Are *Triggers* Needed for Document-Level Event Extraction?

Shaden Shaar*, **Wayne Chen***, **Maitreyi Chatterjee, Barry Wang, Wenting Zhao, and Claire Cardie**
Cornell University, USA
{ss2753, zc272}@cornell.edu

Abstract

Most existing work on event extraction has focused on sentence-level texts and presumes the identification of a *trigger*-span—a word or phrase in the input that evokes the occurrence of an event of interest. Event arguments are then extracted with respect to the trigger. Indeed, triggers are treated as integral to, and trigger detection as an essential component of, event extraction. In this paper, we provide the first investigation of the role of triggers for the more difficult and much less studied task of document-level event extraction. We analyze their usefulness in multiple end-to-end and pipelined transformer-based event extraction models for three document-level event extraction datasets, measuring performance using triggers of varying quality (human-annotated, LLM-generated, keyword-based, and random). We find that whether or not systems benefit from explicitly extracting triggers depends both on dataset characteristics (i.e., the typical number of events per document) and task-specific information available during extraction (i.e., natural language event schemas). Perhaps surprisingly, we also observe that the mere existence of triggers in the input, even random ones, is important for prompt-based in-context learning approaches to the task.

1 Introduction

There is a long history of research in Natural Language Processing (NLP) on the topic of *event extraction* (Grishman, 2019). The vast majority of that research focuses on *sentence-level event extraction* (Ji and Grishman, 2008; Li et al., 2013; Yang et al., 2019; Wadden et al., 2019, i.a.): Given a text and a predefined set of *event types* E, each $e \in E$ with its own set of *arguments* A_e (also known as *roles*), find all role-fillers of these events in the text, identifying for each (1) the *trigger* span t that denotes the occurrence of the event, (2) the event type e, and (3) a span of text s for each instantiated event argument $a \in A_e$ (also called *role fillers*) mentioned **within the sentence containing** t. Prior work treats triggers as integral to accurate event extraction (M'hamdi et al., 2019; Tong et al., 2020; Lin et al., 2022). Specifically, trigger detection is typically the first step of event extraction as it informs event type categorization and provides a lexical anchor for argument extraction. Thus, trigger spans are marked during the data annotation process for sentence-level event extraction (Doddington et al., 2004).

In recent years, however, there has been increasing interest in *document-level event extraction*: Given a document d, a set of event types E, and their associated arguments $A_{e \in E}$, produce one filled event template for each relevant event described in d with event arguments recognized at the *entity level*, i.e., via a single representative mention.[1] Document-level event extraction is important because the vast majority of real-world news and narrative texts describe multiple complex interacting events. Li et al. (2021), for example, find that the arguments of close to 40% of events in their Wikipedia-based dataset appear in sentences other than the event trigger sentence. Unfortunately, the document-level task setting poses significant challenges (Das et al., 2022). Documents can contain zero events or discuss multiple events of the same or similar types. Events and event arguments are likely to be mentioned many times throughout the document, requiring coreference resolution. Arguments can be shared across events. Information about different events can be intermingled.

For example, in Figure 1, ''Heyndal'', and ''Teneda'' are arguments for both the ATTACK and KIDNAPPING events; and their multiple mentions in the text must be recognized as coreferent so that only one mention for each is included in each of

*These authors contributed equally to this work.

[1]Other names for the task include template filling (Jurafsky and Martin, 2000), generalized template extraction (Chen et al., 2023), and N-ary relation extraction (Jain et al., 2020).

Transactions of the Association for Computational Linguistics, vol. 13, pp. 1560–1577, 2025. https://doi.org/10.1162/TACL.a.51
Action Editor: Jonathan Berant. Submission batch: 11/2024; Revision batch: 4/2025; Published 11/2025.
© 2025 Association for Computational Linguistics. Distributed under a CC-BY 4.0 license.

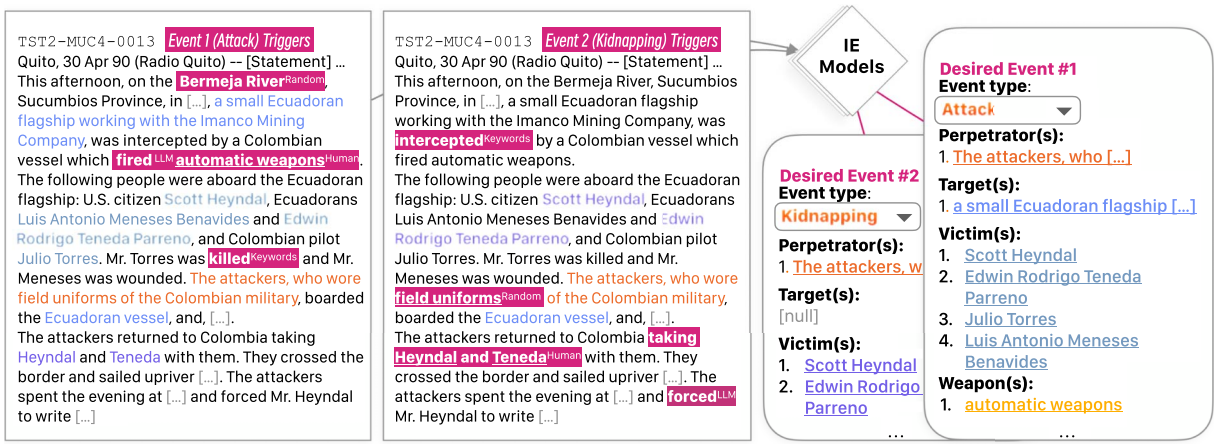

Figure 1: **Document-level event extraction example from *MUC*.** Two events should be extracted from the given text, an ATTACK (doc on left) and a KIDNAPPING (doc on right). Note that only one mention (among many coreferent ones) per event argument is shown. Some event arguments are empty and some might have multiple fillers. The (duplicated) document depicts event triggers obtained from different annotation sources (labeled with pink background and the annotation source); more details in Section 3.

the output events. Additionally, the arguments for each event are spread across multiple sentences.

Note that event trigger-span identification is not required as part of the output; indeed, most datasets for the task do not include trigger-span annotations (Soboroff, 2023; Jain et al., 2020; Pavlick et al., 2016). On one hand, this makes sense: A single event may have multiple triggers (or possibly none), and conversely a particular trigger text span may be associated with more than one event. Also, it is known that event individuation—the problem of distinguishing distinct events—is difficult even for human experts (Gantt et al., 2023). On the other hand, event triggers, which are by definition lexical indicators of an event, should help with event type identification and systems may benefit from extracting them as an intermediate step. In cases where they denote a (verb or noun) predicate for the event, they will serve as an anchor to localize argument extraction (Ebner et al., 2020). Finally, consecutive triggers denoting different event types will be cues to segment the text into event-specific chunks.

In this paper, to the best of our knowledge, we provide the first investigation of the usefulness of triggers for document-level event extraction. As described in Section 3, we treat *Triggers* as ''rationales'' that indicate the presence of an event and guide the system in finding the argument of the event. We analyze whether using triggers improves event extraction performance in multiple end-to-end and pipelined transformer-based (Vaswani et al., 2017) event extraction models for three document-level event extraction datasets. In particular, we measure event extraction performance for a range of event triggers of varying quality (human-annotated, LLM-generated, keyword-based, and random). Additionally, we compare the results to in-context learning prompting baselines. We publicly release all trigger annotations used in our experiments,[2] including our manually annotated trigger-spans for the classic *MUC* dataset.

