruction appears to be written by a human and does not resemble machine-generated instructions.
  - **Contextually Relevant** – The instruction aligns with the context of the document set and reflects a use case that a human might reasonably formulate.
    - Based on general knowledge or expertise (e.g., as an NLP student), the use case appears likely.
  - Example of low naturalness – “*Select content that lists the factors increasing the likelihood of resumption of ovarian cyclicity (ROC) at 36 to 42 days in milk (DIM) from Document #1*”.
2. **Specificity** – The information sought by the instruction is central (1) or anecdotal (5) to the topic of the document set.
  - To assess specificity, imagine a conceptual tree of all information organized hierarchically within the document set. The measure reflects how deep into the tree the instruction seeks information.
  - **Levels:**
    - (a) **Topic Level** – The instruction addresses the broad topic of the document set.
    - (b) **Sub-Topic Level** – The instruction focuses on a specific sub-topic within the broader topic.
    - (c) **Somewhat Central** – The instruction targets a detail that is central but not overly specific.
    - (d) **Anecdotal Point** – The instruction seeks information tied to a specific detail or anecdote.
    - (e) **Atomic Anecdotal Point** – The instruction zeroes in on a singular, highly specific detail related to the topic.

### Content Selection

Select the relevant text spans from the document set based on the given instruction.

Figure 7: Manual annotation guidelines for *naturalness* and *specificity* for rating GENCS instructions.

up-sampling smaller datasets to match the size of the largest dataset in the mix. We found this approach to be inferior and as such abandoned this technique.

The input prompt is shown in Figure 3. The target output is a JSON array of the selected texts.

All tested models accommodated the input size except for two ASPSEL instances on Llama-3-8B and Llama-3-70B without document-level inference. In those cases, we truncated the input to fit the 8K token context size.

All model variants were trained on three A100 GPUs with a maximum sequence length of 4096, batch size of 4, NEFTune noise  $\alpha = 5.0$  (Jain et al., 2023), and a warmup ratio of 0.06 for 3 epochs. Training each model variant typically took several hours on the next token prediction task. We trained each model on the next-token completion task, ignoring system and user prompts in the loss function.

During the decoding inference phase, all models (GPT-4, Claude-3-Opus, Gemini-1.5-Pro, and the three versions of Llama-3) were set to a temperature of 0.0 to ensure reproducibility. We set `max_new_tokens` to 2048 for Llama-3-8B and its trained variants.

For the Llama-3 70B and 405B models, we used Together AI’s API<sup>19</sup> with the model versions `meta-llama/Llama-3-70b-chat-hf` and `meta-llama/Meta-Llama-3.1-405B-Instruct-Turbo`, respectively.

For the Qwen 2.5 family, we used the `instruct` variants (e.g., `Qwen/Qwen2.5-7B-Instruct`) of the four smallest models—0.5B, 1.5B, 3B, and 7B. Similarly, for the SmolLM2 family, we used the `instruct` variants (e.g., `HuggingFaceTB/SmolLM2-1.7B-Instruct`) of all three models: 135M, 360M, and 1.7B.

## G Fuzzy Match

We describe here the fuzzy match grounding algorithm, mentioned in §5, which grounds the output text to a location in the source text. Since LLMs sometimes paraphrase or slightly alter copied text spans, we relax the exact (case-insensitive) textual search. We consider the output text to be matched to a (non-empty) sub-sequence in the source text when there is a token-level Levenshtein distance of up to 15% the length of the output (and no more than 10 tokens) between the two strings.

<sup>19</sup><https://www.together.ai>.

Task	Pearson $r$	Spearman $\rho$	Kendall $\tau$
EVIDSENT	0.966	0.955	0.83
ASPSEL	0.967	0.976	0.886
ASPSUM	0.867	0.85	0.697
ARGMINE	1.0	0.999	0.993

Table 7: System-level correlations between the original evaluation metrics used in four of the benchmark tasks and the proposed token-level  $F_1$  metric (Section 3.3), which we suggest as a standardized metric for IGCS tasks. Overall, the token-level  $F_1$  demonstrates strong to very strong correlation with all other metrics, indicating its reliability for evaluating content selection performance.

