

SmartBook: AI-Assisted Situation Report Generation

Revanth Gangi Reddy¹, Yi R. Fung¹, Qi Zeng¹, Manling Li¹, Ziqi Wang¹, Paul Sullivan and Heng Ji¹

¹Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, 61820, IL, USA.

Contributing authors: revanth3@illinois.com;
yifung2@illinois.edu; qizeng2@illinois.edu; manling2@illinois.edu;
ziqiw9@illinois.edu; hengji@illinois.edu;

Abstract

Emerging events, such as the COVID pandemic and the Ukraine Crisis, require a time-sensitive comprehensive understanding of the situation to allow for appropriate decision-making and effective action response. Automated generation of situation reports can significantly reduce the time, effort, and cost for domain experts when preparing their official human-curated reports. However, AI research toward this goal has been very limited, and no successful trials have yet been conducted to automate such report generation. Pre-existing natural language processing methods, large language model based text generation, and information retrieval techniques are insufficient to identify, locate, and summarize important information, and lack detailed, structured, and strategic awareness. We propose **SmartBook**, a novel task formulation targeting situation report generation, which consumes large volumes of news data to produce a structured situation report with multiple hypotheses (claims) summarized and grounded with rich links to factual evidence. We realize SmartBook for the Ukraine-Russia crisis by automatically generating intelligence analysis reports to assist expert analysts. The machine-generated reports are structured in the form of timelines, with each timeline organized by major events (or chapters), corresponding strategic questions (or sections) and their grounded summaries (or section content). Our proposed framework automatically detects real-time event-related strategic questions, which are more directed than manually-crafted analyst questions, which tend to be too complex, hard

to parse, vague and high-level. Results from thorough qualitative evaluations show that roughly 82% of the questions in Smartbook have strategic importance, with at least 93% of the sections in the report being tactically useful. Further, experiments show that expert analysts tend to add more information into the SmartBook reports, with only 2.3% of the existing tokens being deleted, meaning SmartBook can serve as a useful foundation for analysts to build upon when creating intelligence reports.*

Keywords: Complex Event Understanding, Knowledge Acquisition, Natural Language Generation, Intelligence Analysis Research & Development Toolkit

1 Introduction

In today's fast-paced world, staying informed is more important than ever. With the increasing availability of information sources, it can be challenging to separate fact from fiction. This is particularly true during times of crisis, when rumors and false information can spread quickly, causing unnecessary panic or confusion. In these situations, having access to reliable and up-to-date information can make all the difference. It can help individuals and organizations make informed and safe decisions, take appropriate actions, and even save lives. As such, the ability to understand the situation amidst the overwhelming amount of news and rumors becomes paramount.

Important events and emerging crises often come with a vast amount of news and rumors, which vary in credibility but offer diverse story coverage. Situation understanding amidst this information overload is extremely valuable for decision-makers, stakeholders, and the general public interest, alike. For example, timely updates about the COVID pandemic can help people better prepare for social distancing and prevention of spread. Staying informed about the Russia-Ukraine crisis can help organizations determine where humanitarian efforts are most needed, as well as how energy and food security are impacted across the world.

The key to situation understanding for such events is to extract the common truths across heterogeneous data that meaningfully impact action planning and strategic development. Along this direction, there are many expert analysts today who manually prepare *situation reports* about the state of affairs, potential risks or threats, and perspective, with recommended actions. While playing a foundational role in government policies and organization response, situation reports are time-consuming and resource-intensive to manually develop. They are also often limited to a few topics, scenarios, or

*We deeply regret that one of our project contributors, Paul, passed away before the publication of this paper. His contributions to this research were invaluable, and the loss is felt deeply by all of us who had the privilege of working with him. We honor the memory and express our gratitude for his enduring commitment to the pursuit of knowledge.

regions, and thus not scalable. Further, manual reports can be biased, incomplete, difficult to keep up-to-date, and incapable of consolidating information across multiple data sources and languages. Therefore, it is crucial to explore the automatic generation of situation reports to facilitate the timely, comprehensive, and organized presentation of information for prompt policy actions and tactical operations.

Question: Is Russian positioning forces for possible follow-on actions into Ukraine after the initial invasion?

ChatGPT: As an AI language model, I cannot provide real-time information on current events. However, it is essential for intelligence analysts to monitor the activities of various countries, including troop movements and military exercises, to assess potential threats and anticipate future actions. To determine if Russia is positioning forces for possible actions into Ukraine, one must analyze the latest intelligence reports, satellite imagery, and open-source information. Additionally, understanding the political and historical context of the region can provide valuable insight into the intentions of the countries involved.

SmartBook: The arrival of Russian forces in Slavutych came after several days of shelling against the city, which is strategically located close to the Dnieper River, about 95 miles (150 kilometers) north of Kyiv and close to the border with Belarus. Russia has been using Belarus as a springboard for many of its air operations in Ukraine, according to intelligence collected by NATO surveillance planes flying over the Polish-Ukrainian border and radar seen by CNN. So far, however, US officials have not seen Belarusian troops “being readied to move into Ukraine” or “that they are moving or are in Ukraine”, a senior US defense official said Monday, adding that the forces inside Ukraine are Russian.

Table 1: GPT-4 powered ChatGPT is incapable of generating situation reports because it has limited knowledge of up-to-date events. In comparison, Smartbook is able to generate timely, multi-source, and trustworthy answers.

Automatic situation report generation poses an exciting new problem domain that has not been addressed by previous research work. Existing approaches to news report generation typically lack details, structures, and a high-level strategic awareness for downstream end-user actionability. In particular, current large language model (LLM) based approaches such as ChatGPT [1] may return no answers or incorrect answers to important questions because they cannot process information in real-time or lack fact-checking mechanisms. Current multi-document news summarization tasks aim to condense the information overlap across a news article cluster into a succinct (~5 sentence) paragraph [2]. However, real-world situation reports that policymakers and military experts utilize as actual references for their decision-making require a lot more details. For example, recent human-written situation reports [3] on the Russia-Ukraine crisis, publicized online by the United Nations and the U.S. Congressional Research, aggregate critical information with regards to the prelude societal background; onset of the military invasion; development of subsequent military operation events; economic, political, and humanitarian ramifications; as well as outlook forecast – in over twenty dense pages. Moreover, generating a situation report is *not only* about increasing the sample size of relevant news source retrieval [4] and the length of subsequent summary generation, *but also* about maintaining a structured organization (e.g., chapters and sections) for complex event understanding, with an agenda for

strategic planning in mind. As such, the current application of AI techniques to situation analysis and report generation is far from mature.

In this work, we present **SmartBook**, a generalizable framework designed to assist human analysts in situation report generation. SmartBook ingests data from multiple sources to generate a comprehensive report with real-time information being updated in a timely fashion. The framework automatically discovers questions important for situation analysis and gathers salient information for generating the report. For each question, the report comprises summaries with tactical information coming from relevant claims within news articles, and grounded to the corresponding trustworthy factual evidence.