We find that the usefulness of triggers depends, in part, on characteristics of the extraction task and its data, i.e., the density of relevant events in each document and the degree to which event arguments are localized with respect to the trigger. We also find that lower-quality, automatically identified triggers can be used as an alternative to human annotations. This is good news given the cost of obtaining manual annotations. In addition, we show that performance robustness in the face of degraded trigger quality can be achieved by making descriptive event information available to the extraction model. Perhaps surprisingly, we also find that the mere availability of triggers, even random ones, is important for prompt-based in-context learning approaches to the task. See Section 4.3 for access to our public data and code.

2 Related Work

Document-level Template Extraction. Document-level event extraction was introduced

[2]GitHub link.

in the Message Understanding Conferences (*MUC*) (Sundheim, 1991) as a task to identify occurrences of predefined event types and fill the associated template roles by extracting entity mentions from the text or categorizing aspects of the event. Other event extraction datasets include *RAMS* (Ebner et al., 2020), *CMNEE* (Zhu et al., 2024), *WikiEvents* (Li et al., 2021), and *MAVEN* (Wang et al., 2020). We focus our analysis on *MUC*, *WikiEvents*, and *CMNEE* due to their long-document, multi-event nature and because their limited selection of event types enables clearer analysis of performance differences between our benchmarked systems. Document-level event extraction systems utilize drastically different techniques: Earlier methods relied on pattern matching (Riloff, 1993; Riloff and Jones, 1999), while more modern approaches use deep learning to represent event components such as triggers and entities (Wadden et al., 2019; Wang et al., 2023b; Xu et al., 2021). Most recently, the task has been framed as sequence-to-sequence generation (i.e., TANL [Paolini et al., 2021], GTT [Du et al., 2021] and DEGREE [Hsu et al., 2022]) built upon powerful transformer (Vaswani et al., 2017) language models. Our work focuses primarily on such methods due to the flexibility of sequence-to-sequence frameworks (Paolini et al., 2021) and because we believe future methods will increasingly utilize constantly improving LLMs.

Event Argument Extraction. Argument extraction is a subtask of event extraction that requires identifying fillers for the roles associated with the event given an event type as well as its *trigger* span (Li et al., 2021; Ren et al., 2023; Ma et al., 2022). Argument extractors are commonly employed in event extraction systems that segment the task into trigger/event detection and typing followed by argument extraction. We primarily focus on TANL (Paolini et al., 2021) and GENIE (Li et al., 2021) that model the task via sequence generation, though other methods that rely on careful manipulation and learning of embeddings also exist (Chen et al., 2015; Yang et al., 2019).

3 Approach

Many NLP systems predict *rationales* in conjunction with outputs as a means of explaining them. In this study, we effectively treat *triggers* as rationales: a system is given, or predicts, event triggers that may guide identification of associated arguments. Previous studies have illustrated that NLP systems can leverage rationales in various ways. DeYoung et al. (2020) assert that rationales can function as explanations, helping end users evaluate a system's trustworthiness. Yang et al. (2018b) show that providing rationales as extra supervision enhances a model's capability to perform multi-hop reasoning. Most recently, Chen et al. (2022) demonstrate the potential of rationale-based NLP systems on adversarial inputs. Given these benefits, we investigate whether triggers can offer similar advantages to document-level event extraction systems and prompting-based baselines.

Triggers, as defined in Walker et al. (2005), refer to text spans (verbs, nouns, or adjectives) that most clearly express the occurrence of one or more events. Since trigger annotation is potentially a time-consuming and expensive process and trigger quality is likely to affect their efficacy for document-level event extraction, **we experiment with four sources of triggers**: human-generated, LLM-generated, keyword-based, and randomly selected. Trigger generation and datasets will be explained in Section 4.

Our approach for the evaluation is depicted in Figure 2. Specifically, we study triggers in conjunction with **four recent sequence-to-sequence transformer-based approaches to document-level event extraction**: TANL (Paolini et al., 2021), GTT (Du et al., 2021), DEGREE (Hsu et al., 2022), and GENIE (Li et al., 2021). In addition to adopting different underlying neural architectures, each has input requirements that differ in the amount of task-specific knowledge they expect. The models and input settings are described in Section 5. We consider both **end-to-end** and **pipelined** system architectures (Section 5.2) and provide comparisons to **GPT-4o and GPT-4o-MINI baselines** prompted via in-context learning. Details on the experimental setup can be found in Section 6. Results are presented and analyzed in Section 7.

We specifically investigate the following:

1. Do corpora characteristics, such as the typical number of events and the distribution of event arguments across a document, influence the utility of triggers? For example, triggers might be useful anchors for

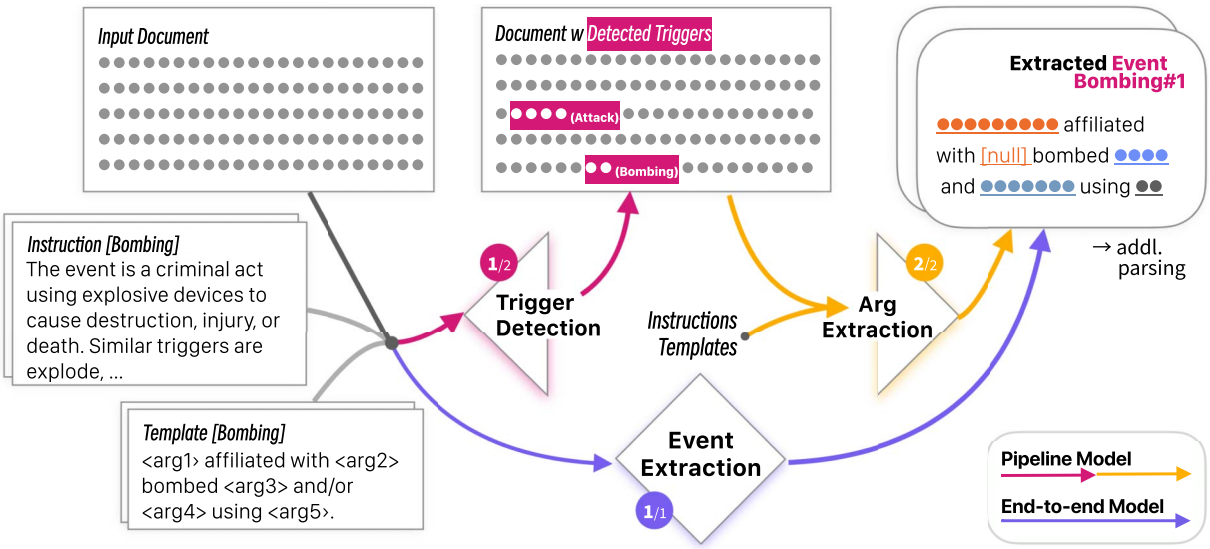

Figure 2: **Approach Overview**. We examine four state-of-the art event extraction systems (e.g., TANL, DEGREE, GTT, GENIE). We run these models in two architectures, when possible, *pipeline* and end-to-end (*E2E*). We also include comparisons to GPT-4o and GPT-4o-MINI prompted via in-context few-shot learning (not shown). Models differ in the task-specific information they provide—document only (TANL, GTT), document+instructions (TANL+GENIE), document+instructions+template (DEGREE).

argument extraction when there are many events with arguments located in close proximity to triggers.