If there are multiple matches, we select the one with the lowest edit distance that appears first. When no such match is found, the text span is discarded from the suggested selection. We tested larger Levenshtein distances, which result in a rapidly growing computational overhead, but have minimal or no positive effect.

## H Correlation of Evaluation Metrics - Complementary Results

In Table 7, we report the system-level correlation coefficients between each task-specific metric and the generic token-level metric (§3.3). These coefficients are computed using system-level scores from 24 model configurations for ARGMINE and ASPSUM, and from 35 configurations for ASPSEL and EVIDSENT. The configurations vary along three dimensions: the choice of fine-tuning training set, the use or omission of in-context learning, and the application of document-level inference (the latter applicable only to ASPSEL and EVIDSENT).

## I Transfer Learning - Complementary Results

In Table 9 we show the original token-level  $F_1$  scores for the six IGCS-BENCH tasks.

## J Significance Testing

As some task-specific metrics are defined at the system level, we used a permutation test to measure the significance (at  $p < 0.05$ ) of the difference between two model scores, performing random sampling with a size of 1,000 for each test (Noreen, 1989).

## K Document-level Inference: Complementary Results

In Table 10, we show the scores of nine models tested on the four multi-document tasks in IGCS-BENCH, from which Figure 2 is derived.

## L Prompt Robustness Results

To further validate our findings in the transfer-only setting (discussed in §6.1.1 and exhibited in Figure 1), we conducted each experiment with two additional prompt variants automatically tuned via meta-prompting using OpenAI’s O4-mini-high.<sup>20</sup> The resulting prompt variants and the default in-context learning prompt are exemplified for a single task in Table 8. Notably, these variants are not mere paraphrases of the original instructions but substantially different in terms of length and details for each of the six tasks. Thus they provide more variability for the analysis. For the in-context learning scenario, one of the prompt variants (V1) used a one-shot example instead of two-shot, to additionally explore different number of samples. The fine-tuned models, namely GENCS<sub>Union</sub> and GENCS<sub>Majority</sub>, were fine-tuned using the default prompt, suggesting adaptability of the fine-tuned models to prompt variations.

In Figure 8, we observe trends similar to those in Figure 1, with significant performance gains across most models under different prompt variants. Moreover, we note variability in model performance across these prompts, a phenomenon reported in prior works (Sclar et al., 2024; Mizrahi et al., 2024).

<sup>20</sup><https://openai.com/index/introducing-o3-and-o4-mini>.

---

**AspSEL's Zero-shot Prompt Variant 1**

---

Select between 1 and 3 sentences from the provided news article that are most relevant to the {topic}'s {aspect.description}.  
Read the news article carefully and identify all sentences. Internally analyze each sentence for relevance to the specified topic and aspect.  
Perform your reasoning steps before arriving at a final conclusion, but only output the final result. Do not modify or alter any of the selected sentences.  
# Steps 1. **Parse the Article:** Break the article into individual sentences. 2. **Internal Reasoning:** Evaluate each sentence for its relevance to the {topic}'s {aspect.description} based on the context and details provided. 3. **Selection:** Choose at least 1 and at most 3 sentences that best capture the required information. 4. **Conclusion:** Prepare your final selection after completing your internal reasoning.  
# Output Format Output a valid JSON array of strings containing the exact selected sentence(s). For example: ["Sentence 1", "Sentence 2"]  
# Notes - Do not change or rephrase any of the copied sentence(s). - Ensure that your internal reasoning process is used to determine the selection, but do not include it in the final output.

---

**AspSEL's Zero-shot Learning Prompt Variant 2**

---

From the news article below, pick between 1 and 3 sentences that best address the {topic}'s {aspect.description}. Return them verbatim as a valid JSON array of strings.

