

Generating Uncontextualized and Contextualized Questions for Document-Level Event Argument Extraction

Md Nayem Uddin[♣] Enfa Rose George[◇] Eduardo Blanco[◇] Steven R. Corman[♣]

[♣]Arizona State University [◇]University of Arizona

[♣]{muddin11, steve.corman}@asu.edu

[◇]{enfageorge, eduardoblanco}@arizona.edu

Abstract

This paper presents multiple question generation strategies for document-level event argument extraction. These strategies do not require human involvement and result in uncontextualized questions as well as contextualized questions grounded on the event and document of interest. Experimental results show that combining uncontextualized and contextualized questions is beneficial, especially when event triggers and arguments appear in different sentences. Our approach does not have corpus-specific components, in particular, the question generation strategies transfer across corpora. We also present a qualitative analysis of the most common errors made by our best model.

1 Introduction

Event argument extraction (Dodding et al., 2004; Aguilar et al., 2014) is about identifying entities participating in events and specifying their role (e.g., the *giver*, *recipient*, and *thing given* in a *giving* event). Transforming natural language into structured event knowledge benefits many downstream tasks such as machine reading comprehension (Han et al., 2021), news summarization (Li et al., 2016), coreference resolution (Huang et al., 2019), and dialogue systems (Su et al., 2022).

Traditionally, corpora are limited to arguments within the same sentence an event belongs to. Inter-sentential arguments are more challenging and have received less attention (Gerber and Chai, 2010; Ruppenhofer et al., 2010). Figure 1 presents an example from RAMS (Ebner et al., 2020), the largest corpus annotating multi-sentence event-argument structures. Two out of four event-argument relations cross sentence boundaries.

The current landscape of argument extraction is dominated by prompt-based methods (Ma et al., 2022; Nguyen et al., 2023) and question-answering

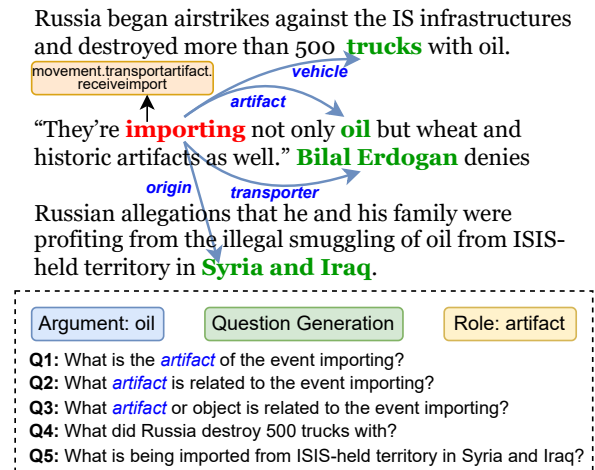


Figure 1: Example from RAMS (top, event trigger *importing* and its arguments). In this paper, we experiment with several strategies to generate questions for event argument extraction. Questions for the *artifact* of *importing* are shown in the dashed box. Q1 is generated following a role-specific template, Q2 and Q3 are generated prompting GPT-4, and Q4 and Q5 are generated by a weakly-supervised T5 model.

techniques (Du and Ji, 2022). However, these approaches heavily rely on rigid templates, neglecting the valuable context provided by the document. Another avenue for improvement involves integrating external knowledge (Zhang et al., 2023; Yang et al., 2023), but the challenge lies in the trial and error process of finding relevant external knowledge. In this paper, we address these issues by generating contextualized questions grounded on the event of interest. Our approaches capture the semantic relevance across different event arguments and non-event argument entities from the document.

As illustrated in Figure 1, for the *artifact* role of event trigger *importing*, we generate five questions and employ transformers to identify answers that pinpoint the argument (i.e., *oil*) or, alternatively, indicate that there is no answer. The first three

questions are uncontextualized (i.e., disregard the document; they only take into account the event and argument role). On the other hand, the last two questions are contextualized (i.e., grounded on the event trigger and document). Contextualized questions are generated with a T5 model trained using weak supervision. As we shall see, they yield substantial improvements even though they are noisy (e.g., the answer to Q4 is not *oil*).

The main contributions of this paper are:

- Several strategies to generate uncontextualized and contextualized questions, including weakly supervised models that require no manual effort beyond defining prompts.
- Experimental results showing that combining uncontextualized and contextualized questions is beneficial with RAMS.
- Experimental results showing that our approach is effective with out-of-domain corpora, event triggers, and roles. In particular, the question generation component transfers seamlessly to other corpora (WikiEvents).

A key benefit of our approach is that it can be easily adapted to any event-argument extraction benchmark—we do not have any benchmark-specific component. The only requirement is a list of argument roles an event may have. Most event-argument annotation efforts satisfy this requirement, including PropBank (Palmer et al., 2005), NomBank (Meyers et al., 2004), FrameNet (Baker et al., 1998), RAMS, ACE (Doddington et al., 2004), and WikiEvents (Li et al., 2021). All examples in the paper are from RAMS, but we also applied our optimal question-answering strategy to WikiEvents to demonstrate its adaptability.

2 Previous Work

Event argument extraction (Ahn, 2006) has a long history in the field (Grishman and Sundheim, 1996; Doddington et al., 2004). Initially, datasets focused on extracting arguments within the same sentence than the event (Palmer et al., 2005; Walker et al., 2006). There are also corpora focused on inter-sentential arguments (Gerber and Chai, 2010; Ruppenhofer et al., 2010; Ebner et al., 2020; Li et al., 2021). Early models were based on handcrafted features (Li et al., 2013; Liao and Grishman, 2010; Hong et al., 2011). Like most NLP tasks, models for event argument extraction experienced a transformation building on word embeddings, RNNs, and CNNs (Chen et al., 2015; Nguyen et al., 2016).

Transformer-based approaches are currently the best performing. Some efforts assume event triggers and argument spans are part of the input and present classifiers to identify the argument role (Ebner et al., 2020; Chen et al., 2020). Unlike them—and like the remaining previous works discussed below—we only assume event triggers. At a high-level, efforts to identify argument spans and their role can be categorized into sequence labeling, casting the problem as a question-answering task, and using generative models. Sequence label classifiers approach the problem with the traditional BIO encoding (Ramponi et al., 2020). Framing the problem in terms of questions and answers is popular (Du and Cardie, 2020; Liu et al., 2020; Li et al., 2020; Uddin et al., 2024). Doing so enables zero-shot (Lyu et al., 2021) and few-shot (Sainz et al., 2022) predictions. Li et al. (2021) and Ma et al. (2022); Du et al. (2021) leverage generative language models (Raffel et al., 2020; Lewis et al., 2020). Language generation facilitates a more flexible extraction by *generating* the arguments rather than identifying spans in the input document. Transfer learning has also been explored, including role overlapping knowledge (Zhang et al., 2023), semantic roles (Zhang et al., 2022), abstract meaning representations (Xu et al., 2022), and frame-aware knowledge distillation (Wei et al., 2021).

We approach the problem as a question answering task by querying a transformer model. Unlike previous work, we explore both uncontextualized as well as contextualized questions grounded on the event and document of interest. The latter questions are generated with a weakly supervised model. This allows us to train the model with one corpus (RAMS) that transfer to other corpus (WikiEvents).

3 Generating Questions for Event Argument Extraction

We transform RAMS event-argument structures into questions and answers. While gold event-argument annotations include gold answers (i.e., the argument), there are no gold questions asking for the arguments filling the roles of event triggers. Hence, we propose five question generation strategies divided into two categories: uncontextualized and contextualized. Unlike uncontextualized questions, contextualized questions are grounded on the event and document of interest.

3.1 Uncontextualized Questions

Uncontextualized questions only have access to the event trigger (e.g., *importing* in Figure 1) and role (e.g., *artifact*). They disregard the document of interest and always generate the same question for an event trigger and role.

