

# CMNEE: A Large-Scale Document-Level Event Extraction Dataset based on Open-Source Chinese Military News

Mengna Zhu<sup>1</sup>, Zijie Xu<sup>2</sup>, Kaisheng Zeng<sup>3</sup>, Kaiming Xiao<sup>1</sup>, Mao Wang<sup>1</sup>,  
Wenjun Ke<sup>2</sup>, Hongbin Huang<sup>1\*</sup>

<sup>1</sup> Laboratory for Big Data and Decision, National University of Defense Technology

<sup>2</sup> School of Computer Science and Engineering, Southeast University

<sup>3</sup> Computer Science and Technology, Tsinghua University

{zhumengna16, kmxiao, wangmao, hbhuang}@nudt.edu.cn

{zijiexu, kewenjun}@seu.edu.cn

{zks19}@mails.tsinghua.edu.cn

## Abstract

Extracting structured event knowledge, including event triggers and corresponding arguments, from military texts is fundamental to many applications, such as intelligence analysis and decision assistance. However, event extraction in the military field faces the data scarcity problem, which impedes the research of event extraction models in this domain. To alleviate this problem, we propose **CMNEE**, a large-scale, document-level open-source **Chinese Military News Event Extraction** dataset. It contains 17,000 documents and 29,223 events, which are all manually annotated based on a pre-defined schema for the military domain including 8 event types and 11 argument role types. We designed a two-stage, multi-turns annotation strategy to ensure the quality of CMNEE and reproduced several state-of-the-art event extraction models with a systematic evaluation. The experimental results on CMNEE fall shorter than those on other domain datasets obviously, which demonstrates that event extraction for military domain poses unique challenges and requires further research efforts. Our code and data can be obtained from <https://github.com/Mzzzhu/CMNEE>.

**Keywords:** Corpus, Information Extraction, Information Retrieval, Knowledge Discovery/Representation

## 1. Introduction

Event extraction refers to the extraction of structured information from unstructured text, which is typically separated into two subtasks: event detection and event argument extraction (Ahn, 2006; Yang et al., 2019; Xu et al., 2023; Wang et al., 2021b; Shi et al., 2023; Yang et al., 2023; Wan et al., 2023; Liu et al., 2020; Chu et al., 2023; Peng et al., 2023; Wang et al., 2023). Current research predominantly focuses on the general news or financial domains, with only relatively fewer studies for military domain. Nevertheless, the importance of military event extraction is undeniable.

The military domain encompasses numerous documents which contain rich event information. Extraction of military events from documents is crucial to downstream applications such as intelligence analysis (Santucci, 2022; Bang, 2016; Freedman, 1983; Ivanov, 2011), decision making assistance (Skryabina et al., 2020), and strategic planning (Schrodt, 2012; Sankar, 2023; Lyu, 2022). Figure 1 provides an example of event extraction based on military news. This military news document contains six events. It's necessary to recognize triggers in the document and determine their corresponding event types, as well as relevant arguments and corresponding roles. For instance,

in the 6th sentence, S06, it's imperative to recognize the trigger *attacking* and the event type *Conflict*. Additionally, the "Subject" *British Naval*, the "Object" *targets in Yugoslavia*, and "Date" *1999* of the event should also be identified.

Existing research on event extraction largely relies on deep neural network-based models, which have a strong dependence on data size and quality. Consequently, the lack of high-quality training data significantly limits their performances. Most of these models use sentence-level text for analysis, and while sentence-level event extraction research is relatively mature (Gao, 2021; Hsu et al., 2021; Huang et al., 2023a), this technology struggles when directly applied to document-level event extraction tasks. This is due to with the increased length of the text, event arguments often scattered across different sentences. Moreover, most existing event extraction datasets are oriented towards general (Li et al., 2020, 2021a) or financial domains (Han et al., 2022a; Zheng et al., 2019). The performance of methods applied to datasets in different languages and domains often varies significantly. Currently, military event data extraction primarily relies on human labor, leading to issues such as low efficiency, inconsistent standards, and incomplete information. These challenges make it difficult to support practical applications based on events. Therefore, constructing a document-