As a case study, we consider the case of automatic situation report generation for the Ukraine-Russia crisis. This scenario is particularly challenging, given the complexity of the conflict, the politicization of the issue, and the need to constantly update the reports as the situation evolves. Overall, manually creating situation reports for the Ukraine-Russia crisis requires analysts to have a combination of technical skills, analytical expertise, and strategic thinking. Our proposed automatic situation report is created from a corpus of news articles and is organized in the form of timelines, with new timelines being added in real-time as the conflict progresses. Each timeline covers a 2-week duration and major events within that duration serve as the corresponding chapters. The report incorporates a logical structure for each chapter by creating section headings in the form of automatically identified strategic questions, and providing grounded query-focused summaries within the individual sections. Figure 1 shows an example from SmartBook for the Ukraine-Russia crisis, with the structured hierarchy in the form of timelines, chapters and corresponding sections.

Our contributions can be summarized as follows:

- We identify an exciting new problem domain and task formulation of *automated situation report generation*, which is not meant to replace human analysts but rather to assist them.
- We propose **SmartBook**, an automated framework that generates comprehensive, real-time situation reports from various sources, by discovering critical strategic questions and grounding the report with trustworthy factual evidence.
- We quantitatively and qualitatively evaluate our automatic situation report generation approach for the Ukraine-Russia crisis and demonstrate that SmartBook is well-received by expert analysts in terms of informativeness and effectiveness for strategic planning.
- We show that SmartBook outperforms a web search + large language model based generation approach (akin to Bing + ChatGPT) by up to 9% in terms of relevance and strategicness of information presented.
- Finally, we contribute a publicly accessible UI¹ and code² for SmartBook so that others can explore and improve upon.

¹<https://blenderdemo.com/smartbook>

²<https://github.com/blender-nlp/SmartBook>

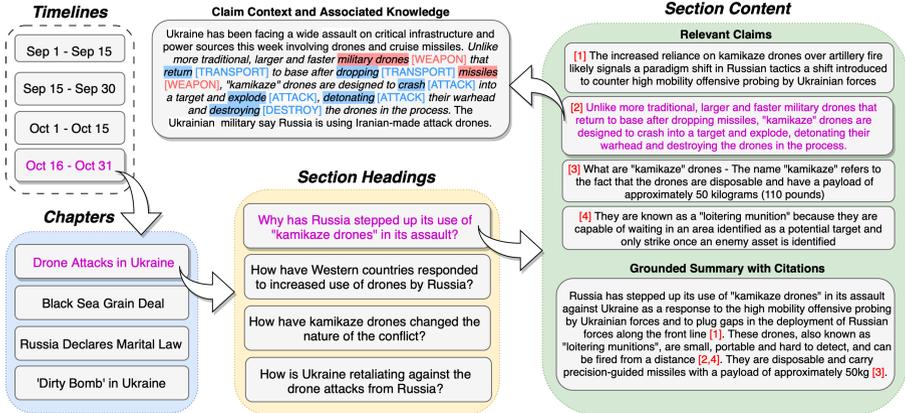


Fig. 1: Figure showing an example from SmartBook for the Ukraine-Russia Crisis. SmartBook is organized by timelines, with each timeline containing chapters and corresponding sections. The section headings are strategic questions with the section content being grounded summaries that are linked to relevant claims, each highlighted with factual evidence and knowledge elements. The entities and events (along with the corresponding types), which comprise the associated knowledge elements, are highlighted in red and blue respectively.

2 Related Work

Intelligence is more than data aggregation and hypothesis formation. Automatic tools assist intelligence analysts in data collection [5–7], data inspection and visualization [8], evidence credibility assessment [9, 10], scenario-based question answering [11], hypotheses generation [12, 13], and quality evaluation [14]. However, most current tools still require expensive crowdsourcing [15, 16] or human-machine collaboration [17–19] in report generation. To save human labor, SmartBook aims to provide analysts with an initial draft of the situation report so that they can focus on thinking critically about information and making in-depth assessments. SmartBook contains automatically generated strategic questions for different events, and helps topic discovery from large-scale data by addressing the conundrum of “you don’t know what you don’t know”. Analysts can leverage Smartbook to discover new patterns and explore alternative hypotheses with potentially conflicting information.

Summarizing multiple news articles into an situation analysis report falls into the scope of multi-document summarization [2, 20–24], which aims to extract critical information across multiple documents with particular guidance, such as sentiments [24], queries [23], and document structures [25], into a few-sentence summary [2]. In comparison, the guidance in our query-based situation report summarization is the identified and organized event knowledge in real-world scenarios (e.g., COVID and Russia-Ukraine Crisis), and the summary is a comprehensive detailed report supporting deep-dive

actions. Our work is also closely related to knowledge-enhanced summarization, which extracts useful information with prior knowledge, such as sentence similarity [26], semantic role [27–30], Yago knowledge base [31], and dependency parsing [32]. Nevertheless, the above methods fail to generate situation reports because their outputs are typically short and unable to organize events described in multiple documents, as opposed to our event-organized detailed report generation.

3 Formulation of Situation Report Generation

A situation report is a document that an intelligence analyst produces as part of their work in collecting, analyzing, and disseminating critical information. The report is typically focused on a specific issue or topic and provides a comprehensive overview of the information available on that topic. The purpose of a situation report is to provide an accurate and objective assessment of a situation or issue, based on the analysis of available data and information. The ultimate goal of such reports is to help policy-makers, military officials, and other experts keep track of developments and provide them with the necessary information to make informed decisions and take appropriate actions.

With this in mind, we aim to generate a situation report, given a collection of documents from a variety of sources, that embody the below characteristics:

- **Information salience:** The report should be able to automatically identify and extract the most relevant and crucial information across multiple documents. Such information may include key events, trends, statistics, or developments related to the subject of interest.
- **Having a logical structure:** A situation report needs to have a clear logical structure to help the reader understand, follow, and easily access information. Such logical structure can be imposed in the form of chapters, with individual section headings within each chapter being descriptive enough to give a high-level idea of the main points of the chapter.
- **Organized as timelines:** Situation reports cover event progressions over considerably long periods of time. Hence, it is beneficial to organize such reports in the form of timelines, which enables seamless report updates with new events. Further, timelines present events in an easy-to-follow chronological order, as well as identify potential trends or patterns that may emerge over time.
- **Grounded factual content:** To be reliable, a situation report must be grounded in verifiable sources. Grounded factual content helps to build credibility and trust in the situation report as grounding provides the ability to cross-check the presented information before the report can be used to inform policy and strategic planning.