Template-Based Questions. Inspired by previous work (Du and Cardie, 2020; Liu et al., 2021), we generate wh-questions for each argument role an event may have according to the event ontology. We use a straightforward template to generate questions: *Wh-word is the [argument role] of event [event trigger]?*, where Wh-word is selected from the following: what, where, who and how. The only requirement is to know a priori the name of the argument roles an event may have, an assumption we share with previous work. Questions are generated regardless of whether the argument is present in the input document. Answers are *No answer* if an argument is not present. For example, in Figure 1, “Q1: *What is the artifact of the event importing?*” is the template-based question for the *artifact* role of event trigger *importing*.

Prompt-Based Questions. Template-based questions rely on templates defined by humans. This is suboptimal when generating questions for new argument roles. We employ large language models to bypass this limitation. Specifically, we use zero- and few-shot prompting (see Appendix C for details) with the GPT-4 model to generate wh-questions for all roles an event may have. The few-shot prompting includes 10 randomly chosen examples and enables in-context learning (Brown et al., 2020). These questions are designed to ask for specific roles without being tied to any specific event document. We refer to these questions as zero- and few-shot prompt-based questions. Q2 and Q3 in Figure 1 are zero- and few-shot prompt-based questions generated for the *artifact of importing*: *What artifact is related to the event importing?* and *What artifact or object is related to the event importing?*

Uncontextualized questions sometimes result in unnatural or ungrammatical questions as exemplified below (*X* is a placeholder for the event trigger):

- *Where is the place of event [X]?* (template)
- Zero Prompt: *Where did the event [X] take place?* (zero-shot prompt)
- Few Prompt: *What is the place or location related to the event [X]?* (few-shot prompt)

Rather than defining complex rules to account for the correct verb tense and formulate better wordings, we work with contextualized questions.

3.2 Contextualized Questions

Uncontextualized questions lack any information included in the document the event trigger belongs to, including other arguments as well as the rest of the document. Contextualized questions address this weakness. We propose a weakly supervised method to generate contextualized questions grounded on the event and document of interest. We explore two strategies to collect data to train models for question generation: (a) leveraging SQuAD (Rajpurkar et al., 2016), a question-answering dataset not related to event argument extraction, and (b) weak supervision from questions automatically obtained by prompting LLMs.

SQuAD-Based Questions. We use the SQuAD (document, questions, and answers) as training data for question generation. The training process involves feeding the model document-answer pairs so that it learns to generate questions. Despite SQuAD is not designed for event argument extraction, this model can generate a (sometimes noisy) question given an event trigger and role. It is worth noting that this model *requires the answers* to generate the questions. Therefore, SQuAD-based questions can only be used at training time.

While these questions are contextualized, and as we shall see it is beneficial to train with them, we found that they are often not grounded on the event of interest or contain misunderstandings. For example, the answer to Q4 from Figure 1, (*What did Russia destroy 500 trucks with?*) is not *oil* (*oil* is the artifact of *importing*).

Weak Supervision from LLMs. To generate more sound contextualized questions, we adopted a targeted approach. We utilized a subset of RAMS (500 random samples, 7% from the training split) to prompt (see details in Appendix C) GPT-4 with documents, event triggers, and the corresponding argument roles. We instruct GPT-4 to generate five wh-questions based on the input prompt, yielding nearly 9,000 questions asking for the arguments of events.

We then train a T5 model to generate questions using weak supervision and the 9,000 questions (output) paired with the documents, event triggers and their roles (input). This model generates role-specific questions not only for RAMS,

but, as we shall see, also for other corpora such as WikiEvents. For example, in Figure 1, Q5: *What is being imported from ISIS-held territory in Syria and Iraq?* is generated for the *artifact* role of the *importing* event. Note that the questions contain other arguments of *importing* (*Syria and Iraq: origin*) as well as additional information from the document (*ISI-held territory*). Additionally, this question does not explicitly mention the role of interest (*artifact*). These questions capture both argument and non-argument cues from the document, facilitating the generation of nuanced and semantically diverse questions. Furthermore, using argument roles as input, as opposed to argument spans, facilitates event-grounded knowledge transfer to other datasets.

4 Experiments

We formulate the problem of event argument extraction as a question-answering task to empirically assess the proposed question generation strategies. Our experiments¹ are designed to determine the optimal strategy as well as the optimal combination of uncontextualized and contextualized questions. Additionally, we conduct zero- and few-shot inference on RAMS to benchmark the GPT-3 model w.r.t our simplest question generation approach.

Models. Template-based uncontextualized questions need no model. We use the GPT-4 model to generate prompt-based uncontextualized questions. For contextualized questions, we fine-tune T5 (Raffel et al., 2020) with either SQuAD or the weakly supervised data obtained by prompting GPT-4 (weak supervision from LLMs). Regardless of how questions are generated, we fine-tune BART (Lewis et al., 2020, base and large versions) to answer them (i.e., find arguments for a given role of an event trigger). We use Pytorch (Paszke et al., 2019) and HuggingFace transformers (Wolf et al., 2020). The only exceptions are GPT-3 and GPT-4, which have their own API.

Combining Uncontextualized and Contextualized Questions. In addition to using a single strategy, we combine uncontextualized and contextualized questions during training. At testing time, we use uncontextualized questions only based on our initial empirical findings (Table 1).

¹Code including dataset transformed into questions and answers are available at <https://github.com/nurakib/event-question-generation>

Datasets. We work with two datasets: RAMS (Ebner et al., 2020) and WikiEvents (Li et al., 2021). The events and arguments in RAMS are annotated based on the AIDA-1² ontology. Documents are sourced from news articles, resulting in 139 event types and 65 argument roles. The event-argument structures in RAMS are annotated within a 5-sentence window; there is one event trigger annotated per document.

WikiEvents is quite different than RAMS. It adopts the KAIROS³ ontology, sources articles from Wikipedia, and encompasses 50 event types and 59 argument roles. WikiEvents offers a more detailed perspective by annotating event-argument structures across entire documents (much longer than five sentences), facilitating the identification of coreferent mentions and allowing annotation of multiple events within the same document.

Table 8 in Appendix B presents basic statistics of the RAMS and WikiEvents datasets.

4.1 Zero- and Few-Shot Inference with GPT-3

Large language models are credited with having emergent abilities (Wei et al., 2022). They are also capable of in-context learning (Brown et al., 2020), meaning that they can (presumably) solve down stream tasks with a small set of data samples when given instructions (Wang et al., 2022). Our methodology uses them to generate questions; these questions are used (a) fine-tune T5 for question generation and (b) BART for event-argument extraction. In order to test the aforementioned abilities when it comes to extracting event-argument structures across sentences, we experiment with GPT-3 with zero- and few-shot prompts:

Zero-Shot. We prompt GPT-3 with the input document (five sentences) and the questions generated with the template-based approach (Section 3.1).

Few-Shot. We prompt GPT-3 with the same zero-shot prompt but preceded by two examples from the training split randomly selected (and using the same format than the expected answer). These examples also include questions without answers.

We present details and examples of the prompts in appendix D. We do not elaborate on them here because they obtain poor results (Section 5.1).

4.2 Evaluation

All our fine-tuned BART models are evaluated using the official evaluation scripts. That is, we

²AIDA-1

³KAIROS

	Base	Large
Uncontextualized Questions		
Template-based	48.5±0.81	52.3±1.72
Prompt-based, zero-shot	48.8±0.88	51.8±1.12
Prompt-based, few-shot	47.9±1.95	53.1±0.68
Contextualized Questions		
SQuAD-based	n/a	n/a
Weak Supervision, LLMs	45.2±1.71	46.5±0.91

Table 1: Results (F1, mean and standard deviation of five runs) obtained with the test split of RAMS using each strategy to generate questions. For uncontextualized questions, zero-shot prompts outperformed the few-shot prompts, but the opposite is true of contextualized questions. Contextualized questions yield substantially lower results than uncontextualized questions, although combining both is the best strategy (Table 2).

use P, R, and F1 after comparing (exact match) predictions and the ground truth based on token spans (i.e., indexes of the first and last tokens, not the text in the token span). Our rationale is that it is the only way to compare with previous work.