---

\* Corresponding author

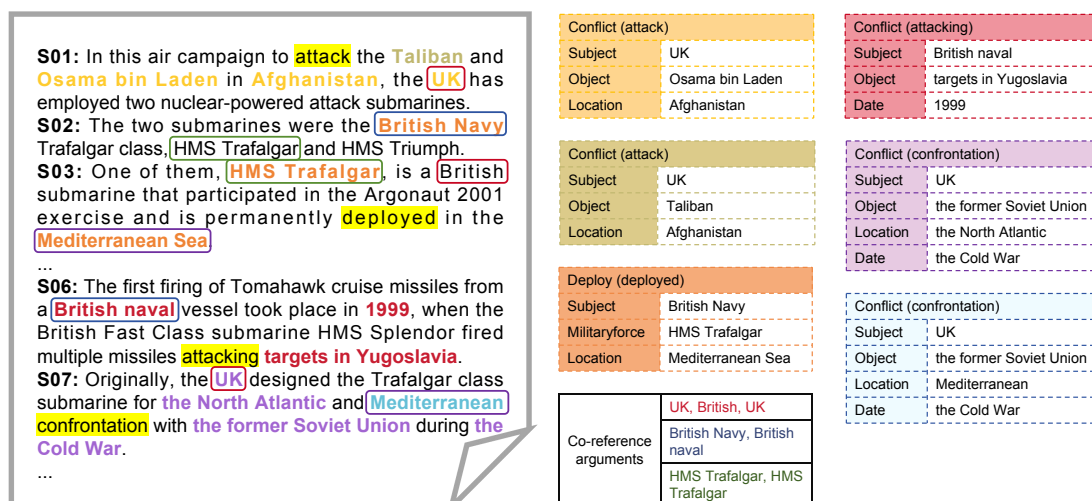


Figure 1: An example from CMNEE. Each document in CMNEE is annotated with event types and involved event arguments. This document contains 6 events. We used different colors to distinguish arguments of these events and colored rectangular boxes to mark different co-reference arguments.

level event extraction dataset specifically for the military domain holds substantial practical significance and application value.

However, due to the confidentiality and sensitivity of military domain data and the difficulty of obtaining event schema, there is a long-term gap in the datasets used for the event extraction task in this domain, which is the main reason that the development of military event extraction is seriously lagging behind mainstream research.

To fill this gap and alleviate relevant problems, we proposed a large-scale document-level open-source **Chinese Military News Event Extraction** dataset (**CMNEE**), which involved corpus from authoritative websites such as Huanqiu<sup>1</sup>, China Military Online<sup>2</sup>, Sina Military<sup>3</sup> and Baidu Encyclopedia<sup>4</sup>. The acquired military news text itself is reliable. In the process of data annotation based on the news, firstly, the event types that may be contained in the document are pre-labeled by matching with constructed triggers dictionary, combined with domain experts' knowledge and the analysis of related news to form the event schema. Based on the event schema and referring to the data annotation criteria, the annotation is completed in a two-stage multi-turns manner. The annotation requirements are to make manual judgments on

the results of pre-labelling, annotate the triggers and event types, event arguments and corresponding argument roles. In addition, considering that one of the reasons for the poor performance of event argument extraction is that the identified arguments are co-reference arguments rather than those labeled in the dataset (Lu et al., 2023), the co-reference arguments involved in the text are also labeled. Quality indicators are designed to evaluate CMNEE. Annotations and data annotation criteria are continuously corrected until all indicators reach the specified thresholds.

CMNEE is currently the only dataset for the document-level event extraction task in the military domain. To better evaluate CMNEE, we implemented several advanced models and conducted a systematic evaluation. The experimental results demonstrate that CMNEE has some unique challenges that need to be overcome and deserves more efforts to study. We hope that CMNEE can facilitate relevant research, and attract more attention into event extraction for military domain.