2. Can we view triggers as grounding information for the event extraction task, and if so, could other task/corpus specific information be utilized in place of triggers such that trigger-free methods can match trigger-based methods? For example, can methods that ingest fine-grained event schema details without explicitly identifying triggers during extraction be as performant as methods explicitly learning trigger identification? If grounding information beyond triggers is effective and can be easily constructed, we may be able to bypass expensive human trigger annotation efforts during event extraction dataset creation.
3. To further minimize manual annotation, perhaps only providing in-context examples to a modern LLM can achieve reasonable performance. Alternatively, the many subtasks subsumed in event extraction (event detection, entity detection, entity and event typing) may still be too complex for LLMs. We also explore how triggers impact in-context learning methods.
4. Event extraction methods are generally either pipelined (extract a trigger, then extract a single event's arguments conditioned on a trigger) or end-to-end (the model learns trigger and argument extraction simultaneously). Which performs better? Does the source of the trigger annotations impact performance between the two paradigms? While our benchmark methods by default fit exclusively in one category, slight modifications allow single systems to model both paradigms such that we can center our analyses around the paradigm used.

4 Datasets and Trigger Creation

Similar to previous work, we perform our analyses on the *MUC* (Sundheim, 1991), *WikiEvents* (Li et al., 2021), and *CMNEE* (Zhu et al., 2024) datasets. These are described next, followed by the event trigger generation process.

4.1 Datasets

MUC is a dataset introduced to study event extraction in English news articles reporting Latin American terrorist incidents. The dataset consists of document-template(s) pairs covering six event types. A document can contain no events or multiple events; similarly, an argument role can have no argument entities or multiple entities. Following previous work (Du et al., 2021; Wang et al., 2023a), we restricted event templates to contain

Dataset	# docs	# tokens / doc	# events / doc	% docs w/ $\geq$ 1 event	# events / relevant doc	# args / event	% filled args / event	# tokens between args
MUC								
– train	1300	2089.0	0.86	57%	1.59	2.13	43%	33.9
– val	200	2153.4	0.95	58%	1.65	2.28	46%	44.8
– test	200	1979.8	1.02	63%	1.63	2.34	47%	38.3
WikiEvents								
– train	206	4162.9	15.73	95%	16.71	1.26	28%	1.6
– val	20	3422.2	17.25	100%	17.25	1.39	25%	1.5
– test	20	3785.0	18.25	95%	19.21	1.08	30%	1.9
CMNEE								
– train	6448	192.2	0.95	58%	1.65	5.02	100%	8.8
– val	1162	193.3	0.97	66%	1.47	4.35	100%	8.4
– test	1370	188.2	1.60	80%	2.00	6.12	100%	9.2

Table 1: **Dataset Statistics.** This shows *# tokens / doc* and *# events / doc* on average; and for documents with events, the average *# events / relevant doc* and *# args / event*. It also includes *% filled args / event* and the standard deviation of # tokens between arguments and the midpoint of the text chunk that covers all of the event's arguments (*# tokens between args*).

6 out of the original 24 roles—EVENT-TYPE, PERPETRATOR, PERPETRATOR-ORGANIZATION, TARGET, VICTIM, WEAPON. *MUC* does not provide any *trigger*-span annotations, but we obtain (manually identified) trigger annotations as part of this work. The annotation process is described in Section 4.3.

WikiEvents is a collection of English Wikipedia texts scraped from its ''event'' pages. The dataset consists of document-template(s) pairs covering 33 event types. In contrast to *MUC*, each role is filled by at most one entity. *WikiEvents* provides HUMAN (manually identified) *trigger*-span annotations.

CMNEE is a Chinese dataset of reports of military activity. The dataset consists of document-template(s) pairs covering eight event types; argument roles always have at least one entity filler. Since the dataset is in Chinese, we translated the documents to English using the Google Translate API. We used HTML tags to delimit entity mentions before passing it through the translator to recover their translated correspondences. Because tags are sometimes dropped during translation, some mentions are unrecoverable. We filter out documents with such issues resulting in retaining about 60% of all documents. *CMNEE* provides HUMAN *trigger*-span annotations.

4.2 Dataset Analysis

The characteristics of the datasets vary considerably. Table 1 shows that *WikiEvents* has many events per document compared to *MUC* and *CMNEE*. *MUC* has few events on average per document (many documents have no associated events) whereas all of the *WikiEvents* documents and 62% of the *CMNEE* documents contain at least one event.

Finally, the *localization* of the argument mentions for event descriptions in the documents of each dataset varies tremendously—for *WikiEvents*, the arguments for each event are 1–2 tokens apart[3] on average; for *CMNEE*, arguments for each event are generally within 8–9 tokens of each other; arguments for *MUC* events, on the other hand, are very spread out, spanning about 40 tokens.

Ideally the documents in our datasets would both be long and have event arguments widely dispersed throughout the texts, but such a dataset does not exist at the time of writing. Our selection of *MUC*, *WikiEvents*, and *CMNEE* as a group nevertheless still allows us to perform meaningful analyses of these characteristics.

4.3 Trigger Creation

We investigated four sources of trigger annotation: HUMAN, LLM, KEYWORD, and RANDOM and provide details on how each is generated below. Examples of each trigger source are provided in Figure 1.

[3]We used nltk.word_tokenize to tokenize documents. For each event, we identify the index of the each of its arguments' first token, then compute the standard deviation of these indices.

We will make publicly available all trigger sets employed in our experiments.[4]

HUMAN Triggers HUMAN trigger spans are obtained by hiring annotators to identify the most informative trigger for each event in the document. They are the most expensive trigger type to obtain. *WikiEvents* and *CMNEE* provide one human-annotated trigger per event; *MUC* does not. As a result, we hired 9 annotators to identify all plausible trigger spans for each document-event pair. 57% of the document-template pairs in the dataset were annotated by at least two annotators. Average α-Krippendorff and F-Kappa agreement scores across annotators of 0.530 and 0.453, respectively, indicate a moderate level of agreement. The annotation guidelines are provided in Appendix A.

To obtain a single trigger for each document-event pair, we take the union of all annotated triggers for the event and select the first, i.e., earliest, span among them. This typically produced better extraction scores than using the trigger span most centrally located among the event's arguments or other plausible trigger selection heuristics.

LLM Triggers Recent work has demonstrated that GPT performs well on some information extraction tasks (Wang et al., 2023c; Wadhwa et al., 2023). We therefore hypothesized LLM-generated triggers might be an acceptable proxy for HUMAN annotated triggers. We use zero-shot prompting of GPT-4-TURBO to identify and rank multiple trigger-span(s) for each event and select the highest ranking span.[5]

KEYWORD Triggers KEYWORD triggers are obtained by identifying a set of event-specific keywords at the corpus level and applying them heuristically for each document. For each event type $e \in E$, we identified documents that contained at least one event of type *e* to form a list L_e of words associated with *e*. We manually filtered each L_e producing lists $\hat{L}_e$ that contain only the words that are plausible triggers for *e* (i.e., ''exploded'' or ''bombed'' for a bombing event), thus forming a list of keywords for each event type. We then use the entries from $\hat{L}_e$ to annotate all events of type *e*.