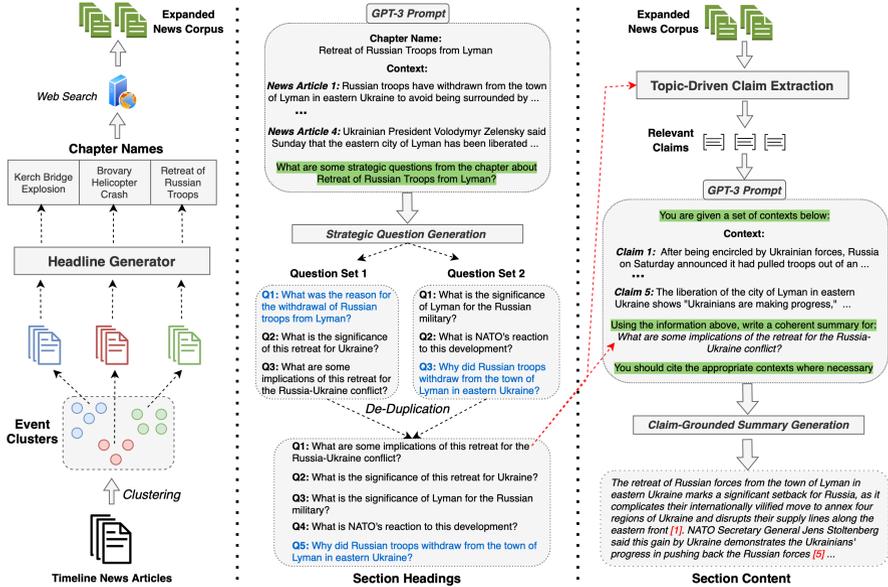


Fig. 2: Overall workflow for constructing *SmartBook*. Given the articles corresponding to a specific timeline, the figure shows the process for obtaining the chapters, their section headings, and the corresponding section content.

For our formulation of the automatic generation of situation reports, we narrow the scope with regard to the following aspects:

- **Restriction to news:** While situation reports can be created from a variety of data sources such as news articles, open-source materials, and classified information, we limit the knowledge sources to be primarily news articles. This restriction is motivated by the widespread availability of news articles and the verifiability of news publishers.
- **Withholding conclusions and recommendations:** A typical situation report is expected to contain the analyst’s conclusions and suggestions, in addition to providing insights and implications of various events. Instead, our focus for automatic report generation is on providing strategic insightful information, and grounding it to reputable sources, while leaving the subjective part of drawing conclusions and making recommendations to the human experts.

4 Method

In this section, we describe our framework for the automatic creation of situation reports given a corpus of news articles. Our automatic situation report is organized in the form of timelines to provide a coherent, chronological representation of event developments, facilitating users tracking and understanding

of situation context. Each timeline spans a duration of 2 weeks, allowing for manageable capturing and focused analysis of significant newsworthy occurrences. Building upon this, major events identified within the 2-week timeline duration serve as the foundation for corresponding chapters (section 4.1). To further guide detailed chapter analysis, we incorporate a logical structure through automatically generating section headings in the form of strategic questions relating to each major event (section 4.2). These individual sections comprise the core content of our situation report and contain grounded, query-focused summaries which address strategic questions (section 4.3), providing readers with a comprehensive understanding of event context and implications. Figure 2 gives an overall workflow diagram for our approach to constructing *SmartBook*, with each of the steps detailed below.

4.1 Major Events in Timelines as Chapters

For each timeline, we first identify major events which are in the form of news clusters within the timeline. To facilitate information navigation, we then introduce how to derive a high-level chapter name for each event cluster, and use it as keyphrase to retrieve additional news articles for enriched situation report generation.

Our approach to identifying major events within a timeline begins with collecting news articles³ for a 2-week duration corresponding to the timeline. As the number of major events is unknown apriori, we cluster the daily news summaries into major event groups utilizing the agglomerative hierarchical clustering algorithm [33], based on their term frequency-inverse document frequency (TF-IDF) scores, which assign higher weights to rare or document-specific words [34]. Finally, we are left with clusters of news snippets, each providing a focused view to a major event. But because news summary snippets are condensed in detail, we want to enrich the comprehensiveness of major event cluster with chapter name and news corpus expansion, as detailed next.

A key to further enriching event cluster representation is to derive a corresponding short chapter name, to facilitate information readability and retrieval, for the timeline within the situation report. To achieve this, we utilize a sequence-to-sequence transformer-based [35] language model, BART [36], that has been trained on the NewsHead dataset [37] for multi-document news headline generation. Each cluster within NewsHead contains up to five news articles and a crowd-sourced headline of up to 35 characters which describes the major information covered by the story in the cluster. Specifically, we take the concatenated title and text from all the news snippets within the event cluster as input sequences into the BART language model to generate a short event heading. With the generated chapter name, we query google news to obtain an expanded set of news articles relevant to the event. In particular, we get news articles corresponding to the given timeline by formatting the query as follows,

³Specifically, we use a running list of daily news snippets provided by CNN. For example <https://edition.cnn.com/europe/live-news/russia-ukraine-war-news-12-21-22/>

query = <Chapter name> after:<Timeline start date> before:<Timeline end date>, with start and end dates in the *yyyy:mm:dd* format.

4.2 Strategic Questions as Section Headings

Beyond simply describing event details in each major event chapter, Smart-Book aims to provide information from a strategic perspective that can help aid decision-making and policy planning. Thus, in this work, we generate chapter content with the use of questions that cover various strategic aspects of the event, which include the possible motivations of the actors in an event and the future implications of the event. These strategic questions are organized in the form of section headings to incorporate structure into the chapter. As we detail next, the event-related questions are generated by prompting GPT-3 [38] with a grounded context in the form of news articles from the event cluster. The generated questions then undergo a post-processing step of de-duplication to ensure clear and unique section headings.

Recent work [39, 40] has shown that GPT-3 is capable of generating natural questions that require long-form and informative answers, in comparison to existing approaches [41, 42] that mainly generate questions designed for short and specific answers. In this work, we prompt GPT-3 to explicitly generate strategic questions about the event. Further, to mitigate hallucinations [43, 44] which are common in large language models, we ground the input context to GPT-3 with news articles from the event cluster (from section ??) to ensure the generated questions are relevant to the event. Figure 2 contains an example with the prompt used to generate strategic questions with GPT-3 for a chapter named *Retreat of Russian Troops from Lyman*. News articles within the context are shown, with the prompt **highlighted** in bold at the end. To ensure diversity in the generated questions, we sample multiple question sets from GPT-3 using nucleus sampling [45].

Within the generated question sets, we observe that questions may sometimes be repeated across different sets. Figure 2 contains an example, with the duplicate questions within the generated question sets marked in blue. The problem is similar to that of detecting duplicate question pairs given a large collection of questions. For this, we leverage a publicly available⁴ RoBERTa-large [46] model trained on the Quora Duplicate Question Pairs⁵ dataset. Given a pair of questions, the model predicts a score between 0 and 1 on how likely the two given questions are duplicate. We follow this approach for filtering out duplicate questions in order to merge the multiple question sets generated by GPT-3 into a single set of diverse and unique strategic questions about the event.

⁴<https://huggingface.co/cross-encoder/quora-roberta-large>

⁵<https://quoradata.quora.com/First-Quora-Dataset-Release-Question-Pairs>

4.3 Grounded Summaries as Section Content

With the strategic questions obtained for each chapter as section headings, we can now incorporate query-focused summarization to generate section content, by treating the strategic question as query and the event-specific news articles as background corpus. Our approach employs an extract-then-summarize framework, where we first identify relevant context for the strategic question corresponding to the section. Specifically, we leverage a question-driven claim extraction approach (section 4.3.1) to identify relevant claims within news articles. Finally, to obtain the section content, we summarize the relevant claims and their corresponding contexts using a novel prompting mechanism, which also generates citations linking each summary fragment to its source (section 4.3.2). This citation linking *not only* helps ground the summary fragment generation to appropriate factual input context *but also* ensures that experts can cross-check and verify information as needed, enhancing the reliability of SmartBook situation report for decision-making assistance.