## 2. Construction of CMNEE

Our main goal is to construct a large-scale dataset that facilitates the development of event extraction in the military domain. During the construction of CMNEE, event trigger, event type, event arguments, argument roles and co-reference arguments need to be labelled. The main process of CMNEE construction is shown in Figure

<sup>1</sup><https://www.huanqiu.com/>

<sup>2</sup><http://www.81.cn/wzsy/index.html>

<sup>3</sup><https://mil.news.sina.com.cn/>

<sup>4</sup><https://baike.baidu.com/>

2. Firstly, we crawled military news texts from websites with authority and then pre-processed data with the help of existing mature tools. Then manually annotated them in groups after completing automatic event types pre-labeling based on predefined candidate triggers dictionary. Annotation is constantly corrected by using a two-stage multi-turns annotation approach until all the quality indicators reached the specified thresholds. We finally obtained 17,000 pieces of high-quality labelled documents data.

### 2.1. Data Collection and Preprocessing

To ensure the quality of the dataset, our data sources are military news websites with authority, relevance, and domain value, such as Huanqiu, China Military, Sina Military, and Baidu Encyclopedia. We used Python's Requests library and BeautifulSoup library to crawl nearly 100,000 news data from above websites. Based on the 7 predefined event types in the sentence-level military event detection dataset MNEE (Huang et al., 2023b), we used ChatGPT<sup>5</sup> to generate synonyms and related words for these event types, which were manually filtered to form a keyword dictionary, and used words in the dictionary to match with news text, according to the number of matching words to filter out 30,000 pieces of data with high relevance. Then the Python Pandas library was used to filter the missing samples and remove the duplicated samples, and regular expressions were used to remove the garbled codes and the meaningless information such as the html of the web pages. Finally 17000 high quality news documents were retained.

### 2.2. Event Schema Construction and Pre-labeling

Following MNEE, combining the analysis of collected news and referring to domain experts' knowledge, 8 event types other than non-event are defined, including Experiment, Manoeuvre, Deploy, Support, Accident, Exhibit, Conflict, and Injure. We constructed a trigger dictionary and pre-labeled event types that each sample may contain, and used the pre-labeled results to conduct preliminary statistical analysis to obtain event type proportion, and referred to domain experts' knowledge in combination with the real situation to assess the reasonableness of the distribution of event types.

Then, event argument roles are determined by analyzing documents corresponding to each event type obtained after pre-labelling. The final event schema can be seen in Table 1.

### 2.3. Human Annotation and Quality Evaluation

After completing data preprocessing, the pre-labeled event type information is used as a reference to assist human annotation and labelers need to judge them first of all. In the process of data labelling, two-stage multi-turns iterative approach is adopted to continuously revise the annotation so as to enhance the quality of the dataset.

**Stage 1:** In order to improve the efficiency of labelling while keeping the quality as high as possible, we divide labelers into two groups, 1) Normal Group, members of which are employees of the labelling company engaged in the labelling work for a long time, with partial background knowledge. and 2) Criteria Group, members of which are domain experts with relatively complete professional knowledge and involved in the formulation of data annotation criteria. All the data are shuffled and processed in batches, personnel of these two groups respectively labelled documents by batch at the same time. When the acquired documents have been labelled, the annotation should be checked, and labels of disagreement should be discussed and amended in time. Meanwhile, data annotation criteria is continuously refined.

**Stage 2:** The Normal Group was divided into several smaller groups to examine annotation results of Stage 1 in multiple turns by group separately, and labels were corrected according to the final data annotation criteria. At the same time, Criteria Groups randomly select samples to examine and discuss corrections within the group.