[4]GitHub link.

[5]The prompts used are part of the GitHub release.

RANDOM Triggers We also collected triggers by randomly selecting a span of 1–3 words from the document. They are obviously of poor quality and are used as a control to identify whether the mere presence of triggers in the system input is helpful.

5 Models

We study the performance of four state-of-the-art sequence-to-sequence style systems for document-level event extraction: TANL (Paolini et al., 2021), GTT (Du et al., 2021), DEGREE (Hsu et al., 2022), and GENIE (Li et al., 2021) and compare them against prompt-based baselines using GPT-4O and GPT-4O-MINI.

5.1 Model Overview

GPT-4O Baselines We prompt GPT-4O and GPT-4O-MINI to extract event templates in a JSON format giving it six in-context examples. We extract all events for all event types using a single query. We use event descriptions and templates as instructions. An example of the prompt can be found in Appendix B.

TANL (Paolini et al., 2021) is a pipeline-based model that takes only the document as input. As shown in Figure 2, TANL fine-tunes a T5-large transformer (Raffel et al., 2020) to complete event extraction in two phases: 1) event detection, and 2) argument extraction. First, the document is passed through T5 to tag/delimit a trigger for each detected event (the entire document is produced as output with event type tags associated with recognized triggers). Then, each identified {trigger, event type} pair (tagged/delimited in the input document) is passed through T5 once to identify all event arguments and assign them a role type.

GTT (Du et al., 2021) is an end-to-end model that takes a document and an ordered list of the possible event types and generates filled templates for all event types in a single pass through BERT (Devlin et al., 2019). Templates are linearized to sequences with special argument and event delimiters.

DEGREE (Hsu et al., 2022) is an end-to-end model that takes a document, a natural language template for one event type *E* and extra information about events of type *E* (i.e., type definitions, example triggers); it generates filled

natural language templates (that include identified triggers) for all detected events for type E in a single pass. The model is called once for each event type of interest for each document. For this reason, training is augmented with multiple negative examples (i.e., document-(empty)template pairs for event types that do not appear in the text).

GenIE (Li et al., 2021) was originally proposed for argument extraction. Specifically, a document is provided with the trigger delimited by special tokens and a natural language template (similar to that of Degree) corresponding to the event type to extract arguments for. The model output is the natural language template filled with arguments. We use GenIE for the full event extraction task by using TANL as the trigger extractor and GenIE as the argument extractor.

5.2 Model Type: Pipeline vs End-to-End (e2e), No Trigger

We explore both pipeline and end-to-end model architectures and modify TANL and Degree to run in both settings. As shown in Figure 2, a pipeline system consists of a trigger extraction and an argument extraction component. The trigger extraction phase identifies all event triggers and their corresponding event types. The argument extraction phase then operates on each {trigger, event type} pair separately. This design allows argument extraction to anchor in an informative part of the text, i.e., in the vicinity of a known trigger. End-to-End systems, on the other hand, jointly extract triggers, event types, and arguments, allowing the two extraction tasks to inform one another, thus alleviating cascading errors that can occur in a pipelined setting.

We next describe changes made to GTT, Degree and TANL to run in various settings.

Pipeline We modify Degree to fit the pipeline architecture. We train distinct models for trigger and argument extraction. Specifically, the trigger extraction input and output templates only contain slots for a trigger and event type; the argument extraction input templates have an additional slot for a trigger (identified during trigger extraction) and the output templates only remove the slot for the trigger, retaining only slots for each argument type. The model is called once to fill natural language templates for all identified triggers (trigger extraction) and once for each {identified trigger, event type} pair to identify arguments (argument extraction). We did not use GTT for pipelined experiments. By default, TANL, TANL + GenIE are pipelined models.

No Triggers We modify all models to run without the use of triggers. All *no trigger* experiments were conducted in an e2e manner. For TANL, the raw text is provided and the model tags arguments with an event type and number and the argument role (e.g., ''perpetrating individual for attack 1''). A single span can have multiple such tags if it is a shared entity mention for multiple events. The spans associated with each unique {event type, number} pair are processed into fillers for single events. This scheme only requires one pass of T5 for full event extraction. For Degree, we only remove the slot for the trigger in input/output templates. By default, GTT is an e2e system that does not use triggers.

End-to-End (with triggers) For TANL, the event type and number and the argument role tag is supplemented with a tag for the trigger (e.g., ''perpetrating individual for attack 1 | trigger = killed''). There is still only one pass for full event extraction. For GTT, a slot is added at the start of the output event templates for a trigger. By default, Degree is an e2e system using triggers.

6 Experimental Setup

In our study of the utility of triggers on document-level event extraction, we tested the performance of fine-tuned (i.e., pipeline and end-to-end) and prompt-based approaches on the three datasets described in Section 4. We consider the prompt-based approaches as a baseline and compare them to fine-tuning the four state-of-the-art models, described in Section 5; each model was tested in a pipeline and an end-to-end manner as applicable. We train on $1 \times$ A100 and $4 \times$ 1080Ti using the default learning rate suggested for each model using the first 10% of the steps as warmup. The reported score on the test data is obtained using the system from the best performing epoch using the validation set.

Due to the varying sizes of the datasets, with some very small, we perform 5-cross validation on the test data by sampling 80% of the data for each evaluation run. Each model is trained once and evaluated once on each of the 5 validation splits. We then report average and standard deviation scores across the 5 splits. Our evaluation focuses

on Micro-F1—combined F1 score on both event type detection and argument role extraction. Since the datasets are multi-event, we have to align the predicted and gold templates. We follow the matching algorithm of Das et al. (2022) to align and score the predicted templates.[6]

7 Results and Discussion

In this section, we discuss our experimental results shown in Table 2 (follows Conclusions). We first analyze some high level performance trends independent of trigger usage before commenting on how triggers should be used depending on various system and corpora specifications.

7.1 General Observations for Document-Level Event Extraction

Result #1 – Prompt-based in-context learning approaches for LLMs do not outperform smaller-scale fine-tuned models. As shown in Table 2, the best prompt-based systems achieve 43.8% (with human triggers), 38.3% (keyword), and 66.7% (human) F1 while the best fine-tuning methods achieve 59.7% (DEGREE$_{e2e}$ no-trigger), 47.1% (TANL$_{pipeline}$ human), and 78.9% F1 (DEGREE$_{e2e}$ no trigger) on *MUC*, *WikiEvents* and *CMNEE*, respectively.

Result #2 – End-to-end systems perform similarly to pipeline architectures for the same model types. As shown in Figure 3, e2e models trained with triggers perform roughly on par or better than their pipelined counterparts with five explainable exceptions. All exceptions involve the TANL system applied to *WikiEvents* and *CMNEE*. We hypothesize that since event arguments are localized around the trigger for both datasets (see Table 1), dedicating a separate pipeline component towards trigger extraction, which helps disambiguate events and anchor argument extraction, is particularly beneficial. (This is explored in more detail in Section 7.2 where the propensity of DEGREE to ignore triggers is also discussed.) This is important because we can generally build e2e systems that are faster to train and quicker during inference than pipelined systems without sacrificing performance.