---

**AspSEL's In-context Learning Default Prompt**

---

Given the following news article, select at least 1 and at most 3 sentences that are the most relevant to the given aspect. Output the exact sentences from the given document as a valid json array of strings. Do not change the copied text. Below {is an example | are examples} of an input and the correct selected content:  
Aspect: {example1\_topic}'s {example1\_aspect\_description}  
Input Document(s):  
{example1\_documents}  
...  
— END OF EXAMPLES —  
Now, select content from the below document(s):  
Aspect: {topic}'s {aspect.description}  
Input Document(s):  
{documents}

---

**AspSEL's In-context Learning Prompt Variant 1**

---

From the news article below, pick between 1 and 3 sentences that best address the aspect. Return them verbatim as a valid JSON array of strings.  
Below {is an example | are examples} of an input and the correct selected content:  
Aspect: {example1\_topic}'s {example1\_aspect\_description}  
Input Document(s):  
{example1\_documents}  
...  
— END OF EXAMPLES —  
Now, select content from the below document(s):  
Aspect: {topic}'s {aspect.description}  
Input Document(s):  
{documents}

---

Table 8: Zero-shot and In-context learning prompt variants for AspSEL. The default zero-shot prompt template can be found in Figure 3.

Models		ASPSUM		ARGMINE		EVIDSENT		ASPSSEL		SALIENCE	EVIDPROP
		F1	O	F1	O	F1	O	F1	O	F1/O	F1/O
transfer- only	Llama-3-8B <sub>ICL</sub>	59.2	34.6	46.6	46.9	58.0	44.8	27.1	28.6	42.9	13.5
	Llama-3-8B	56.1	29.4	42.4	42.4	46.6	47.5	42.3	41.9	36.6	27.3
	+ L <sub>OO</sub>	41.6	34.0	29.3	29.1	23.8	25.8	32.0	30.4	41.9	10.2
	+ GENCS <sub>Union</sub>	63.3	37.0	36.6	36.7	51.3	42.6	52.0	49.3	37.5	33.6
	+ GENCS <sub>Majority</sub>	57.6	35.6	25.2	25.2	58.9	48.1	49.7	47.1	32.4	35.6
prompt-based+ supervision+ transfer	Llama-3-8B <sub>Sup</sub>	69.5	40.6	64.7	63.5	74.4	66.0	–	–	–	–
	+ L <sub>OO</sub>	72.4	42.3	65.3	64.1	78.9	70.0	–	–	–	–
	+ GENCS <sub>Union</sub>	72.7	42.7	64.9	63.7	81.1	72.1	–	–	–	–
	+ GENCS <sub>Majority</sub>	75.4	43.2	64.7	63.6	79.8	68.8	–	–	–	–
	prompt-based models	Llama-3-70B	50.5	29.3	41.1	40.7	65.2	58.5	51.9	56.8	33.2
Llama-3-405B		53.7	30.0	45.5	45.4	61.8	56.2	57.8	59.8	35.1	42.1
GPT-4		60.4	32.8	39.5	39.0	63.6	58.6	55.4	57.4	39.1	50.1
GPT-4 <sub>ICL</sub>		63.8	33.9	46.0	45.5	63.2	57.3	52.9	55.0	39.9	47.5
Claude3-Opus		51.7	31.3	50.0	49.6	57.2	54.5	53.9	52.2	43.7	28.2

Table 9: Comparison between our proposed token-level  $F_1$  (§3.3) and the original metric defined for each task, denoted as **O**. For SALIENCE and EVIDPROP the original metric is the same. Model configurations and fine-tuning details are described in §5 and follow the same notations as Table 4.