4.3.1 Question-Driven Claim Extraction with Validation

Traditionally, a situation report requires foraging for different claims and hypotheses from the source documents (i.e., news articles) that help explain a situation [19], which is expensive to obtain through manual crowd-sourcing. However, recent work [47, 48] has shown that directed queries, such as strategic questions in our case, can be used to automatically extract claims from news articles relevant to a particular topic. Following [48], we adopt a Question Answering (QA) formulation to identify claims relevant to a given strategic question. Specifically, we design a QA pipeline, utilizing a transformer-based RoBERTa-large encoder model [46] variant⁶ that has been trained on SQuAD 2.0 [49] and Natural Questions [50]. The pipeline takes as input the news corpus split into snippets along with the strategic question, and outputs short answer extractions to these questions. The identified short answers are then expanded, by including the 3-sentence window around it to provide additional context.

However, there is still a risk of false positives [51, 52] being identified as candidate answers with high confidence, thus necessitating the validation [53, 54] of extracted answers. To this end, we employ an answer sentence selection model [55] that validates each of the extracted contexts (from Section 4.3.1) separately against the strategic question. We concatenate the question and extracted context as input to a binary classification model⁷, with an underlying RoBERTa-large backbone, that is trained on Natural Questions [50] and WikiQA [56]. The output of the model is a validation score, between 0 (incorrect answer selection) and 1 (correct answer selection), used to select the top-5 relevant contexts for summarization.

⁶We use the [question-answering](#) pipeline provided by huggingface.

⁷Model is available at [TANDA github](#).

What is NATO's position on the Ukraine-Russia conflict?

Summary

NATO has not imposed a no-fly zone in the Ukraine-Russia conflict because neither Ukraine nor Russia are members of NATO. NATO's focus has been on the "tactical approach" to the war in Ukraine, responding to the developments on the ground. However, the foreign minister of Lithuania, a NATO member state, has said that NATO needs to think strategically about the longer-term and be prepared for the possibility that Russia "might war again, and not excluding NATO countries".

Claims

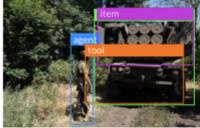
Claim Sentence	Context
<p>The US and NATO oppose creating a no-fly zone in Ukraine; such a move, they have warned, could lead to "full-fledged war in Europe." Source</p>	<p>His message was delivered amid ongoing conversations between the US and European allies about the possibility of providing Ukraine with fighter jet aircrafts from Eastern European countries, five sources familiar with the discussions told CNN. Officials say there is a divide among countries about whether or not individual nations should provide Ukraine aircraft, given the associated risks. The US and NATO oppose creating a no-fly zone in Ukraine; such a move, they have warned, could lead to "full-fledged war in Europe." See More</p> 

Fig. 3: Figure showing an example for how multimodal information (in the form of images) supports and provides additional context to the claims presented in SmartBook. In this example, the presence of anti-aircraft weapons (as seen in the image) in Ukraine provides background for the discussion in NATO on whether to impose a no-fly zone.

Situation reports often rely on a variety of sources, including text, images, videos, and audio recordings, to provide a holistic view of events. Different modalities contain complementary information, which can help gain a more comprehensive understanding or reveal novel insights. Figure 3 shows an example of how images provide additional evidence to the claims in SmartBook. Hence, for each relevant claim, we also link the associated multimedia knowledge elements by running our publicly available multimedia knowledge extraction system GAIA [57].

4.3.2 Grounded Summary Generation with LLMs

Given the set of relevant claim contexts for a strategic question, we aim to generate a concise summary as **section content** within SmartBook chapters. Observing that humans overwhelmingly prefer summaries from prompt-based large language models (LLMs) [38, 58] over models fine-tuned on article-summary pairs [36, 59, 60] due to better controllability and easier extension to novel scenarios [61–63], we leverage the Instruct-tuned 175B GPT-3 model (text-davinci-003) [38, 64] for zero-shot summarization. Specifically, we concatenate the top-5 most relevant contexts (from Section 4.3.1) as input, along with an instruction along the lines of “*summarize the above, regarding [strategic question X], with citations*”, to form the input prompt for the summary generation. It is noteworthy that the use of citations or links from the input text context in the language model prompting is novel and impactful, as it ensures the generated summary is accurate and trustworthy. Moreover, this approach also enables users to easily verify the summary information against sources used in the input. Ultimately, the intuitive and structured incorporation of

timeline chunking, major event chapter clustering, and query-focused section summaries within chapters enables our proposed SmartBook formulation to generate insightful reports for time-sensitive, emerging situations.

5 Experiments & Results

In this section, we perform quantitative and qualitative evaluations on the different generative content within SmartBook, such as the strategic questions (*section headings*) and the corresponding summaries (*section content*). Specifically, we evaluate both how informative the content is as well as how much it can benefit analysts. For strategic questions, the evaluation is based on how useful these might be to analysts (Section 5.1). The summaries for these questions, which corresponds to the section content, are evaluated based on how much relevant information is contained within them (Section 5.2).

For our experiments, the news articles that are used to identify the major events are collected daily from the CNN website⁸ related to the Ukraine-Russia Crisis. Using this data source, we organized SmartBook into chapters according to major events within a biweekly time span. The overall time period considered ranges from Sept 1, 2022, to Jan 31, 2023.

5.1 Strategic Importance of Questions

We first evaluate the quality of the strategic questions, which are used as section headings in our SmartBook. The evaluation is done by drawing parallels to the construction of a situation report, which involves the participation of both senior and junior analysts. A senior analyst’s role is to come up with strategic questions which are then passed onto a junior analyst. The junior analyst now gathers tactical information that can help answer or provide more background to the strategic questions.

Specifically, our human evaluation of the questions within SmartBook measures the following aspects:

- **Strategic Importance:** This evaluation is from the perspective of a senior analyst. Evaluating the strategic importance of a question requires introspection on whether the question provided within SmartBook would have been asked as a strategic question by a senior analyst. The scoring rubric involves marking the question as one of three categories, namely, *Not Strategic*, *Some Strategic Value*, and *Definitely Strategic*.
- **Relevant Tactical Information:** Another dimension of evaluating the quality of a question is by measuring how much relevant tactical information useful to a junior analyst can be gathered using that question. In this case, the question is evaluated based on its corresponding SmartBook summary, in terms of how much of the question-relevant tactical information in the summary is helpful for an intelligence analyst to gain deeper insights into the

⁸Here is a URL corresponding to news for March 1: [link](#)

<p>Question: What were the motives behind the attack on Russian missile in Dnipro?</p>	<p>Question: What is the capability of Ukraine's air defense systems in repelling the missile attacks from Russian missiles</p>
<p>Summary: The Russian missile attack on the city of Dnipro on Saturday was part of the Kremlin's efforts to limit Ukraine's ability to heat and power itself in the middle of winter. The attack was reportedly carried out with a Kh-22 missile, an older and less accurate weapon than most modern missiles. It is believed that the attack was part of the Kremlin's attempts to retake a town, Soledar, and the nearby city of Bakhmut, as a way to grind down the best Ukrainian forces and prevent them from launching counterattacks elsewhere. Nine people were killed, including a 15-year old girl, Kyrylo Tymoshenko, and Russia denied responsibility for the attack, claiming it was caused by Ukrainian air defences.</p>	<p>Summary: Ukraine's air defense systems have demonstrated their capability in repelling the missile attacks from Russian forces in Kyiv. [1] reported that Ukrainian air defense forces shot down 37 of "about 40" rockets that entered the city's airspace, with one person injured. [2] revealed that a commander, who asked to be identified only by the call sign "Hera" for security reasons, stated that a cruise missile was shot down by a machine gun from the city's territorial defense mobile group. [3] mentioned that U.S. President Joe Biden had pledged to deliver one Patriot surface-to-air missile battery system to Ukraine, which is one of the most advanced U.S. air defense systems and is capable of intercepting threats such as aircraft and ballistic missiles.</p>

Table 2: Table showing summaries for two strategic questions corresponding to a SmartBook chapter on Russian missile attacks. The tactically useful and relevant information has been highlighted in blue. Tactically useful but irrelevant information has been highlighted in red.