In order to assess the quality of the dataset, four types of quality indicators, namely, Event Type Accuracy, Event Type Recall, Event Argument Recall, and Co-reference Argument Recall, are used with reference to MNEE (Huang et al., 2023b), definitions can be seen in Table 2. Detailed descriptions are as follows.  $Cor_{es}$  represents the number of samples whose events are all correctly labelled.  $S$  represents the total number of selected samples.  $Cor_e$ ,  $Cor_a$ ,  $Cor_c$  represent the number of correctly labelled events, arguments, co-reference arguments in the sample respectively.  $Act_e$ ,  $Act_a$ ,  $Act_c$  represent the number of actually existed events, arguments, co-reference arguments in the sample respectively. Correctly labelled events refer to events whose triggers with their offsets and corresponding event types are all correctly labelled. Correctly labelled event arguments refer to arguments that have the correct texts, roles, and offsets under the condition that their event types are correct. Actual events, arguments, co-reference arguments refer to those that are expected to be labelled in accordance with the data annotation criteria.

After annotation is completed, the Criteria Group

---

<sup>5</sup><https://chat.openai.com/>

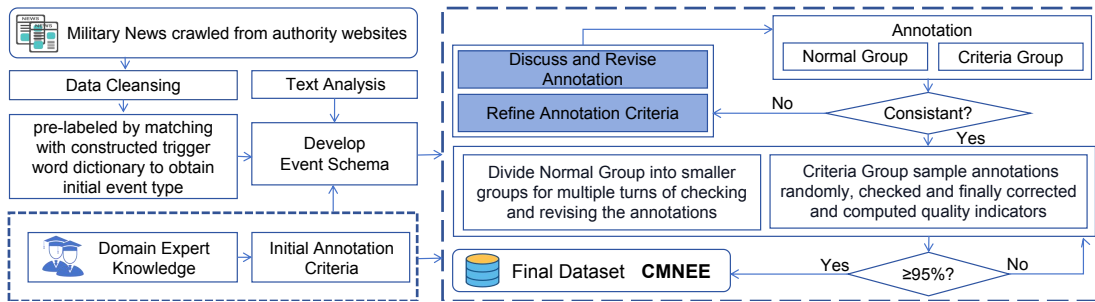


Figure 2: Main process of CMNEE construction. Firstly, crawled military news texts from websites with authority and pre-processed data. Then manually annotated them in groups after completing pre-labeling. Annotation is constantly corrected by using a two-stage multi-turns annotation approach until all the indicators reached the specified thresholds.

Event type	Definition	Argument Roles
Experiment	Countries, organizations, institutions, or companies verify whether the equipment is qualified based on existing standards, including but not limited to flights, launching, testing, etc.	Subject, Equipment, Date, Location
Manoeuvre	Exercises conducted in the course of campaigns and battles, in the context of the situation envisaged, and are the highest and most centralized form of military training.	Subject, Content, Date, Area
Deploy	The movement of military forces, i.e., personnel or equipment, within a country or organization.	Subject, Militaryforce, Date, Location
Support	Subject provides material help or relief actions to the object.	Subject, Object, Materials, Date
Accident	Accidents that happen unexpectedly.	Subject, Result, Date, Location
Exhibit	Countries, organizations, institutions, companies, etc. exhibit or publicize equipment, products, etc. through airshows or other forms.	Subject, Equipment, Date, Location
Conflict	Violence acts, such as attack, that causes damage or injury, or a conflict or confrontation between two parties, such as protest or condemnation.	Subject, Object, Date, Location
Injure	Person entity suffered physical injury.	Subject, Quantity, Date, Location

Table 1: Event schema of CMNEE

checks a certain percentage of randomly selected samples many turns. If all indicators of the samples are higher than 95%, the dataset is qualified, or else the Normal Group will continue to correct the annotations. Final quality indicator results are 95.1%, 96.4%, 96.8%, and 97.2% respectively.

### 3. Data Analysis of CMNEE

In this section, we analyze various aspects of CMNEE to provide a deep understanding of the dataset. Overall, CMNEE defines 8 event types and 11 argument roles. It contains 17,000 valid documents sourced from authoritative websites, 29,223 non-empty events and 93,348 event arguments. The distribution of data sources is shown

in Figure 3, which shows that the vast majority of documents are from reliable sources. On average, there are 330 tokens, 6.7 sentences, 1.8 events and 6.6 event arguments per document in CMNEE. The longest number of document tokens in the CMNEE is 496 and the largest number of sentences is 17. Tokens here refer to Chinese characters.