For zero- and few-shot inference with GPT-3, we use a more lenient evaluation. Rather than checking for spans, we consider a prediction correct if the generated text includes the text of the expected answer. Since there can be multiple spans in the document matching the output of GPT-3, this evaluation is more lenient and better accommodates generative models. Despite being more lenient, GPT-3 obtains substantially lower results.

5 Quantitative Results and Analyses

5.1 Results with RAMS

Table 1 outlines the results (F1) with our BART models (base and large) with RAMS and using question from a single strategy. For uncontextualized questions, zero-shot prompt-based questions performed best with the base model (F1: 48.8 vs. 45.2–48.5), while few-shot prompt-based questions excelled with the large model (F1: 53.1 vs. 46.5–52.3). Weakly supervised questions negatively impact performance (F1: 45.2 (base) and 46.5 (large)), possibly due to high variation and the requirement of resolving multiple contextual cues. This lead us to recommend uncontextualized questions when using a single strategy to generate questions.

We report the results with BART models finetuned with uncontextualized and contextualized questions in Table 2. Combining template-based

	Base	Large
Template-based		
+ SQuAD-based	49.0±0.86	53.1±0.34*
+ Weak Supervision, LLMs	51.5±0.26*	54.5±0.68*
Prompt-based, zero-shot		
+ SQuAD-based	49.6±0.93	52.7±1.87*
+ Weak Supervision, LLMs	50.3±0.46*	53.4±1.48*
Prompt-based, few-shot		
+ SQuAD-based	49.3±0.53	52.6±0.12
+ Weak Supervision, LLMs	49.7±0.77	53.5±1.31*
Previous Work (Top 5)		
APE (2023)	51.6	54.3
APE (Single) (2023)	49.6	51.7
SPEAE (2023)	51.1	53.3
TabEAE (2023)	n/a	52.7
TARA (2023)	48.0	52.5
PAIE (2022)	49.5	52.2

Table 2: Results (F1, mean and standard deviation of five runs) obtained with the test split of RAMS using combination of uncontextualized and contextualized questions. The best strategy is to blend template-based (uncontextualized) and weak supervision with LLMs (contextualized) questions. We indicate statistically significantly better results (McNemar test (McNemar, 1947), $p < 0.01$) with respect to *Template-Based Questions* with an asterisk (*).

and weakly supervised questions from RAMS yields the best results for both base and large models, showcasing a 3.0 F1 improvement (48.5 vs. 51.5) for the base model and a 2.2 F1 improvement (52.3 vs. 54.5) for the large model. This approach slightly surpasses state-of-the-art models for the large variant. We hypothesize that contextualized (i.e., questions grounded on the event and document of interest) force the model to learn semantic interactions between event arguments and contextual cues due to the framing of the question. Additionally, combining SQuAD-based questions (also contextualized) at training time also benefit results across the board. This is true despite their inherent noise and wording discrepancies related to events other than the event of interest (Section 3.2). The results in Table 1 confirm the advantage of asking uncontextualized questions, and Table 2 shows the consistent benefit of combining uncontextualized and contextualized questions.

Comparison with Previous Work on RAMS.

Our best model slightly outperforms the best published results with RAMS (Table 2, bottom block),

	Base		Large	
	F1	% Δ F1	F1	% Δ F1
2 before	26.89	+53.3	30.53	+6.8
1 before	35.04	+19.5	40.00	+19.1
same	56.19	+3.7	59.38	+11.5
1 after	24.82	+19.3	27.84	+3.5
2 after	15.69	+37.2	21.05	+12.6

Table 3: Results with RAMS and our best model (boldfaced in Table 2) broken down by distance (# sentences) between arguments and events. % Δ F1 indicates the relative improvement with respect to training only with template-based questions (first supervised system in Table 2). Our approach benefits all arguments, especially those that are not in the same sentence than the event.

APE (Zhang et al., 2023), which achieves an F1 of 54.3 (ours: 54.5). APE relies on overlapping role annotations from multiple event argument extraction corpora (ACE05, RAMS, and WikiEvents). Notably, when APE is trained only with RAMS, it achieves 51.7 F1 (ours: 54.5), indicating a significant performance drop compared to other models. Nguyen et al. (2023) proposed generating contextualized soft prompts for event argument extraction. This work achieves the second-closest F1-score to our model. Generating soft prompts depends on a set of relevant documents and manually initialized prompts, whereas our event-grounded question generation approach learns from a small portion of RAMS documents and as we shall see, it transfers to other domains (WikiEvents, Section 5.2).

Are Inter-Sentential Arguments Harder? Table 3 details the results of our best models (combining template-based and weak supervision from LLMs, boldfaced in Table 2) broken down by distance (number of sentences) between the event trigger and argument. The improvements (% Δ F1) with respect to the simplest question-answering approach (only template-based questions; first system in Table 2) are substantial regardless of distance between the event trigger and argument. For the base model, we observe a substantial 37.2% improvement when the argument appears two sentences after the event trigger. Notably, improvements remain considerable for arguments appearing in sentences before (53.5% F1, 19.5% F1) or after (19.3% F1) the event trigger. For the large model, the most significant improvement of 19.1% F1 occurs when arguments are in the sentence before the event trigger, followed by those two sentences




Figure 2: F1 per argument of our best system (boldfaced in Table 2). Frequency in training (between parenthesis) is only a weak indicator of F1, leading to the conclusion that some arguments are easier to learn (e.g., *passenger* is 70% less frequent than *participant* yet the former obtains twice the F1 (0.70 vs. 0.33)).




Figure 3: Average F1 per event (top 15 most frequent events) by our best system (boldfaced in Table 2). There is no clear relation between event frequency in training (between parenthesis) and F1, leading to the conclusion that arguments of some events are easier to learn (e.g., *selfdirectedbattle* vs. *payforservice*)

after (12.6% F1). Arguments from other sentences also benefit (6.8% and 3.5% F1). In essence, our approach demonstrates robustness in extracting event-argument structures regardless of the distance between arguments and event triggers. The benefits, however, are most noticeable when arguments are not in the same sentence than the event trigger.

Are Frequent Arguments Easier? It is a common belief that the more training data the easier it is to learn. Figure 2 provides empirical evidence showing that this is not the case when predicting event-argument structures in RAMS. We observe that per-argument F1 scores range from 0.33 to 0.70, but there is no pattern indicating that frequency correlates with F1 score. For example, infrequent events such as *employee* and *passen-*

	place	recip	part	commu	artif	victi	giver	targe	attac	benef	instr	trans	origi	desti	defen
place	87	4	1	1	1	0	1	3	0	0	0	0	0	0	0
recipient	5	76	0	4	3	0	1	0	0	10	0	0	0	0	0
participant	6	0	88	0	0	0	0	0	0	4	0	0	0	0	0
communicator	2	1	0	96	0	0	0	0	0	0	0	0	0	0	0
artifact	2	5	0	0	86	0	0	2	0	0	1	1	0	0	0
victim	0	0	0	0	0	100	0	0	0	0	0	0	0	0	0
giver	0	8	0	0	0	0	80	0	0	12	0	0	0	0	0
target	16	0	0	0	2	0	0	81	0	0	0	0	0	0	0
attacker	0	0	0	0	0	0	2	97	0	0	0	0	0	0	0
beneficiary	0	27	9	0	0	0	9	0	0	54	0	0	0	0	0
instrument	0	0	0	0	0	0	0	0	0	0	100	0	0	0	0
transporter	0	0	0	0	0	0	0	0	0	0	0	80	20	0	0
origin	0	0	0	0	0	0	0	0	0	0	0	19	76	4	0
destination	0	0	0	0	0	0	0	0	0	0	0	2	5	92	0
defendant	0	0	0	0	0	0	0	0	0	0	0	0	0	0	100

Figure 4: Confusion matrix comparing gold (rows) and predicted (columns) argument roles for correctly predicted argument spans (top 15 most frequent types). Most errors are plausible (at face value) but semantically wrong argument roles (e.g., mislabeling the *beneficiary* as the *recipient*; note that both are usually people).

ger (2%) obtain results as high as those obtained with communicator (6%) and victim (5%).