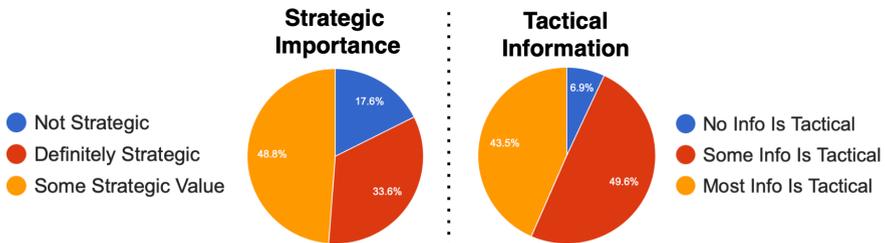


Fig. 4: Results from the human evaluation of the questions (section headings) in SmartBook. Judgments are based on the strategic importance of the question and tactical information that can be gathered using the question.

situation. The scoring rubric is based on the summary, categorized into *No information is tactical*, *Some information is tactical*, and *Most information is tactical*. Table 2 shows some examples with tactical information relevant to the question highlighted in blue.

For the evaluation, we randomly selected 25 chapters from SmartBook, which amounted to 125 strategic questions in total. Each example was annotated by 3 human evaluators. Figure 4 shows the results. We observe that most questions are strategic, with at least 82% of the questions having some strategic value. Further, we can see that these questions can help gather relevant tactical information in roughly 93% of the cases, as judged by the annotators that most summaries to these questions have information that is tactically relevant.

<i>Summary-Type</i>	<i>Background</i>	<i>Coherence</i>	<i>Relevance</i>	<i>Strategicness</i>
(1) Query-Focused	News	4.39	3.85	3.77
(2) Web Search + LLM	Web	4.44	4.07	3.83
(3) SmartBook	News	4.56	4.23	4.17

Table 3: Results of human evaluation comparing the summaries within SmartBook (3) against those either generated directly from news articles via query-focused summarization (1) or obtained from the internet by web search combined with generation from a large language model (2). Results are on a 1-5 scale (worst to best).

5.2 Quality of Generated Section Content

We also evaluate the quality of summaries generated for the strategic questions. Summaries within SmartBook are generated by using question-driven claim extraction and validation (section 4.3.1) before generating the final summary. For comparison, we introduce a *query-focused summarization* baseline that directly uses the news articles as input and uses a large language model (LLM) to generate a query-focused summary⁹. Further, SmartBook uses news articles (relevant to the event) as the background corpus. We compare this against using the web as the background corpora, by obtaining relevant web pages from internet search and summarizing them using a large language model. This baseline simulates a competitive strategy of directly obtaining information from an LLM-enabled web search engine¹⁰ (the *web search + LLM* baseline is similar to Bing + ChatGPT).

We randomly selected 50 strategic question sections to perform the human evaluation. Each data sample consists of a question and the generated summaries from three approaches: Smartbook, Web Search + LLM, and Query-focused summarization. Three annotators are asked to access the anonymous summaries in terms of coherence, relevance, and strategicness with the scores ranging from 1 (worst) to 5 (best). For coherence and relevance, we follow the guidelines as in [65]. Simply, *coherence* measures the quality of all sentences collectively regardless of the query and *relevance* quantifies whether the summary contains key points that answer the given question. We define *strategicness* as an indication of whether the summary provides non-trivial and insightful information for analysts, and suggests the breadth and depth of the provided key points.

Table 3 shows the results from the human evaluation of the summaries. We observe that SmartBook outperforms alternative competitive strategies along with coherence, relevance, and strategicness. The benefit of our question-driven claim extraction step (3 vs 1) can be seen from the considerably more relevant summaries within SmartBook, compared to the direct query-focused

⁹We use the same Instruct-GPT davinci-003 model as in Section 4.3.2. The prompt used is the same as in the summary generation example in Figure 2, except the entire news article texts are passed as context.

¹⁰Specifically, we leverage perplexity.ai, which combines web-search with LLM summarization. The query used is the strategic question, along with a phrase “with respect to the Ukraine-Russia crisis between <timeline>” to ground the query.

Question: How have kamikaze drones changed the nature of the Russia-Ukraine conflict?

Query-Focused	Web Search + LLM	SmartBook
<p>Kamikaze drones, or suicide drones, have been used by Russia in its attacks on Ukrainian cities such as Kyiv, Vinnytsia and Odesa. They are known as a “loitering munition” because they are capable of waiting in an area identified as a potential target and only strike once an enemy asset is identified. Russia is using Iranian-made Shahed-136 drones, which have a payload of approximately 50 kilograms (110 pounds) and are hard to detect due to their small size. Ukraine has responded by developing an app that civilians can use to report sightings of incoming drones and missiles, and increasing the proportion shot down before they hit the ground.</p>	<p>Russia and Ukraine have been using Iranian-made Shahed-136 drones in the conflict. These drones are known as ‘kamikaze’ drones due to their ability to self-destruct on impact. The use of these drones has been highlighted by Russia’s growing reliance on them and their use as both military weapons and psychological weapons. On October 13, 2022, kamikaze drones and missiles hit north and south Ukraine, though the Kremlin denied any involvement from President Putin in talks between Russia and the West to end the war in Ukraine.</p>	<p>The Russia-Ukraine conflict has seen a paradigm shift in the nature of Russian tactics, with the increased reliance of kamikaze drones over artillery fire. These drones, which are designed to crash into a target and explode, can spray potentially deadly debris and are being used to plug gaps in the deployment of Russian forces along the front line. Russia is reportedly using Iranian-made Shahed series drones, which are showcased at Khashan Airfield south of Tehran. This shift in tactics has seen an escalation in the aerial bombardment of densely populated urban areas and an attack on critical infrastructure and power sources.</p>

Table 4: An example showing the summaries produced by different approaches evaluated in Table 3. Question-relevant information in the summaries is highlighted in blue. It can be seen that the SmartBook summary contains considerably more relevant information. The *Web Search + LLM* baseline output also contains hallucinated content (highlighted in red) that suggests *both* Russia and Ukraine have been using the ‘kamikaze’ drones.

summarization of the news articles without an explicit extraction step. Further, we see that directly obtaining information from the internet (2 vs 3) can give strategically less important content compared to the focused news-driven summarization within SmartBook. A qualitative example comparing the three approaches is shown in Table 4.