Model	EVIDPROP			EVIDSENT			SALIENCE			ASPSSEL		
	Score	$\Delta$	$\Delta$ Sel.	Score	$\Delta$	$\Delta$ Sel.	Score	$\Delta$	$\Delta$ Sel.	Score	$\Delta$	$\Delta$ Sel.
Llama-3-8B	29.6	-2.4	+64.9	44.8	+2.7	+7.2	31.7	+4.9	+142	18.4	+23.4	+741
Llama-3-70B	42.8	+2.1	+12.7	57.5	+1.0	+0.8	28.9	+4.3	+126	24.6	+32.2	+478
Llama-3-405B	40.4	+1.6	+34.8	56.7	-0.5	+1.0	31.7	+3.4	+192	43.8	+16.0	+454
Llama-3 <sub>GENCS<sub>Union</sub></sub>	32.1	+1.5	+21.9	42.0	+0.6	+4.3	31.0	+6.6	+282	29.5	+19.8	+643
Llama-3 <sub>GENCS<sub>Majority</sub></sub>	33.6	+2.0	+21.3	47.8	+0.2	+3.3	21.3	+11.0	+146	18.7	+28.4	+578
Llama-3 <sub>LOO+GENCS<sub>Union</sub></sub>	27.9	-8.3	+213.7	70.6	+0.2	+0.8	37.9	+3.4	+260	20.9	+21.9	+992
Llama-3 <sub>LOO+GENCS<sub>Majority</sub></sub>	25.9	-10.6	+289.5	69.6	+0.5	+0.1	32.6	+8.4	+352	21.8	+23.3	+795
GPT-4	51.7	-1.6	+21.5	59.1	-0.5	+2.4	32.9	+6.2	+246	39.3	+18.1	+485
Claude3-Opus	36.8	-8.6	+79.7	53.7	+0.8	+7.1	36.6	+7.1	+247	35.9	+16.3	+564
Average		-2.7	+84.4		+0.6	+3.0		+6.1	+222		+22.2	+636

Table 10: The performance of the four multi-document tasks, detailed in §6.2, is presented as task-specific metric scores. **Score** refers to the task score without applying document-level inference (§5.3);  $\Delta$  indicates the performance gain when employing the document-level inference method;  $\Delta$  Sel. indicates the difference in the size of the model’s output, measured as the average number of tokens in the selection. The last row shows the average difference across all models for each task.

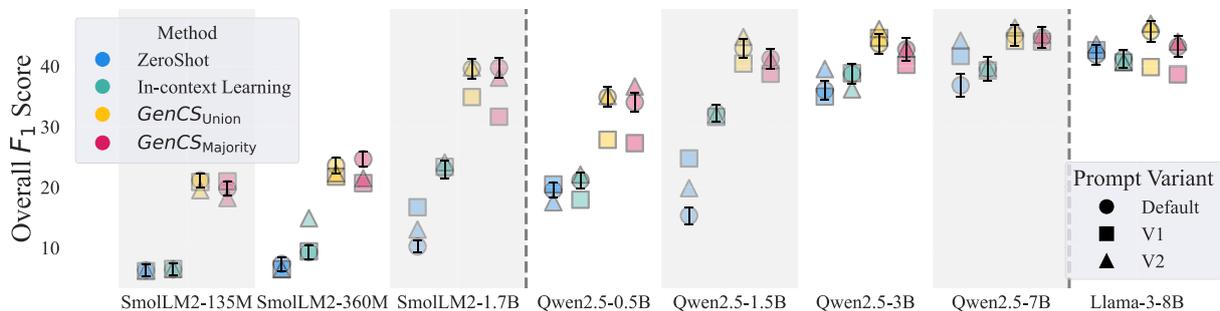


Figure 8: Overall  $F_1$  scores on IGCS-BENCH for eight models, evaluated across four methods and three prompt variants (see examples for the prompt variants in Table 8). For comparison, see Figure 1, which presents results using only the default prompt. Confidence intervals ( $\alpha = 0.05$ ) are shown for the default prompt. While performance varies across prompt variants, the overall trends remain consistent with those reported in Section 6.1: most models benefit from GENCS fine-tuning, with the smallest models exhibiting the largest gains.