Are Frequent Events Easier? No, they are not. Surprisingly, more training data for an event does not always lead to better results. Figure 3 shows the average F1 scores for the 15 most frequent events. The graph shows no clear pattern between event frequency in training and F1 scores. Indeed, arguments of events with 2% frequency obtain F1 scores ranging from 0.33 and 0.60, a large range overlapping with the F1 scores of more frequent events.

Which Arguments are Mislabeled? Our best model obtains 54.5 F1. This number is low, but the evaluation is as strict as it gets: it expects predictions to match the exact argument span and argument role. Figure 4 compares gold (rows) and predicted (columns) argument roles when our best model (boldfaced in Table 2) predicts the correct argument spans. We observe two main trends. First, the model mislabels arguments with labels that are (in principle) plausible but wrong given the input document. For example, *recipient*, *beneficiary*, and *giver* are often people but they have different semantics given an event trigger. Second, our model mislabels arguments that could be considered correct but do not follow the RAMS annotations. For example, the *transporter* of a *transport-*

	P	R	F1
Ours, Base	57.1±0.48	42.1±1.26	48.5±0.81
Ours, Large	56.0±1.31	49.1±2.14	52.3±1.76
GPT3			
Zero-shot	27.3	21.4	24.0
Few-shot	32.6	29.1	30.7

Table 4: Results obtained with our simplest question generation strategy (*template-based*) and GPT-3 (*text-davinci-003*) with zero- and few-shot prompts.

	Base	Large
Template-Based Questions + Weak Supervision, LLMs	63.6±1.24	64.7±0.76
	66.3±0.24*	67.7±1.13*
Previous Work (Top 5)		
APE (2023)	66.0	68.7
APE (Single) (2023)	62.1	65.8
TabEAE (2023)	n/a	66.5
SPEAE (2023)	66.2	66.1
PAIE (2022)	63.4	65.3
AHR (2023)	46.1	63.4

Table 5: Results with WikiEvents (F1, mean and standard deviation of five runs) using template-based questions and the best strategy with RAMS (boldfaced in Table 2). We reuse the weakly supervised question generation model trained on a sample of RAMS. Results are statistically significantly better (McNemar test (McNemar, 1947), $p < 0.01$, indicated with *).

ing event (i.e., the person moving something) could be the *origin* or the event, but RAMS reserves that argument role for the location where *transporting* started. We hypothesize that our model is leveraging the knowledge acquired about *transporter* and *origin* prior to our training with RAMS (and it never overcame this knowledge during training).

Comparison with GPT-3. Zero-shot and few-shot prompting with GPT-3 falls short in regards to performance with respect to our question-answering approach.⁴ The results in Table 4 show that GPT-3 obtains better results in a few-shot in-context learning setting compared to a zero-shot setting. Yet, the performance lags behind our template-based question answering model. It is important to acknowledge the potential of prompt engineering (Gu et al., 2022) to improve performance. We reserve the avenue of prompt engineering for future work.

⁴We only report one run for zero- and few-shot with GPT-3 as we do not tune it.

Error Type	Example
Alternatives (42.17%)	[...] The military operation is the most complex carried out in [Iraq] ^{PREDICTED_PLACE} since US forces withdrew from [the country] ^{GOLD_PLACE} in 2011. Last week, the UN said it was bracing itself for the world’s biggest and most complex humanitarian effort following the [battle] ^{EVENT_TRIGGER} , which it expects will displace up to one million people [...]
Partial spans (33.73%)	[...] US Secretary of State Hillary Clinton who allowed this Gulf State nation to secure the 2022 World Cup — and that [the [Qataris] ^{PREDICTED_GIVER}] ^{GOLD_GIVER} were so appreciative of they [donated] ^{EVENT_TRIGGER} millions to [the [Clinton Foundation.] ^{PREDICTED_BENEFICIARY}] ^{GOLD_BENEFICIARY} [...]
Wrong spans (15.66%)	[...] Mr. Assange spoke from the [Ecuadorean Embassy] ^{GOLD_PREVENTER} in London, where he has been holed up for four years. [Sweden] ^{PREDICTED_PREVENTER} is seeking his extradition for an investigation into sexual misconduct allegations; his supporters have expressed fear that if he is [arrested] ^{EVENT_TRIGGER} , he could be sent to the United States and prosecuted for publishing leaked documents. [...]
Two or more arguments (8.04%)	[...] On Monday, the [U.S.] ^{GOLD_PARTICIPANT_2} and [Russia] ^{GOLD_PARTICIPANT_1} predicted_participant entered a [ceasefire] ^{EVENT_TRIGGER} agreement in Syria, in part to get humanitarian aid to cities like Aleppo. The ceasefire does not apply to strikes targeting ISIS. [...]

Table 6: Most common errors made by our best performing model on RAMS (combining template-based questions with questions obtained with a weakly supervised model, boldface in Table 2).

5.2 Results with WikiEvents

We also evaluate the effectiveness of our proposed question generation strategy (best performing approach with RAMS, boldfaced in Table 2) on the WikiEvents dataset. To formulate uncontextualized questions, we employed a template-based approach utilizing the event ontology provided with the dataset. We reuse the question generation models used with RAMS, including the weakly supervised model trained with questions obtained via prompting with documents, event triggers and arguments from RAMS. In other words, there is no WikiEvents-specific component or fine-tuning.

The results (Table 5) demonstrate that our approach transfers across corpora—recall that WikiEvents contains documents in different domains than RAMS as well as different events and argument roles. Indeed, the best strategy to combine uncontextualized and contextualized questions yields statistically significantly higher results. While we do not outperform all state-of-the-art models, particularly for larger models, our results remain competitive despite the question generation component was trained with RAMS.

6 Qualitative Error Analysis

We close our analyses by examining the errors made by the best model (boldfaced in Table 2). We discuss linguistic commonalities in either the input documents or model predictions observed

in a manual analysis of 83 errors from 100 documents. The most common error type (42.17%, Table 6) is predicting *alternative spans*. In the example, our model predicts a coreferent mention that is counted as an error. This allowed us to raise the issue of rethinking the evaluation of event argument extraction; only one gold span matching may not be ideal. The second most common error type (33.73%, Table 6) is predicting *partial spans* (either shorter or longer than the gold). The differences include articles, conjunctions, numbers, and detailed descriptions complementing entities. Completely *wrong spans* are much less likely (15.66%). Our model is limited to predicting one span per argument role, thus it will always make errors with events that have two instances of the same argument role (8.04% of errors). A previously reported by Zhang et al. (2020), we found that some model errors (7.3%) appear to be due to annotation errors—no annotations are perfect, and RAMS is not an exception. For example, the test set includes the following: *he raised the [funds]^{RECIPIENT} privately and [reimbursed]^{EVENT_TRIGGER} the city [...]*.