5.3 Editing by Expert Analyst

SmartBook aims to assist analysts in situation report generation, by serving as an initial draft that contains potentially strategic questions and summaries with tactical content. To quantitatively assess how beneficial the automatically generated chapter summaries from our proposed SmartBook framework can be, we asked an expert intelligence analyst to edit the summaries until they deemed the summary acceptable to them.

We aim to measure the extent to which the expert analyst makes changes to the summaries within SmartBook. Specifically, such edits can be quantified in the form of (a) the percentage of tokens inserted and deleted, and (b) Levenshtein edit distance [66], which measures character level changes required to transform one string (old summary) to the other (new summary). Further, we also evaluate the commonly used token-overlap-based metrics such

SmartBook Summary: Russia has reportedly stepped up its use of kamikaze drones in its assault against Ukraine. The increased reliance on kamikaze drones over artillery fire likely signals a paradigm shift in Russian tactics a shift introduced to counter high mobility offensive probing by Ukrainian forces.

Analyst-edited Summary: Russia has reportedly stepped up its use of kamikaze drones in its assault against Ukraine. **The aircraft are called kamikaze drones because they attack once and don't come back.** The increased reliance on kamikaze drones over artillery fire likely signals a paradigm shift in Russian tactics - a shift introduced to counter high mobility offensive probing by Ukrainian forces. **Their low price means the drones can be deployed in large numbers and they hover before they strike, so they have a psychological effect on civilians as they watch and wait for them to strike. These drones allow Russia to target Ukrainians far away from the front line, away from the primary battle space. The emergence of swarms of drones in Ukraine is part of a shift in the nature of the Russian offensive, which some speculate indicates that Moscow may be running low on long-range missiles.**

Table 5: Example showing intelligence analyst edits for an automatic machine generated SmartBook summary on the use of “kamikaze” drones in the Ukraine-Russia crisis. Text that has been added by the expert is colored, with blue corresponding to additional tactical information, whereas red corresponds to insights/conclusions added by the analyst.

as BLEU [67] and ROUGE [68] scores, by considering the edited summary as the reference. The expert editing was performed over 94 randomly sampled summaries.

Empirical results indicate that the token overlap between the generated and post-edited summaries is high, with a BLEU score of 59.0% and a Rouge-L score of 74.1%. This suggests that the generated summaries are sufficiently good-quality such that not much human expert revision is necessary. We acknowledge that a gap still exists between the automatically generated summaries and human expert summary curation, as the Levenshtein edit distance computed at the character level is 34.4%. Interestingly, 15% of the generated summaries are determined as perfect, with no edits made by the expert analyst. The general observation is that more content tends to be inserted rather than deleted by the expert analyst. The percentage of tokens inserted is 49.6%, whereas the percentage of tokens deleted is only 2.3%, pointing to the importance of summary generation to be more detailed.

Table 5 shows an example of edits (shown in color) made by an intelligence analyst for a machine-generated summary in SmartBook. We can see that human analyst mainly tends to add additional tactical information (in blue) in order to elaborate on certain aspects (e.g. what is special about the “kamikaze” type of drone). Further, the analyst can also draw some interesting insights (in red) based on the information in the summary. Overall, this shows that SmartBook can act as a good starting point for analysts to expand upon for the generation of situation reports.

5.4 Error Categorization for Summaries

In Section 5.3, we saw that 15% of the summaries had no errors and were approved by an expert analyst as needing no edits. To further understand the different types of errors in the remaining summaries, we asked the analyst to

categorize the errors within them. The analyst was also shown the strategic question and the corresponding extracted contexts that were used to generate the summary, with the summary error categories being as follows:

- *No relevant contexts*: None of the extracted contexts are relevant to the question (and thereby the summary is expected to be fully irrelevant too).
- *Inaccurate information in summary*: Summary has incorrect information, that is not reflective of the underlying input contexts.
- *Incoherent summary*: Summary is incomprehensible and unclear.
- *Incomplete summary*: Important information in the input contexts is missing in the summary.
- *Irrelevant information in summary*: Summary has material that is not relevant to the question, despite some extracted contexts being relevant.

Figure 5 shows the distribution of error categories for the summaries. It can be seen that incompleteness of summaries is a predominant error, with more than 50% of the summaries missing important information or not being sufficiently complete. Also, very few summaries are incoherent, which is expected, since large language models such as GPT-3 have been shown [61, 69] to generate fluent and easy to read output. However, the hallucination problem [43, 70] of such LLMs is evident from the considerable number of summaries with inaccurate information.

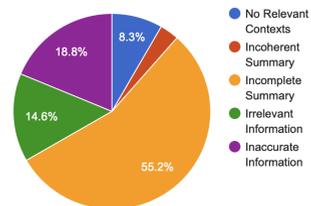


Fig. 5: Distribution of different error categories for the summaries within SmartBook.

6 Conclusion and Future Extensions

In this paper, we present the novel task of automated situation report generation, which is pertinent for timely and scalable complex event understanding of emergent crises. This problem formulation differs from the traditional multi-document news summarization domain in that it requires comprehensive, broad coverage of important information, which is to be presented in a structured format that is spread across a timeline. In addition, the vast amount of information needs to be organized into chapters and sections with a strategic structure to facilitate analysis and planning. To address these challenges, we propose *SmartBook*, a generalizable framework for clustering news topics and summarizing news claims automatically. It is important to note that while *SmartBook* offers automated assistance in generating situation reports, it is not meant to replace human analysts, but rather to assist them in the process.

For future work, we aim to explore the following avenues to improve the reliability and expand the scope of *SmartBook*:

- **Controlling the bias of sources:** Different news sources present information from varying angles, with different levels of detail and interpretation. Thereby, we plan to control for the bias of news sources used in the creation

of SmartBook to help cross-check information and avoid the potential pitfalls of relying on a single news source or a single interpretation of events. Further, alternate perspectives/hypotheses can be identified by independently considering news sources on the left, in the center, and on the right.

- **Adding more languages:** The world is diverse, and news articles from different languages allow capturing different perspectives, insights, and information relating to global events, that might not be available in a single language. For this reason, we aim to add news articles from more languages into SmartBook to present a more comprehensive and accurate picture. Further, the ability to ingest information from different languages also enables analysts to understand local customs, cultures, and nuances that may influence the interpretation of events.
- **Verifying the extracted claims:** Given that situation reports contribute towards strategic planning, it is important to present accurate information to prevent misguided decisions and actions. Hence, we intend to provide a verification score, for each claim presented within SmartBook, which can be interpreted as a measure of trustworthiness. Claims with lower scores signal the need for additional oversight by the intelligence analysts before incorporating into them the final report.