[6]We modify their proposed scorer by improving its compute time (using the Hungarian Matching algorithm [Kuhn, 1955]) and by removing span-based scoring as not all models predict spans.

7.2 Observations on the Importance of *triggers* for Document-Level Event Extraction

Result #3 – The presence of triggers, even random ones, is generally *needed* for prompt-based approaches. Across both GPT models and multiple datasets, we observe that in-context prompting with triggers (from *any* trigger source) regularly outperforms the associated no-trigger variations. This is especially drastic for GPT-4o on *WikiEvents*, where we see a 5 point F1 gain after incorporating keyword triggers in the in-context examples as shown in Table 2. Interestingly, even adding random spans as triggers boosts performance by 2 points for GPT-4o. Upon manual inspection we observe that GPT-4o actually identifies legitimate trigger-like phrases in spite of the nonsense triggers in the in-context examples. This follows findings in recent work (Min et al., 2022) and suggests that learning the concept of a trigger, even without proper examples, meaningfully supports extraction.

Result #4 – Triggers are *not* needed for corpora that have few events to extract in each document. We notice for *CMNEE* and (especially) *MUC* that systems trained without triggers perform comparably to systems trained with triggers (Figure 4a) (i.e., most points fall below the $x = y$ diagonal). The key commonality between the two datasets is their relatively low number of events per document (compared to *WikiEvents*). We hypothesize that event disambiguation is therefore easier in *MUC* and *CMNEE*; therefore the use of triggers (to anchor argument extraction) for these two datasets is less useful. Exceptions to this occur when applying TANL to *CMNEE* where fine-tuning in the pipelined setting performs better with triggers (e.g., 54.3%, 59.6%, and 64.0% F1) than without (43.8% F1). We believe this may be due to localized events for *CMNEE* as noted in Result #2 and Result #5.

Result #5 – Higher quality triggers are preferred for corpora that have many *localized* events to extract in each document. This result is almost a negation of the previous claim, except for the additional conditions of *localized* events and *higher quality* triggers. Aside from the typical number of events per document, *WikiEvents* also differs from *MUC* and *CMNEE* in how

Experiment	*MUC*	*WikiEvents*	*CMNEE*
		no *Triggers*	
GPT-4o-MINI$_{prompting}$	$.434_{\pm.01}$	$.325_{\pm.02}$	$.565_{\pm.00}$
GPT-4o$_{prompting}$	$.391_{\pm.01}$	$.339_{\pm.02}$	$.635_{\pm.00}$
TANL$_{e2e}$	$.559_{\pm.01}$	$.361_{\pm.02}$	$.438_{\pm.00}$
DEGREE$_{e2e}$	$.597_{\pm.00}$	$.248_{\pm.01}$	$.789_{\pm.00}$
GTT$_{e2e}$	$.454_{\pm.01}$	—	$.318_{\pm.00}$
		Human *Triggers*	
GPT-4o-MINI$_{prompting}$	$.438_{\pm.02}$	$.364_{\pm.02}$	$.564_{\pm.00}$
GPT-4o$_{prompting}$	$.405_{\pm.01}$	$.379_{\pm.02}$	$.667_{\pm.00}$
TANL$_{pipeline}$	$.460_{\pm.01}$	$.471_{\pm.02}$	$.596_{\pm.01}$
DEGREE$_{pipeline}$	$.524_{\pm.01}$	$.357_{\pm.01}$	$.723_{\pm.00}$
TANL+GENIE$_{pipeline}$	$.464_{\pm.01}$	$.382_{\pm.01}$	$.599_{\pm.00}$
TANL$_{e2e}$	$.530_{\pm.01}$	$.337_{\pm.02}$	$.410_{\pm.00}$
DEGREE$_{e2e}$	$.566_{\pm.01}$	$.339_{\pm.02}$	$.782_{\pm.00}$
GTT$_{e2e}$	$.429_{\pm.01}$	—	$.326_{\pm.00}$
		LLM *Triggers*	
GPT-4o-MINI$_{prompting}$	$.416_{\pm.02}$	$.368_{\pm.02}$	$.580_{\pm.00}$
GPT-4o$_{prompting}$	$.414_{\pm.01}$	$.383_{\pm.02}$	$.659_{\pm.00}$
TANL$_{pipeline}$	$.305_{\pm.01}$	$.207_{\pm.01}$	$.543_{\pm.00}$
DEGREE$_{pipeline}$	$.567_{\pm.01}$	$.335_{\pm.01}$	$.715_{\pm.00}$
TANL+GENIE$_{pipeline}$	$.490_{\pm.01}$	$.388_{\pm.01}$	$.558_{\pm.00}$
TANL$_{e2e}$	$.507_{\pm.00}$	$.267_{\pm.02}$	$.440_{\pm.00}$
DEGREE$_{e2e}$	$.557_{\pm.00}$	$.313_{\pm.02}$	$.703_{\pm.00}$
GTT$_{e2e}$	$.439_{\pm.02}$	—	$.332_{\pm.00}$
		Keyword *Triggers*	
GPT-4o-MINI$_{prompting}$	$.429_{\pm.02}$	$.361_{\pm.02}$	$.582_{\pm.00}$
GPT-4o$_{prompting}$	$.420_{\pm.01}$	$.389_{\pm.02}$	$.659_{\pm.00}$
TANL$_{pipeline}$	$.525_{\pm.02}$	$.306_{\pm.02}$	$.640_{\pm.00}$
DEGREE$_{pipeline}$	$.532_{\pm.01}$	$.241_{\pm.01}$	$.532_{\pm.00}$
TANL+GENIE$_{pipeline}$	$.470_{\pm.01}$	$.316_{\pm.01}$	$.628_{\pm.00}$
TANL$_{e2e}$	$.495_{\pm.00}$	$.292_{\pm.02}$	$.443_{\pm.00}$
DEGREE$_{e2e}$	$.523_{\pm.01}$	$.273_{\pm.01}$	$.601_{\pm.00}$
GTT$_{e2e}$	$.455_{\pm.01}$	—	$.324_{\pm.00}$
		Random *Triggers*	
GPT-4o-MINI$_{prompting}$	$.412_{\pm.01}$	$.369_{\pm.02}$	$.533_{\pm.01}$
GPT-4o$_{prompting}$	$.402_{\pm.00}$	$.359_{\pm.02}$	$.533_{\pm.01}$
TANL$_{pipeline}$	$.116_{\pm.01}$	$.059_{\pm.01}$	$.223_{\pm.00}$
DEGREE$_{pipeline}$	$.538_{\pm.01}$	$.255_{\pm.01}$	$.589_{\pm.00}$
TANL+GENIE$_{pipeline}$	$.322_{\pm.01}$	$.120_{\pm.01}$	$.487_{\pm.00}$
TANL$_{e2e}$	$.473_{\pm.01}$	$.343_{\pm.02}$	$.394_{\pm.00}$
DEGREE$_{e2e}$	$.514_{\pm.01}$	$.270_{\pm.02}$	$.595_{\pm.01}$
GTT$_{e2e}$	$.304_{\pm.01}$	—	$.310_{\pm.00}$

Table 2: **Event Extraction Experiment.** Average F1-micro scores on 5-fold cross validation for the prompting baselines, end-to-end, and pipeline systems on *MUC*, *WikiEvents*, and *CMNEE*. The colored cells indicate the best performing model, for a given dataset, with the darkest color indicating the highest F1 score.