7 Weakly Generated Question Analysis

Our empirical assessment across two corpora for event argument extraction indicates that training with weakly supervised questions yields positive results. To further understand this, we conduct a manual analysis of 100 random questions generated

Question about the event and the argument role	Question about the event	Question about neither the event nor the argument role
<p>Document: This is why Turkey became anxious when Russia began delivering airstrikes against the IS infrastructure and destroyed more than 500 trucks with oil already. This really got on Erdogan and his company's nerves. They're importing not only oil, but wheat and historic artefacts as well." Bilal Erdogan denies Russian allegations Bilal Erdogan this week denied continuous Russian allegations that he and his family were profiting from the illegal smuggling of oil from ISIS - held territory in Syria and Iraq, after Russia's defence ministry said it had proof that the Erdogan family was benefiting directly from this trade.</p>	<p>Document: WikiLeaks says it has access to NSA source code and will be posting it for free, unlike an outfit known as the "Shadow Brokers" that had announced plans to auction off code they allegedly stole. Univision won the bidding auction for Gawker and will buy all seven of the company's media properties for \$135 million. Pay TV providers posted their highest subscriber loss ever in this year's second quarter, with cable, satellite and telecom companies losing at least 665,000 subscribers.</p>	<p>Document: People who had just arrived at the airport in Istanbul after the blasts. Desperate passengers embrace as they gather in shock outside the terminal after the airport was evacuated and all flights were grounded. Twelve months after that, he was found in Georgia, in the Lopota mountain valley following a battle between police and rebels which saw three policeman and 11 rebels killed. Chataev was injured and later detained - but said he had nothing to do with it. Chataev is also said to have led Chechen Islamist Doku Umarov's armed underground movement of Northern Caucasus from Turkey.</p>
<p>Event: importing</p>	<p>Event: auction</p>	<p>Event: battle</p>
<p>Argument: Oil (Artifact)</p>	<p>Argument: \$135 million (Money)</p>	<p>Argument: No Answer (Instrument)</p>
<p>Generated Question: What is being imported from ISIS-held territory in Syria and Iraq?</p>	<p>Generated Question: What was the key component in the auction that led to the purchase of all seven of Gawker's media properties?</p>	<p>Generated Question: What led to the capture of Chataev in the Lopota mountain valley in Georgia?</p>
<p>Expected Question: What is being imported from ISIS-held territory in Syria and Iraq?</p>	<p>Expected Question: How much money did Univision bid to win the auction for Gawker, acquiring all seven of the company's media properties?</p>	<p>Expected Question: What instrument was used in the battle between police and rebels take place?</p>

Figure 5: Examples of three categories of questions generated by the weakly supervised T5 model. Text highlighted in green indicates the events, while red text indicates the arguments within the RAMS documents. Argument roles are mentioned in parenthesis. *Question about the event and the argument role* are grounded on both the event and the argument role, and *Question about the event* are only grounded on the event. *Question about neither the event or the argument role* are irrelevant to the event and the argument role. *Generated Question* are the outputs generated by the weakly supervised model; and *Expected Question* refers to human-written questions exemplifying perfect questions (hypothetical) we would like to generate.

by the weakly supervised model. This analysis evaluates the generated questions for their relevance and accuracy in relation to the events and argument roles detailed within the document. Overall, we found three different categories after analyzing 100 questions. The first category of questions are targeting both the event and the argument role, while the second category of questions focuses solely on the event. The third category of questions lacks relevance to both the event and the argument role. Figure 5 exemplifies the 3 categories using the output from the weakly supervised T5 model. We observe that 73% questions generated are grounded in both the event and the argument role (*first category*). Another 23% questions are grounded in the event of interest (*second category*) but aim to extract different information than the role of interest. Only 4% questions are grounded in neither the event nor argument role of interest (*third category*).

8 Conclusions

We have presented several question generation strategies for asking questions about the arguments of an event. Our approach can be used seamlessly with any corpora; the only requirement is a list of events and the arguments they may take, information that existing corpora come with. Combining uncontextualized and contextualized questions obtains the best results. Contextualized questions incorporate cues about other arguments of the event of interest as well as other parts of the document, providing our BART-based models with information that is not otherwise available. Crucially, our strategies to generate questions do not require manual intervention. Additionally, the weakly supervised models to generate questions transfer across corpora (from RAMS to WikiEvents; different corpora, events, and arguments roles).

Limitations

This work presented in this paper has several limitations. Our weakly supervised approach brings certain advantages in event argument extraction. This has been assessed through empirical analysis. While we have observed the effectiveness of incorporating questions during the training phase, we have not reported an in-depth analysis into why this method is successful across all samples. It remains unclear whether the addition of all contextualized questions is universally beneficial. To address this, employing an adaptive selection strategy could prove valuable in determining the instances where incorporating questions is most advantageous.

Our model predicts at most one argument per event trigger. Thus, any event trigger that has two arguments with the same role is guaranteed to yield an error (Section 6). Addressing this limitation would necessitate further research, particularly in the context of multi-turn question answering, to enhance the model’s capabilities.

We compare our fine-tuned model and GPT-3 with template-based questions. It can be argued that there might be a better alternative way of prompting the GPT-3 model.

Acknowledgments

This research was supported by a grant from the U.S. Office of Naval Research (N00014-22-1-2596). We would like to thank the anonymous reviewers for their insightful comments and suggestions.

Ethics Statement

The authors state that this work is in accordance with the ACL Code of Ethics and does not raise ethical issues.

References

Jacqueline Aguilar, Charley Beller, Paul McNamee, Benjamin Van Durme, Stephanie Strassel, Zhiyi Song, and Joe Ellis. 2014. [A comparison of the events and relations across ACE, ERE, TAC-KBP, and FrameNet annotation standards](#). In *Proceedings of the Second Workshop on EVENTS: Definition, Detection, Coreference, and Representation*, pages 45–53, Baltimore, Maryland, USA. Association for Computational Linguistics.

David Ahn. 2006. [The stages of event extraction](#). In *Proceedings of the Workshop on Annotating and Rea-*

soning about Time and Events, pages 1–8, Sydney, Australia. Association for Computational Linguistics.

Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. [The Berkeley FrameNet project](#). In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pages 86–90, Montreal, Quebec, Canada. Association for Computational Linguistics.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Yubo Chen, Liheng Xu, Kang Liu, Daojian Zeng, and Jun Zhao. 2015. [Event extraction via dynamic multi-pooling convolutional neural networks](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 167–176, Beijing, China. Association for Computational Linguistics.

Yunmo Chen, Tongfei Chen, and Benjamin Van Durme. 2020. [Joint modeling of arguments for event understanding](#). In *Proceedings of the First Workshop on Computational Approaches to Discourse*, pages 96–101, Online. Association for Computational Linguistics.

George Doddington, Alexis Mitchell, Mark Przybocki, Lance Ramshaw, Stephanie Strassel, and Ralph Weischedel. 2004. [The automatic content extraction \(ACE\) program – tasks, data, and evaluation](#). In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC’04)*, Lisbon, Portugal. European Language Resources Association (ELRA).

Xinya Du and Claire Cardie. 2020. [Event extraction by answering \(almost\) natural questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 671–683, Online. Association for Computational Linguistics.

Xinya Du and Heng Ji. 2022. [Retrieval-augmented generative question answering for event argument extraction](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4649–4666, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Xinya Du, Alexander Rush, and Claire Cardie. 2021. [GRIT: Generative role-filler transformers for document-level event entity extraction](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 634–644, Online. Association for Computational Linguistics.