Acknowledgments

We are grateful to Dr. Clare Voss from the Army Research Laboratory for her helpful insights and feedback. We would also like to acknowledge Lisa Ferro and Brad Goodman from MITRE for their valuable comments and help with expert evaluation. This research is based upon work supported by U.S. DARPA AIDA Program No. FA8750-18-2-0014, DARPA KAIROS Program No. FA8750-19-2-1004, DARPA SemaFor Program No. HR001120C0123, DARPA INCAS Program No. HR001121C0165 and DARPA MIPS Program No. HR00112290105. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of DARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [1] OpenAI: GPT-3.5. Computer software. <https://openai.com/research/>
- [2] Fabbri, A., Li, I., She, T., Li, S., Radev, D.: 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, pp. 1074–1084. Association for Computational Linguistics, Florence, Italy (2019). <https://doi.org/10.18653/v1/P19-1102>. <https://aclanthology.org/P19-1102>
- [3] Bowen, A.: Russia’s war in ukraine: Military and intelligence aspects. Congressional Research Service **47068** (2022)
- [4] Brody, S., Kantor, P.: Automatic assessment of coverage quality in intelligence reports. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, pp. 491–495. Association for Computational Linguistics, Portland, Oregon, USA (2011). <https://aclanthology.org/P11-2086>
- [5] Stottlemire, S.A.: Humint, osint, or something new? defining crowd-sourced intelligence. *International Journal of Intelligence and CounterIntelligence* **28**(3), 578–589 (2015)
- [6] Billman, D., Convertino, G., Shrager, J., Pirolli, P., Massar, J.: Collaborative intelligence analysis with cache and its effects on information gathering and cognitive bias. In: Human Computer Interaction Consortium Workshop, vol. 35 (2006). Citeseer
- [7] Pohl, D., Bouchachia, A., Hellwagner, H.: Automatic sub-event detection in emergency management using social media. In: Proceedings of the 21st International Conference on World Wide Web, pp. 683–686 (2012)
- [8] Prunckun, H.: Handbook of scientific methods of inquiry for intelligence analysis **11** (2010)
- [9] Toniolo, A., Cerutti, F., Oren, N., Norman, T.J., Sycara, K.: Making informed decisions with provenance and argumentation schemes. In: Proceedings of the Eleventh International Workshop on Argumentation in Multi-Agent Systems, vol. 2014 (2014)
- [10] Barnes, A.: Making intelligence analysis more intelligent: Using numeric probabilities. *Intelligence and National Security* **31**(3), 327–344 (2016)
- [11] Small, S., Strzalkowski, T., Liu, T., Ryan, S., Salkin, R., Shimizu, N., Kantor, P., Kelly, D., Rittman, R., Wacholder, N., Yamrom, B.: HITIQA: Scenario based question answering. In: Proceedings of the Workshop on Pragmatics of Question Answering at HLT-NAACL 2004, pp. 52–59. Association for Computational Linguistics, Boston, Massachusetts, USA (2004). <https://aclanthology.org/W04-2507>
- [12] Heuer, R.J.: Psychology of intelligence analysis (1999)
- [13] Cerutti, F., Norman, T.J., Toniolo, A., Middleton, S.E.: Cispaces. org:

- From fact extraction to report generation. *COMMA* **2018**, 269–280 (2018)
- [14] Brody, S., Kantor, P.: Automatic assessment of coverage quality in intelligence reports. In: *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 491–495 (2011)
- [15] van Gelder, T., Kruger, A., Thomman, S., de Rozario, R., Silver, E., Saletta, M., Barnett, A., Sinnott, R.O., Jayaputera, G.T., Burgman, M.: Improving analytic reasoning via crowdsourcing and structured analytic techniques. *Journal of Cognitive Engineering and Decision Making* **14**(3), 195–217 (2020)
- [16] Xia, H., Osterlund, C., McKernan, B., Folkestad, J., Rossini, P., Boichak, O., Robinson, J., Kenski, K., Myers, R., Clegg, B., et al.: Trace: A stigmergic crowdsourcing platform for intelligence analysis (2019)
- [17] Pirolli, P., Card, S.: The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In: *Proceedings of International Conference on Intelligence Analysis*, vol. 5, pp. 2–4 (2005). McLean, VA, USA
- [18] Pioch, N.J., Everett, J.O.: Polestar: collaborative knowledge management and sensemaking tools for intelligence analysts. In: *Proceedings of the 15th ACM International Conference on Information and Knowledge Management*, pp. 513–521 (2006)
- [19] Toniolo, A., Cerutti, F., Norman, T.J., Oren, N., Allen, J.A., Srivastava, M., Sullivan, P.: Human-machine collaboration in intelligence analysis: An expert evaluation. *Intelligent Systems with Applications* **17**, 200151 (2023)
- [20] Abdi, A., Idris, N., Alguliyev, R.M., Aliguliyev, R.M.: Query-based multi-documents summarization using linguistic knowledge and content word expansion. *Soft Computing* **21**(7), 1785–1801 (2017)
- [21] Li, C., Liu, Y., Liu, F., Zhao, L., Weng, F.: Improving multi-documents summarization by sentence compression based on expanded constituent parse trees. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 691–701 (2014)
- [22] Xiong, C., Li, X., Li, Y., Liu, G.: Multi-documents summarization based on textrank and its application in online argumentation platform. *International Journal of Data Warehousing and Mining (IJDWM)* **14**(3), 69–89 (2018)

- [23] Abdi, A., Shamsuddin, S.M., Aliguliyev, R.M.: Qmos: Query-based multi-documents opinion-oriented summarization. *Information Processing & Management* **54**(2), 318–338 (2018)
- [24] Abdi, A., Shamsuddin, S.M., Hasan, S., Piran, J.: Machine learning-based multi-documents sentiment-oriented summarization using linguistic treatment. *Expert Systems with Applications* **109**, 66–85 (2018)
- [25] DeYoung, J., Beltagy, I., van Zuylen, M., Kuehl, B., Wang, L.L.: Ms2: Multi-document summarization of medical studies. arXiv preprint arXiv:2104.06486 (2021)
- [26] Wan, X.: An exploration of document impact on graph-based multi-document summarization. In: *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pp. 755–762 (2008)
- [27] Yan, S., Wan, X.: Srrank: leveraging semantic roles for extractive multi-document summarization. *IEEE/ACM Transactions on audio, speech, and language processing* **22**(12), 2048–2058 (2014)
- [28] Khan, A., Salim, N., Kumar, Y.J.: A framework for multi-document abstractive summarization based on semantic role labelling. *Applied Soft Computing* **30**, 737–747 (2015)
- [29] Khan, A., Salim, N., Kumar, Y.J.: Genetic semantic graph approach for multi-document abstractive summarization. In: *2015 Fifth International Conference on Digital Information Processing and Communications (ICDIPC)*, pp. 173–181 (2015). IEEE
- [30] Alshaina, S., John, A., Nath, A.G.: Multi-document abstractive summarization based on predicate argument structure. In: *2017 IEEE International Conference on Signal Processing, Informatics, Communication and Energy Systems (SPICES)*, pp. 1–6 (2017). IEEE
- [31] Baralis, E., Cagliero, L., Jabeen, S., Fiori, A., Shah, S.: Multi-document summarization based on the yago ontology. *Expert Systems with Applications* **40**(17), 6976–6984 (2013)
- [32] Kurisinkel, L.J., Zhang, Y., Varma, V.: Abstractive multi-document summarization by partial tree extraction, recombination and linearization. In: *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 812–821 (2017)
- [33] Jain, A.K., Dubes, R.C.: *Algorithms for clustering data* (1988)
- [34] Sparck Jones, K.: A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation* **28**(1), 11–21 (1972)