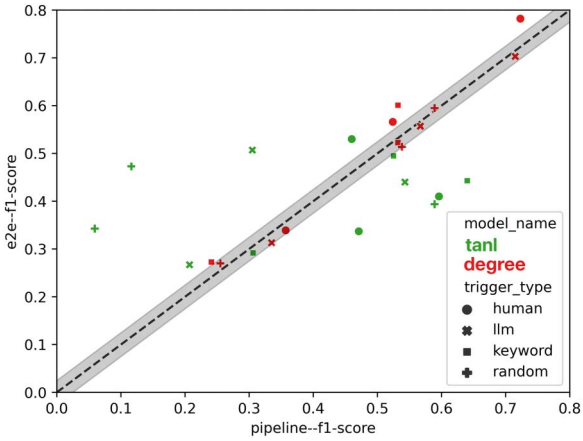

Figure 3: **E2E vs. Pipeline.** F1-score of each e2e experiment vs. its pipeline version. If the data point is above $y = x$ (dashed line), then the e2e architecture outperformed the pipeline architecture for that given (model, trigger type, dataset) setting. (Datasets cannot be distinguished in the figure.)

closely grouped together event arguments are (see Table 1). We unfortunately do not have a dataset with documents with many events with dispersed arguments to explore the many event per document setting more generally. In Figure 4b and Table 2, we see that e2e and pipelined systems for *WikiEvents* fine-tuned with human triggers perform better than their fine-tuned counterparts without triggers. Our conclusion regarding the number of events per document is similar to Result #2: Many events require careful event disambiguation and identifying relevant triggers serves this purpose well. Regarding event localization, the trigger's position is a strong clue to inform systems where to look for associated arguments. In fact, we see that pipelined TANL and TANL + GenIE, the only systems that have access to the exact location of triggers within a document prior to argument extraction, are the best models for *WikiEvents* (see Table 2).

Result #6 – Lower quality trigger annotations that are less expensive to obtain *can* be used as an alternative to HUMAN triggers. We see in Figure 5a that trigger quality does not significantly impact overall performance for fine-tuning methods. This suggests that models can learn textual cues beyond triggers during extraction. An exception to this is the pipelined TANL system applied to *WikiEvents*, though we believe this relates to the discussion for Result #5.

Result #7 – Robustness with respect to degraded trigger quality can be achieved by giving the event extraction system access to additional event information (e.g., templates, instructions). Evidence for this can be seen by examining the performance of Degree, which does not exhibit any significant swings in overall performance when comparing across trigger annotation sources. For *MUC* and *CMNEE*, Degree actually performs best (and outperforms all other systems) in the *no triggers* setting (59.7% and 78.9% F1, respectively). We believe that this robustness arises from the background information it requires as part of its input, including a natural language statement of the template for the target event type, a definition of the event type, a short list of likely triggers, short definitions of expected event arguments. As additional evidence for this claim, one can observe the *random triggers* results for $\text{TANL}_{pipeline}$ in Table 2, which are significantly worse than its e2e variation (11.6% versus 47.3%, 5.9% versus 34.3%, 22.3% versus 39.4% F1 for *MUC*, *WikiEvents*, and *CMNEE*, respectively). This difference, however, is remedied by providing a natural language template as input for argument extraction, as in the case of TANL + GenIE: The addition results in +20.6, +6.1, and +26.4 F1 % points for *MUC*, *WikiEvents*, and *CMNEE*, respectively. These results suggest that additional relevant inputs such as natural language templates that relate argument types and descriptions of events benefit fine-tuning methods to the extent of possibly overcoming systematic flaws such as conditioning extraction on noise. Overall, we see the trend outlined in Figure 5b: TANL exhibits very poor performance, adding event template information (TANL + GenIE) improves system performance, and the inclusion of event-specific instructions (Degree) allows further robustness when training on random triggers.

8 Limitations

Although many current state of the art approaches employ some sequence-to-sequence language modeling approach, there are alternative models that we did not investigate. For example, IterX (Chen et al., 2023) formulates event extraction as a Markov decision process while Procnet builds a graph over entity representations to be parsed into events (Wang et al., 2023b). Additionally,

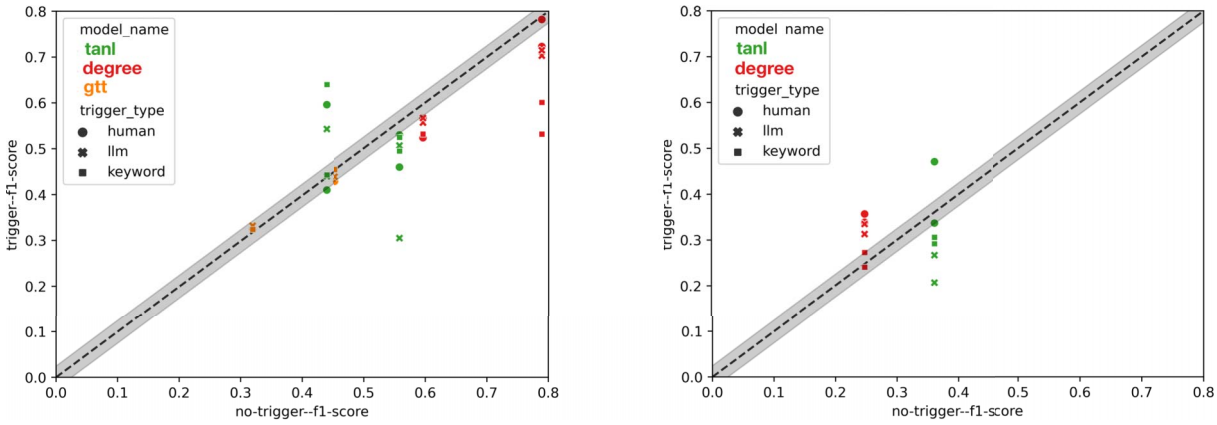

(a) *MUC & CMNEE*

(b) *WikiEvents*

Figure 4: **Triggers vs. no-Triggers.** F1-score of each fine-tuning experiment with triggers (annotated by humans, LLMs, or the keyword heuristic) against its e2e version without triggers. A data point above $y = x$ (dashed line) implies that fine-tuning with triggers performed better for that (model, architecture, dataset).