- Seth Ebner, Patrick Xia, Ryan Culkin, Kyle Rawlins, and Benjamin Van Durme. 2020. **Multi-sentence argument linking**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8057–8077, Online. Association for Computational Linguistics.
- Matthew Gerber and Joyce Chai. 2010. **Beyond Nom-Bank: A study of implicit arguments for nominal predicates**. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1583–1592, Uppsala, Sweden. Association for Computational Linguistics.
- Ralph Grishman and Beth Sundheim. 1996. **Message Understanding Conference- 6: A brief history**. In *COLING 1996 Volume 1: The 16th International Conference on Computational Linguistics*.
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2022. **PPT: Pre-trained prompt tuning for few-shot learning**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8410–8423, Dublin, Ireland. Association for Computational Linguistics.
- Rujun Han, I-Hung Hsu, Jiao Sun, Julia Baylon, Qiang Ning, Dan Roth, and Nanyun Peng. 2021. **ESTER: A machine reading comprehension dataset for reasoning about event semantic relations**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7543–7559, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yuxin He, Jingyue Hu, and Buzhou Tang. 2023. **Re-visiting event argument extraction: Can EAE models learn better when being aware of event co-occurrences?** In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12542–12556, Toronto, Canada. Association for Computational Linguistics.
- Yu Hong, Jianfeng Zhang, Bin Ma, Jianmin Yao, Guodong Zhou, and Qiaoming Zhu. 2011. **Using cross-entity inference to improve event extraction**. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1127–1136, Portland, Oregon, USA. Association for Computational Linguistics.
- Yin Jou Huang, Jing Lu, Sadao Kurohashi, and Vincent Ng. 2019. **Improving event coreference resolution by learning argument compatibility from unlabeled data**. 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 785–795, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. **BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Fayuan Li, Weihua Peng, Yuguang Chen, Quan Wang, Lu Pan, Yajuan Lyu, and Yong Zhu. 2020. **Event extraction as multi-turn question answering**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 829–838, Online. Association for Computational Linguistics.
- Qi Li, Heng Ji, and Liang Huang. 2013. **Joint event extraction via structured prediction with global features**. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 73–82, Sofia, Bulgaria. Association for Computational Linguistics.
- Sha Li, Heng Ji, and Jiawei Han. 2021. **Document-level event argument extraction by conditional generation**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 894–908, Online. Association for Computational Linguistics.
- Wei Li, Lei He, and Hai Zhuge. 2016. **Abstractive news summarization based on event semantic link network**. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 236–246, Osaka, Japan. The COLING 2016 Organizing Committee.
- Shasha Liao and Ralph Grishman. 2010. **Using document level cross-event inference to improve event extraction**. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 789–797, Uppsala, Sweden. Association for Computational Linguistics.
- Jian Liu, Yubo Chen, Kang Liu, Wei Bi, and Xiaojiang Liu. 2020. **Event extraction as machine reading comprehension**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1641–1651, Online. Association for Computational Linguistics.
- Jian Liu, Yufeng Chen, and Jinan Xu. 2021. **Machine reading comprehension as data augmentation: A case study on implicit event argument extraction**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2716–2725, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Qing Lyu, Hongming Zhang, Elicor Sulem, and Dan Roth. 2021. **Zero-shot event extraction via transfer learning: Challenges and insights**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*

- (*Volume 2: Short Papers*), pages 322–332, Online. Association for Computational Linguistics.
- Yubo Ma, Zehao Wang, Yixin Cao, Mukai Li, Meiqi Chen, Kun Wang, and Jing Shao. 2022. [Prompt for extraction? PAIE: Prompting argument interaction for event argument extraction](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6759–6774, Dublin, Ireland. Association for Computational Linguistics.
- Quinn McNemar. 1947. [Note on the sampling error of the difference between correlated proportions or percentages](#). *Psychometrika*, 12(2):153–157.
- Adam Meyers, Ruth Reeves, Catherine Macleod, Rachel Szekely, Veronika Zielinska, Brian Young, and Ralph Grishman. 2004. [Annotating noun argument structure for NomBank](#). In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal. European Language Resources Association (ELRA).
- Chien Nguyen, Hieu Man, and Thien Nguyen. 2023. [Contextualized soft prompts for extraction of event arguments](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4352–4361, Toronto, Canada. Association for Computational Linguistics.
- Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. 2016. [Joint event extraction via recurrent neural networks](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 300–309, San Diego, California. Association for Computational Linguistics.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. [The Proposition Bank: An annotated corpus of semantic roles](#). *Computational Linguistics*, 31(1):71–106.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- 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.
- Alan Ramponi, Rob van der Goot, Rosario Lombardo, and Barbara Plank. 2020. [Biomedical event extraction as sequence labeling](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5357–5367, Online. Association for Computational Linguistics.
- Yubing Ren, Yanan Cao, Ping Guo, Fang Fang, Wei Ma, and Zheng Lin. 2023. [Retrieve-and-sample: Document-level event argument extraction via hybrid retrieval augmentation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 293–306, Toronto, Canada. Association for Computational Linguistics.
- Josef Ruppenhofer, Caroline Sporleder, Roser Morante, Collin Baker, and Martha Palmer. 2010. [SemEval-2010 task 10: Linking events and their participants in discourse](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 45–50, Uppsala, Sweden. Association for Computational Linguistics.
- Oscar Sainz, Itziar Gonzalez-Dios, Oier Lopez de Lacalle, Bonan Min, and Eneko Agirre. 2022. [Textual entailment for event argument extraction: Zero and few-shot with multi-source learning](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2439–2455, Seattle, United States. Association for Computational Linguistics.
- Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, and Yi Zhang. 2022. [Multi-task pre-training for plug-and-play task-oriented dialogue system](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4661–4676, Dublin, Ireland. Association for Computational Linguistics.
- Md Nayem Uddin, Enfa Rose George, Eduardo Blanco, and Steven R. Corman. 2024. Asking and answering questions to extract event-argument structures. In *Proceedings of Joint International Conference on Computational Linguistics, Language Resources and Evaluation*, Turin, Italy.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. Ace 2005 multilingual training corpus. *Linguistic Data Consortium, Philadelphia*, 57:45.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krима Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sampat, Siddhartha Mishra, Sujan Reddy A, Sumanta Patro, Tanay Dixit, and Xudong

- Shen. 2022. [Super-NaturalInstructions: Generalization via declarative instructions on 1600+ NLP tasks](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5085–5109, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.
- Kaiwen Wei, Xian Sun, Zequn Zhang, Jingyuan Zhang, Guo Zhi, and Li Jin. 2021. [Trigger is not sufficient: Exploiting frame-aware knowledge for implicit event argument extraction](#). 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)*, pages 4672–4682, Online. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Runxin Xu, Peiyi Wang, Tianyu Liu, Shuang Zeng, Baobao Chang, and Zhifang Sui. 2022. [A two-stream AMR-enhanced model for document-level event argument extraction](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5025–5036, Seattle, United States. Association for Computational Linguistics.
- Yuqing Yang, Qipeng Guo, Xiangkun Hu, Yue Zhang, Xipeng Qiu, and Zheng Zhang. 2023. [An AMR-based link prediction approach for document-level event argument extraction](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12876–12889, Toronto, Canada. Association for Computational Linguistics.
- Kaihang Zhang, Kai Shuang, Xinyue Yang, Xuyang Yao, and Jinyu Guo. 2023. [What is overlap knowledge in event argument extraction? APE: A cross-datasets transfer learning model for EAE](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 393–409, Toronto, Canada. Association for Computational Linguistics.
- Zhisong Zhang, Xiang Kong, Zhengzhong Liu, Xuezhe Ma, and Eduard Hovy. 2020. [A two-step approach for implicit event argument detection](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7479–7485, Online. Association for Computational Linguistics.
- Zhisong Zhang, Emma Strubell, and Eduard Hovy. 2022. [Transfer learning from semantic role labeling to event argument extraction with template-based slot querying](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2627–2647, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

A QA Model and Hyperparameters

In this section, we provide an overview of the model and hyperparameters used for finetuning question answering models. We leverage the BART model to generate representations for the question and document pairs. By feeding the question and the document as input to the BART model, we obtained the contextualized representation of the combined text. Then, we employ a task-specific layer that operates on top of these representations. This layer is responsible for predicting the start and end offsets of the argument spans.