- [35] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
- [36] Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., Zettlemoyer, L.: Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7871–7880 (2020)
- [37] Gu, X., Mao, Y., Han, J., Liu, J., Yu, H., Wu, Y., Yu, C., Finnie, D., Zhai, J., Zukoski, N.: Generating Representative Headlines for News Stories. In: *Proc. of the the Web Conf. 2020* (2020)
- [38] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., *et al.*: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020)
- [39] Sharma, S.: Generative models for question generation. Available at SSRN 3970263 (2021)
- [40] Wang, Z., Valdez, J., Basu Mallick, D., Baraniuk, R.G.: Towards human-like educational question generation with large language models. In: *International Conference on Artificial Intelligence in Education*, pp. 153–166 (2022). Springer
- [41] Murakhovs’ka, L., Wu, C.-S., Laban, P., Niu, T., Liu, W., Xiong, C.: MixQG: Neural question generation with mixed answer types. In: *Findings of the Association for Computational Linguistics: NAACL 2022*, pp. 1486–1497. Association for Computational Linguistics, Seattle, United States (2022). <https://doi.org/10.18653/v1/2022.findings-naacl.111>. <https://aclanthology.org/2022.findings-naacl.111>
- [42] Du, X., Shao, J., Cardie, C.: Learning to ask: Neural question generation for reading comprehension. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1342–1352. Association for Computational Linguistics, Vancouver, Canada (2017). <https://doi.org/10.18653/v1/P17-1123>. <https://aclanthology.org/P17-1123>
- [43] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y., Madotto, A., Fung, P.: Survey of hallucination in natural language generation. *ACM Computing Surveys*
- [44] Maynez, J., Narayan, S., Bohnet, B., McDonald, R.: On faithfulness and factuality in abstractive summarization. In: *Proceedings of the 58th*

- Annual Meeting of the Association for Computational Linguistics, pp. 1906–1919 (2020)
- [45] Holtzman, A., Buys, J., Du, L., Forbes, M., Choi, Y.: The curious case of neural text degeneration. In: International Conference on Learning Representations (2019)
- [46] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019)
- [47] Reddy, R.G., Chinthakindi, S.C., Wang, Z., Fung, Y., Conger, K., Elsayed, A., Palmer, M., Nakov, P., Hovy, E., Small, K., *et al.*: Newsclaims: A new benchmark for claim detection from news with attribute knowledge. In: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, pp. 6002–6018 (2022)
- [48] Reddy, R.G., Chinthakindi, S.C., Fung, Y.R., Small, K., Ji, H.: A zero-shot claim detection framework using question answering. In: Proceedings of the 29th International Conference on Computational Linguistics, pp. 6927–6933 (2022)
- [49] Rajpurkar, P., Jia, R., Liang, P.: Know what you don’t know: Unanswerable questions for squad. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pp. 784–789 (2018)
- [50] Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Devlin, J., Lee, K., *et al.*: Natural questions: A benchmark for question answering research. Transactions of the Association for Computational Linguistics **7**, 452–466 (2019)
- [51] Tan, C., Wei, F., Zhou, Q., Yang, N., Lv, W., Zhou, M.: I know there is no answer: Modeling answer validation for machine reading comprehension. In: CCF International Conference on Natural Language Processing and Chinese Computing, pp. 85–97 (2018). Springer
- [52] Chakravarti, R., Sil, A.: Towards confident machine reading comprehension. arXiv preprint arXiv:2101.07942 (2021)
- [53] Reddy, R.G., Sultan, M.A., Kayi, E.S., Zhang, R., Castelli, V., Sil, A.: Answer span correction in machine reading comprehension. In: Findings of the Association for Computational Linguistics: EMNLP 2020, pp. 2496–2501 (2020)
- [54] Zhang, Z., Vu, T., Moschitti, A.: Joint models for answer verification in question answering systems. In: Proceedings of the 59th Annual Meeting

- of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pp. 3252–3262 (2021)
- [55] Garg, S., Vu, T., Moschitti, A.: Tanda: Transfer and adapt pre-trained transformer models for answer sentence selection. Proceedings of the AAAI Conference on Artificial Intelligence **34**(05), 7780–7788 (2020). <https://doi.org/10.1609/aaai.v34i05.6282>
- [56] Yang, Y., Yih, W.-t., Meek, C.: WikiQA: A challenge dataset for open-domain question answering. In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, pp. 2013–2018. Association for Computational Linguistics, Lisbon, Portugal (2015). <https://doi.org/10.18653/v1/D15-1237>. <https://aclanthology.org/D15-1237>
- [57] Li, M., Zareian, A., Lin, Y., Pan, X., Whitehead, S., Chen, B., Wu, B., Ji, H., Chang, S.-F., Voss, C., *et al.*: Gaia: A fine-grained multimedia knowledge extraction system. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations, pp. 77–86 (2020)
- [58] Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H.W., Sutton, C., Gehrmann, S., *et al.*: Palm: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311 (2022)
- [59] Zhang, J., Zhao, Y., Saleh, M., Liu, P.: Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In: International Conference on Machine Learning, pp. 11328–11339 (2020). PMLR
- [60] Liu, Y., Liu, P., Radev, D., Neubig, G.: Brio: Bringing order to abstractive summarization. In: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 2890–2903 (2022)
- [61] Goyal, T., Li, J.J., Durrett, G.: News summarization and evaluation in the era of gpt-3. arXiv preprint arXiv:2209.12356 (2022)
- [62] Bhaskar, A., Fabbri, A.R., Durrett, G.: Zero-shot opinion summarization with gpt-3. arXiv preprint arXiv:2211.15914 (2022)
- [63] Reddy, R.G., Elfardy, H., Chan, H.P., Small, K., Ji, H.: Sumren: Summarizing reported speech about events in news. arXiv preprint arXiv:2212.01146 (2022)
- [64] Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C.L., Mishkin,

- P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al.: Training language models to follow instructions with human feedback. arXiv preprint arXiv:2203.02155 (2022)
- [65] Fabbri, A.R., Kryscinski, W., McCann, B., Xiong, C., Socher, R., Radev, D.: Summeval: Re-evaluating summarization evaluation. *Transactions of the Association for Computational Linguistics* **9**, 391–409 (2021)
- [66] Navarro, G.: A guided tour to approximate string matching. *ACM computing surveys (CSUR)* **33**(1), 31–88 (2001)
- [67] Papineni, K., Roukos, S., Ward, T., Zhu, W.-J.: Bleu: a method for automatic evaluation of machine translation. In: *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318 (2002)
- [68] Lin, C.-Y.: Rouge: A package for automatic evaluation of summaries. In: *Text Summarization Branches Out*, pp. 74–81 (2004)
- [69] Zhang, T., Ladhak, F., Durmus, E., Liang, P., McKeown, K., Hashimoto, T.B.: Benchmarking large language models for news summarization. arXiv preprint arXiv:2301.13848 (2023)
- [70] Tam, D., Mascarenhas, A., Zhang, S., Kwan, S., Bansal, M., Raffel, C.: Evaluating the factual consistency of large language models through summarization. arXiv preprint arXiv:2211.08412 (2022)