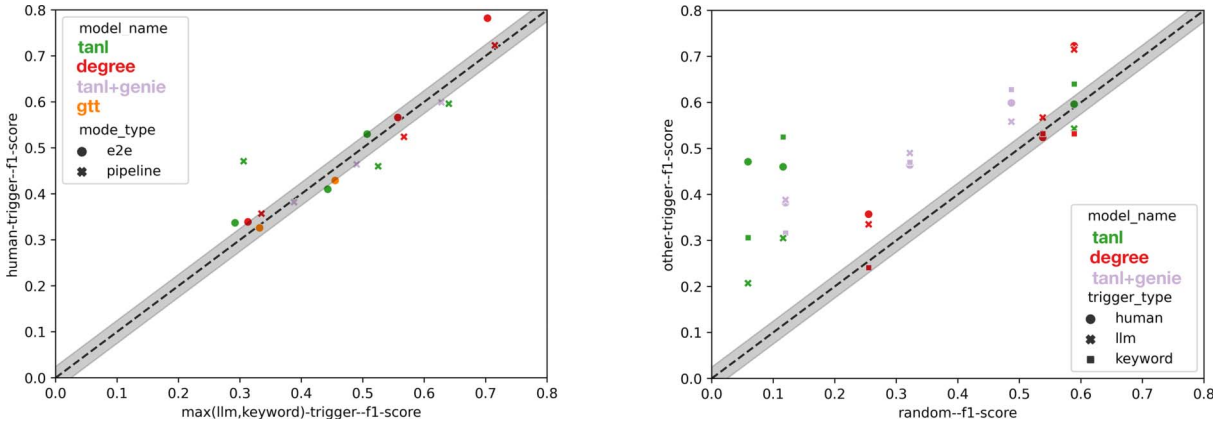

(a) Against/ HUMAN

(b) Against/ RANDOM

Figure 5: **Trigger vs. Trigger.** This plots the F1-scores of experiments fine-tuning with one trigger source against another. In 5a, a point above $y = x$ (dashed line) implies that using human annotations performed better than using any other annotation source, while in 5b and using either human, LLM, or keyword annotations performed better than random annotations.

some more elaborate prompting-based approaches such as CODE4STRUCT (Wang et al., 2022) warrant testing given the rapidly developing capabilities of LLMs.

9 Conclusions and Future Work

In this work, we explore the use of triggers in document-level event extraction. Although typically regarded as an integral part of sentence-level event extraction, we showed human annotated triggers are only necessary given certain dataset and modeling parameters (many localized events in a document, providing relevant conditioning information such as natural language templates).

Future work could investigate a wider scope of models, as discussed just above in Section 8 and/or incorporate additional datasets where available. Existing datasets such as *CMNEE* and *DCFEE* (Yang et al., 2018a) are in Chinese (therefore likely require translation before use) while others such as *RAMS* (Ebner et al., 2020) and *Doc-EE* (Tong et al., 2022) do not exhibit traits of document level datasets (*RAMS* has 1 event per document and ‘‘document’’ are restricted to 5 sentence windows surrounding the trigger while *Doc-EE* has only one event per document). Efforts creating new document level event extraction datasets would significantly expand work like ours

that can inform future event extraction system design choices.

References

Howard Chen, Jacqueline He, Karthik Narasimhan, and Danqi Chen. 2022. Can rationalization improve robustness? In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3792–3805, Seattle, United States. Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.naacl-main.278

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. https://doi.org/10.3115/v1/P15-1017

Yunmo Chen, William Gantt, Weiwei Gu, Tongfei Chen, Aaron White, and Benjamin Van Durme. 2023. Iterative document-level information extraction via imitation learning. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1858–1874, Dubrovnik, Croatia. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.eacl-main.136

Aliva Das, Xinya Du, Barry Wang, Kejian Shi, Jiayuan Gu, Thomas Porter, and Claire Cardie. 2022. Automatic error analysis for document-level information extraction. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3960–3975, Dublin, Ireland. Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.acl-long.274

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. 2020. ERASER: A benchmark to evaluate rationalized NLP models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4443–4458, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.408

George R. Doddington, Alexis Mitchell, Mark A. Przybocki, Lance A. Ramshaw, Stephanie M. Strassel, and Ralph M. Weischedel. 2004. The automatic content extraction (ace) program-tasks, data, and evaluation. In *Lrec*, volume 2, pages 837–840. Lisbon.

Xinya Du, Alexander Rush, and Claire Cardie. 2021. Template filling with generative transformers. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 909–914, 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. https://doi.org/10.18653/v1/2020.acl-main.718

William Gantt, Reno Kriz, Yunmo Chen, Siddharth Vashishtha, and Aaron White. 2023. On event individuation for document-level information extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12938–12958, Singapore. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.findings-emnlp.862

Ralph Grishman. 2019. Twenty-five years of information extraction. *Natural Language Engineering*, 25(6):677–692. https://doi.org/10.1017/S1351324919000512

I-Hung Hsu, Kuan-Hao Huang, Elizabeth Boschee, Scott Miller, Prem Natarajan, Kai-Wei Chang, and Nanyun Peng. 2022. DEGREE: A data-efficient generation-based event extraction model. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1890–1908, Seattle, United States. Association for Computational Linguistics.

Sarthak Jain, Madeleine van Zuylen, Hannaneh Hajishirzi, and Iz Beltagy. 2020. SciREX: A challenge dataset for document-level information extraction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7506–7516, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.670

Heng Ji and Ralph Grishman. 2008. Refining event extraction through cross-document inference. In *Proceedings of ACL-08: HLT*, pages 254–262, Columbus, Ohio. Association for Computational Linguistics.

Daniel Jurafsky and James H. Martin. 2000. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition* (3rd edition draft).

H. W. Kuhn. 1955. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1–2):83–97. https://doi.org/10.1002/nav.3800020109

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. https://doi.org/10.18653/v1/2021.naacl-main.69

Po-Wei Lin, Shang-Yu Su, and Yun-Nung Chen. 2022. TREND: Trigger-enhanced relation-extraction network for dialogues. In *Proceedings of the 23rd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 623–629, Edinburgh, UK. Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.sigdial-1.58

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.

Meryem M'hamdi, Marjorie Freedman, and Jonathan May. 2019. Contextualized cross-lingual event trigger extraction with minimal resources. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 656–665, Hong Kong, China. Association for Computational Linguistics. https://doi.org/10.18653/v1/K19-1061

Sewon Min, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the role of demonstrations: What makes in-context learning work? https://doi.org/10.18653/v1/2022.emnlp-main.759

Giovanni Paolini, Ben Athiwaratkun, Jason Krone, Jie Ma, Alessandro Achille, Rishita Anubhai, Cicero Nogueira dos Santos, Bing Xiang, and Stefano Soatto. 2021. Structured prediction as translation between augmented natural languages. In *9th International Conference on Learning Representations, ICLR 2021*.

Ellie Pavlick, Heng Ji, Xiaoman Pan, and Chris Callison-Burch. 2016. The gun violence database: A new task and data set for NLP. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1018–1024, Austin, Texas. Association for Computational Linguistics. https://doi.org/10.18653/v1/D16-1106

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(1).

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.

Ellen Riloff. 1993. Automatically constructing a dictionary for information extraction tasks. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, AAAI'93, pages 811–816. AAAI Press.

Ellen Riloff and Rosie Jones. 1999. Learning dictionaries for information extraction by multi-level bootstrapping. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence and the Eleventh Innovative Applications of Artificial Intelligence Conference Innovative Applications of Artificial Intelligence*, AAAI '99/IAAI '99, pages 474–479, USA. American Association for Artificial Intelligence.