The output layer is trained using the cross-entropy loss function to minimize the discrepancy between the predicted offsets and the ground truth offsets. Initially, we train all the models using the same random seed, 42. We use five different learning rates [1e-5, 2e-5, 3e-5, 4e-5, 5e-5] and dropout values [0.1, 0.2, 0.3] to tune the performance of the models. In order to determine the optimal hyperparameters, we evaluate the models’ performance on the validation dataset. This allows us to select the final hyperparameters that yield the best result. Based on the optimized values for learning rate and dropout, we retrain the model on four new random seeds (4, 13, 52, 57) and report the F1-scores based on the average and standard deviation of the five runs. We list all the hyperparameters in Table 7.

To mitigate the risk of overfitting and ensure efficient training, we incorporated the technique of early stopping. If the loss function fails to show improvement over 10 consecutive epochs, training is terminated before completing 50 epochs.

B Datasets Statistics

RAMS and WikiEvents datasets share the common objective of annotating events and their associated arguments within a document. RAMS is licensed

Name	Value
learning rates	1e-5, 2e-5, 3e-5, 4e-5, 5e-5
epochs	50
patience	10
dropout	0.1, 0.2, 0.3
training batch	8
validation batch	4
test batch	4
max length	512
loss function	cross-entropy
optimizer	Adam

Table 7: Hyperparameters to finetune the question answering model.

Dataset	Split	#Docs	#Events	#Args
RAMS	Train	3194	7329	17026
	Dev	399	924	2188
	Test	400	871	2023
WikiEvents	Train	206	3241	4542
	Dev	20	345	428
	Test	20	365	566

Table 8: Statistics of the RAMS and WikiEvents dataset.

under *Apache License 2.0* and WikiEvents is licensed under *MIT License*. Table 8 presents basic statistics of the RAMS and WikiEvents datasets.

C Question Generation Prompts

We utilize the GPT-4 model to create event argument role specific questions in two scenarios: uncontextualized and contextualized. In the uncontextualized scenario, we generate one question for each argument role. For contextualized questions, we prompt the model using 500 samples from RAMS and train a T5-base model to generate questions for the remaining data samples.

Uncontextualized Question Generation Prompts.

We generate uncontextualized questions using the the GPT-4 model in OpenAI playground.⁵ Initially, we prompt the model with only argument role information, which we term as zero-shot prompt-based questions. Figure 6 presents the zero-shot prompt-based question generation. We start with a task description, giving the model an overview of the task at hand. The all the argument roles are included; followed by task specific instruction. We

⁵OpenAI Playground.

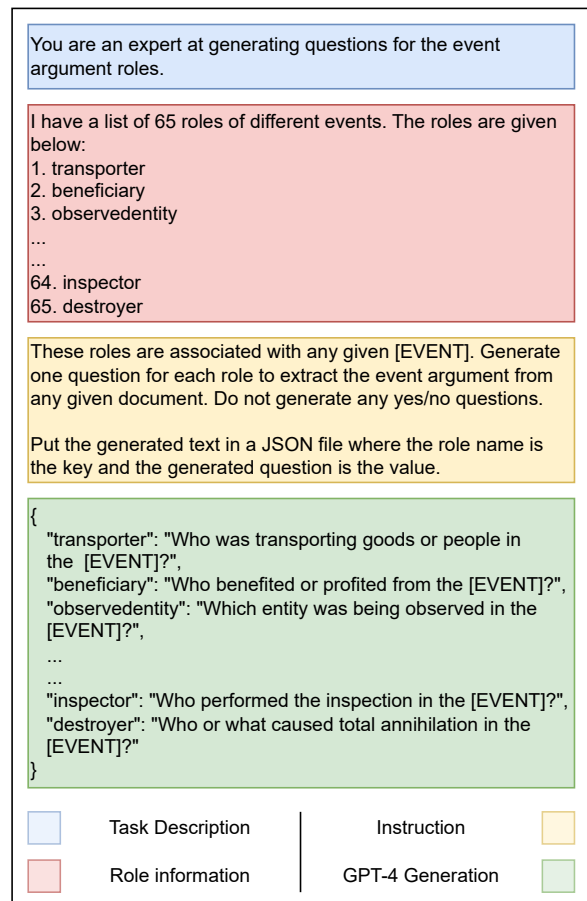


Figure 6: Example of a zero-prompt question generation. *Task Description*, *Role information*, *Instruction* are all together considered the prompt (*input*) and *GPT-4 Generation* is the model output.

instruct the model to produce only one question per role and to avoid generating yes/no questions. Finally, we conclude the instructions by asking the model to format the generated questions in JSON.

For few-shot prompt-based questions generation, we add ten samples from zero-shot prompt-based questions in the prompt. Figure 7 presents the zero-shot prompt-based question generation.

Contextualized Question Generation Prompt.

We prompt the GPT-4 model to collect around 9000 event-grounded question samples. We begin with a task description, outlining the information required for generating event-grounded questions. Next, we provide the event document from RAMS along with the task instruction. Specifically, we instruct the model to produce five distinct questions for a given event and an argument role. Although the actual event argument span is highlighted in the document, our instruction is focused on the event

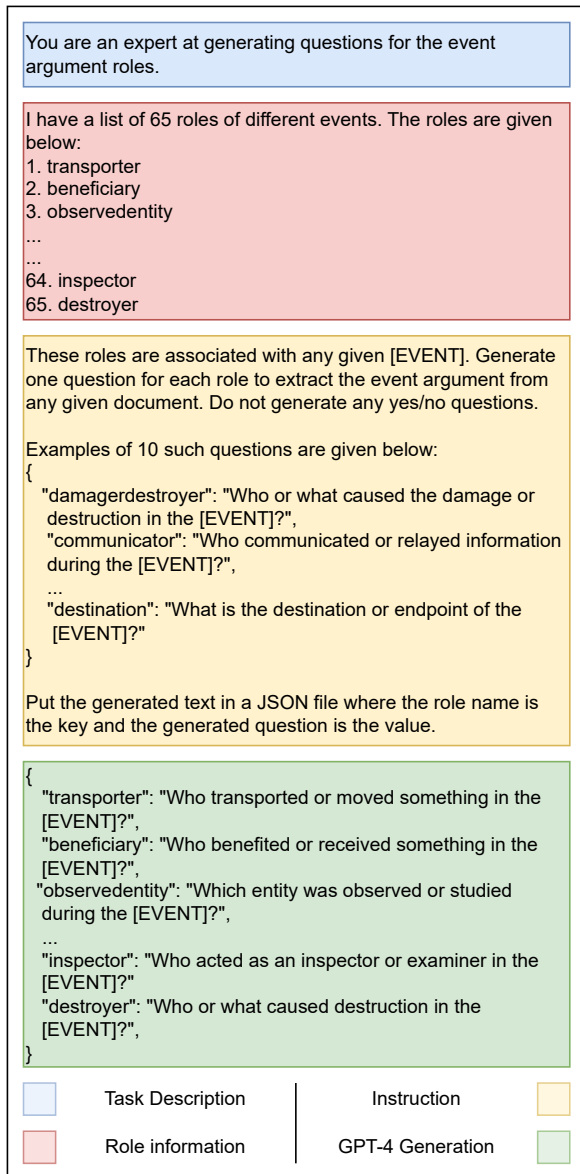


Figure 7: Example of a few-prompt question generation. We added some role-specific question samples in the *Instruction*. *Task Description*, *Role information*, *Instruction* are all together considered the prompt (*input*) and *GPT-4 Generation* is the model output.

and argument role. The model then generates five different event-grounded questions as instructed.

D Zero- and Few-Shot Inference Prompts

To conduct our experiments with GPT-3 in zero and few-shot settings, we utilized the OpenAI API. Specifically, we use the ‘text-davinci-003’ model and the ‘Completion’ endpoint provided by the API. To ensure consistency with the inputs used in the supervised model, we designed the prompts for the GPT-3 model in a similar manner. However,

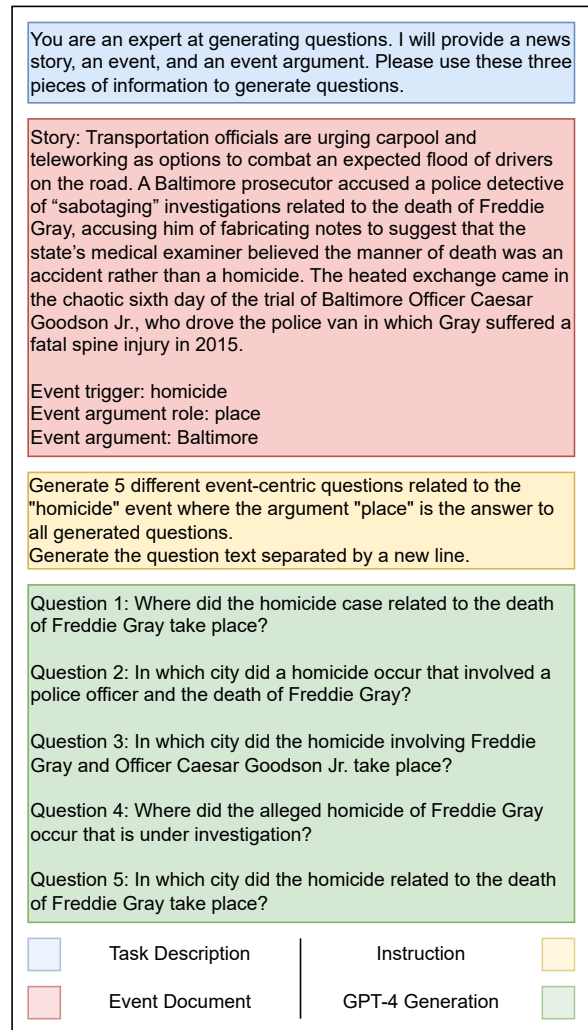


Figure 8: Example of a contextualized question generation using GPT-4. *Task Description*, *Event Document*, *Instruction* are all together considered the prompt (*input*) and *GPT-4 Generation* is the model output.

there were slight differences in how the prompts were handled between the supervised settings and the zero and few-shot experiments. In the supervised settings, we posed one question per iteration, whereas, for the GPT-3 zero-shot and few-shot experiments, we included all the questions simultaneously. We conducted a small-scale study using a subset of samples from the RAMS dataset to validate the impact of asking all questions at once compared to asking one question per iteration. Our study did not reveal any difference between the two approaches. Hence, we proceeded with asking all questions at once for the zero and few-shot experiments. This streamlined the experimentation process and also helped to reduce the costs of querying the API.

Example of Zero-Shot Inference Prompt. In the zero-shot setting, our objective is to extract event arguments without any training examples. To accomplish this, we construct prompts with template-based questions. The GPT-3 model generates answers to these questions, which are then mapped back to the provided document to extract matched event argument spans. Figure 9 presents a snapshot of a zero-shot prompt.

Context: They all fly the Maltese flag. In addition to Russian accusations, Syrian Information Minister Omran Zoabi also recently alleged that Turkey downed the Russian bomber over Syria in November in response to the destruction of hundreds of truck oil tankers sent to Turkey from Syria by the ISIS. The information minister alleged that oil **smuggled** into Turkey was bought by the Turkish president's son, who owns an oil company. Mr al - Zoubi said in an interview, "All of the oil was delivered to a company that belongs to the son of Recep [Tayyip] Erdogan. This is why Turkey became anxious when Russia began delivering airstrikes against the IS infrastructure and destroyed more than 500 trucks with oil already.

Answer these 5 questions based on the given context. Output a text span from the context only. If any of the questions is not answerable from the context information, output "No Answer" for that question.

Question 1: Who is the passenger of the event smuggled?
 Question 2: Where is the origin of the event smuggled?
 Question 3: Where is the destination of the event smuggled?
 Question 4: Who is the transporter of the event smuggled?
 Question 5: What is the vehicle of the event smuggled?

Answer 1: No Answer
 Answer 2: Syria
 Answer 3: Turkey
 Answer 4: ISIS
 Answer 5: Oil truck tankers

 Test Sample	 Instruction
 Asking Questions	 GPT-3 Generation

Figure 9: Example of a zero-shot GPT-3 prompt. *Test sample*, *Instruction* and *Asking Questions* are all together considered the prompt (*input*) and *GPT-3 Generation* is the model output.

Example of Few-Shot Inference Prompt. In the few-shot setting, we leverage a limited amount of training data. We randomly select two training samples from the RAMS dataset. These examples are formatted to match the inputs used during supervised training. By incorporating two training samples, we enhance the model's ability to capture event arguments and generate accurate responses. Figure 10 presents a snapshot of a few-shot prompt.

Example 1:
 Context: A senior member of Saudi Arabia's royal family **bought** a £452 million yacht before helping push through drastic austerity measures within the country. Deputy Crown Prince Mohammed bin Salman picked out a Russian tycoon's 440 ft ship while holidaying in the south of France, according to the New York Times.
 Question 1: Who is the giver of the event bought?
 Question 2: How much is the money of the event bought?
 Question 3: Who is the recipient of the event bought?
 Question 4: Who is the beneficiary of the event bought?
 Question 5: Where is the place of the event bought?
 Answer 1: senior member of Saudi Arabia's royal family
 Answer 2: £ 452 million
 Answer 3: No Answer
 Answer 4: No Answer
 Answer 5: No Answer

Example 2:
 Context: A vote for Hillary Clinton is a vote for more meddling, intervention, and war, with more dead Americans and wasted dollars, and ultimately even more meddling, intervention, and war. She cloaks her constant push for war with praise of "American exceptionalism" and America's role as "the indispensable nation." In her speech to the American Legion she cited Ronald Reagan's belief in America as a "shining city on a hill," even though he **urged** the U.S. to lead by example, not by becoming an international dominatrix. In fact, Reagan was a veritable peacenik in comparison to Clinton, embracing missile defense out of his horror at the prospect of war.
 Question 1: Who is the communicator of the event urged?
 Question 2: Who is the recipient of the event urged?
 Question 3: Where is the place of the event urged?
 Answer 1: he
 Answer 2: U.S.
 Answer 3: No Answer

Context: They all fly the Maltese flag. In addition to Russian accusations, Syrian Information Minister Omran Zoabi also recently alleged that Turkey downed the Russian bomber over Syria in November in response to the destruction of hundreds of truck oil tankers sent to Turkey from Syria by the ISIS. The information minister alleged that oil **smuggled** into Turkey was bought by the Turkish president's son, who owns an oil company. Mr al - Zoubi said in an interview, "All of the oil was delivered to a company that belongs to the son of Recep [Tayyip] Erdogan. This is why Turkey became anxious when Russia began delivering airstrikes against the IS infrastructure and destroyed more than 500 trucks with oil already.

Answer these 5 questions based on the given context. Output a text span from the context only. If any of the questions is not answerable from the context information, output "No Answer" for that question.

Question 1: Who is the passenger of the event smuggled?
 Question 2: Where is the origin of the event smuggled?
 Question 3: Where is the destination of the event smuggled?
 Question 4: Who is the transporter of the event smuggled?
 Question 5: What is the vehicle of the event smuggled?

Answer 1: No Answer
 Answer 2: Syria
 Answer 3: Turkey
 Answer 4: ISIS
 Answer 5: Oil tankers

 Training Samples	 Test Sample	 Instruction
 Asking Questions	 GPT-3 Generation	

Figure 10: Example of a few-shot GPT-3 prompt. *Training samples*, *Test sample*, *Instruction* and *Asking Questions* are all together considered the prompt (*input*) and *GPT-3 Generation* is the model output.