Ian Soboroff. 2023. The better cross-language datasets. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '23, pages 3047–3053, New York, NY, USA. Association for Computing Machinery. https://doi.org/10.1145/3539618.3591910

Beth M. Sundheim. 1991. Overview of the third message understanding evaluation and conference. In *Third Message Understanding Conference (MUC-3): Proceedings of a Conference Held in San Diego, California, May 21–23, 1991*. https://doi.org/10.3115/1071958.1071960

Meihan Tong, Bin Xu, Shuai Wang, Yixin Cao, Lei Hou, Juanzi Li, and Jun Xie. 2020. Improving event detection via open-domain trigger knowledge. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5887–5897, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.522

MeiHan Tong, Bin Xu, Shuai Wang, Meihuan Han, Yixin Cao, Jiangqi Zhu, Siyu Chen, Lei Hou, and Juanzi Li. 2022. DocEE: A large-scale and fine-grained benchmark for document-level event extraction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3970–3982, Seattle, United States. Association for Computational Linguistics. https://doi.org/10.18653/v1/2022.naacl-main.291

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

David Wadden, Ulme Wennberg, Yi Luan, and Hannaneh Hajishirzi. 2019. Entity, relation, and event extraction with contextualized span representations. *ArXiv*, abs/1909.03546. https://doi.org/10.18653/v1/D19-1585

Somin Wadhwa, Silvio Amir, and Byron Wallace. 2023. Revisiting relation extraction in the era of large language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15566–15589, Toronto, Canada. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.acl-long.868, PubMed: 37674787

Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2005. Ace 2005 multilingual training corpus.

Barry Wang, Xinya Du, and Claire Cardie. 2023a. Probing representations for document-level event extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12675–12683, Singapore. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.findings-emnlp.844

Xinyu Wang, Lin Gui, and Yulan He. 2023b. Document-level multi-event extraction with event proxy nodes and hausdorff distance minimization. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10118–10133, Toronto, Canada. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.acl-long.563

Xingyao Wang, Sha Li, and Heng Ji. 2022. Code4struct: Code generation for few-shot structured prediction from natural language. *arXiv preprint arXiv:2210.12810*. https://doi.org/10.18653/v1/2023.acl-long.202

Xingyao Wang, Sha Li, and Heng Ji. 2023c. Code4Struct: Code generation for few-shot event structure prediction. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3640–3663, Toronto, Canada. Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.acl-long.202

Xiaozhi Wang, Ziqi Wang, Xu Han, Wangyi Jiang, Rong Han, Zhiyuan Liu, Juanzi Li, Peng Li, Yankai Lin, and Jie Zhou. 2020. MAVEN: A massive general domain event detection dataset. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1652–1671, Online. Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.emnlp-main.129

Runxin Xu, Tianyu Liu, Lei Li, and Baobao Chang. 2021. Document-level event extraction via heterogeneous graph-based interaction model with a tracker. 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 3533–3546, Online. Association for Computational Linguistics.

Hang Yang, Yubo Chen, Kang Liu, Yang Xiao, and Jun Zhao. 2018a. DCFEE: A document-level Chinese financial event extraction system based on automatically labeled training data. In *Proceedings of ACL 2018, System Demonstrations*, pages 50–55, Melbourne, Australia. Association for Computational Linguistics. https://doi.org/10.18653/v1/P18-4009

Sen Yang, Dawei Feng, Linbo Qiao, Zhigang Kan, and Dongsheng Li. 2019. Exploring pre-trained language models for event extraction and generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5284–5294, Florence, Italy. Association for Computational Linguistics. https://doi.org/10.18653/v1/P19-1522

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018b. HotpotQA: A dataset for diverse, explainable multi-hop question answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. https://doi.org/10.18653/v1/D18-1259

Mengna Zhu, Zijie Xu, Kaisheng Zeng, Kaiming Xiao, Mao Wang, Wenjun Ke, and Hongbin Huang. 2024. CMNEE: A large-scale document-level event extraction dataset based on open-source Chinese military news. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3367–3379, Torino, Italia. ELRA and ICCL.

A Annotation Guidelines

Similar to Walker et al. (2005), we define a trigger as a word that indicate the occurrence of an event. A trigger can be any part of speech—e.g., verb, noun or adjective; but the majority are verbs and nouns. Common triggers for the different event-types in the *MUC* dataset include ‘‘attack‘‘, ‘‘bomb’’, ‘‘blew up‘‘, ‘‘killed’’, etc. In Table 3, we show the additional annotation guidelines we curated to annotate trigger-span(s) for *MUC*.

	Rule Description	Example
1	*Keep trigger spans as long as possible, but divide spans if uninformative words fall between keywords*	dynamite charge$_{\text{trigger1}}$ today exploded$_{\text{trigger2}}$ $\cdots$
2	*Do not include statistics in triggers*	attack$_{\text{trigger1}}$ in which 20 persons were killed$_{\text{trigger2}}$ and approximately 100 were injured$_{\text{trigger3}}$ $\cdots$
3	*Do no include adjectives in triggers*	threatening$_{\text{trigger1}}$ them with hand grenades$_{\text{trigger2}}$ $\cdots$
4	*Do not include recipients/direct objects in triggers*	murder$_{\text{trigger1}}$ of Jecar Neghme $\cdots$
5	*Do not annotate words such as ‘‘body’’, ‘‘bodies’’, ‘‘injury’’, ‘‘injured’’, ‘‘wounded’’ unless there are no other triggers present in the paragraph*	
6	*Do not include words that indicate verb tense such as ‘‘will’’, ‘‘is’’, ‘‘are’’, ‘‘have’’*	

Table 3: **Annotation Guidelines**. We describe the additional rules and definitions we used to trigger in conjecture with Walker et al. (2005). We give and example to each rule and indicate which tokens are selected as triggers.

B Prompting

In this section we show the prompts there were used for our experiments. In the following table, we show the prompt that was used to perform document-level event extraction for *MUC*. For different trigger sources, we would change the trigger used and remove the entry completely when running the no *Triggers* experiment.

```
----------- System Prompt ----------

You are journalist reading news articles and identifying the main events that the
↪ article is reporting...


-------- In Context Examples -------

Perform event extraction on the provided document and identify key details (such as
↪ event-type, involved organizations or individuals, target of the event, victim
↪ of the event, and weapon used) and output it in a json format.
```

```
Document:
   san salvador, 12 nov 89 (domestic service) -- [communique] [salvadoran armed
  ↪ forces press committee] [text] the salvadoran armed forces inform the people
  ↪ of el salvador that fmln [farabundo marti national liberation front] ...

Events:
'''[
 {
 ''trigger'': ''vandals'',
 ''event type'': ''attack'',
 ''perpetrator individual'': [''groups of vandals''],
 ''perpetrator organization'': [''fmln''],
 ''target'': [''shopping centers''],
 ''victim'': [],
 ''weapon'': []
 }
]'''

...

Perform event extraction on the provided document ...

Document:
 lima, 10 jan 90 (efe) -- [text] the national police reported today that over
 ↪ 15,000 people have been arrested in lima over the past few hours ...

Events:

----------- GPT response -----------
'''[
 {
 ''trigger'': ''was riddled with bullets'',
 ''event_type'': ''attack'',
 ''perpetrating individual'': [''gerardo olivos silva''],
```

```
 ''perpetrating organization'': [''shining path''],
 ''target'': [''enrique lopez albujar trint''],
 ''victim'': [''enrique lopez albujar trint''],
 ''weapon'': [''firearm'']
 }
 ...
]'''
```