We compare CMNEE with some existing widely-used EE related datasets and show the main statistics in Table 3, including sentence-level event extraction datasets, ACE 2005 (Christopher et al.), MAVEN (Wang et al., 2020), DuEE (Li et al., 2020), MNEE (Huang et al., 2023b) and document-level event extraction datasets, RAMS (Ebner et al., 2020), WikiEvents (Li et al., 2021b), Duce-fin (Han et al., 2022b), Chfi-



Indicators	Abbreviations	Calculation Methods
Event Type Accuracy	$TypeAcc$	$(Cor_{es}/S) \times 100\%$
Event Type Recall	$TypeR$	$(Cor_e/Act_e) \times 100\%$
Event Argument Recall	$ArgR$	$(Cor_a/Act_a) \times 100\%$
Co-reference Argument Recall	$Co-refR$	$(Cor_c/Act_c) \times 100\%$

Table 2: Definitions of Quality Indicators




Figure 3: Data Sources

nAnn (Zheng et al., 2019), DocEE (Tong et al., 2022). CMNEE is comparable with existing datasets in the number of documents and events. Despite its relatively small number of event types and argument roles, CMNEE basically covers the main event types in the military domain.

### 3.1. Event Types Distribution

Figure 4 shows the event types distribution of CMNEE by their instance numbers. We can observe that the inherent data imbalance problem also exists in CMNEE. It is obvious that CMNEE is long-tailed and consistent with the characteristics of the data in the actual scenario. It is beneficial to use the models trained based on CMNEE for developing real-world applications. Although CMNEE is long-tailed, as CMNEE is large-scale and event types are not particularly numerous, all event types have more than nearly 500 instances respectively. Conflict has the lowest number of instances as 477. The ‘‘Experiment’’ event type has the largest number of instances in the dataset, which is 13,027. CMNEE is relatively less affected by data sparsity.

### 3.2. Events Analysis

Since the actual news may usually contain multiple events, during the annotation process, we try to annotate all the events information in the document as much as possible. Figure 5 displays the distribution of the number of events contained in each instance within the dataset. One instance corresponds to a single document. The inner circle numbers in the donut chart represent the quantity of events, while the outer circle numbers in-

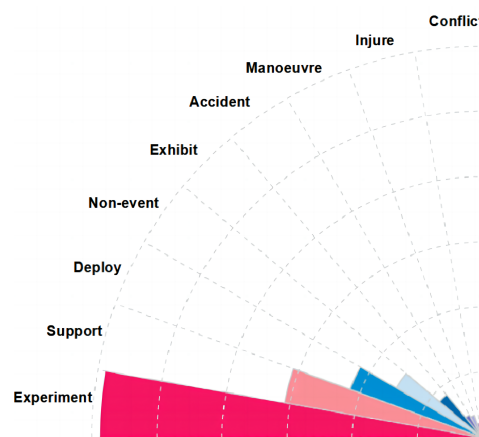


Figure 4: Event Type Distribution

dicating the respective proportions. For example, the green segment denotes the proportion of documents that contain only one event. From this figure, it is evident that nearly half of the documents include two or more events, and approximately a quarter contain three or more events. So it is hopeful that CMNEE will facilitate the research on multi-event extraction and promote the extraction more independently, completely and accurately.




Figure 5: Multi-events Distribution

In addition, due to the concise and condensed nature of military texts, there are a large number of overlapping events in CMNEE, and the percentage of instances containing overlapping events is 42%. Overlapping events mean that the same word acts as a trigger for different events in the same instance, the same argument plays different roles in multiple events or the same argument plays differ-

Level	Dataset	Domain	Language	Docs	EvTyps	ArgRs	Events	ED	EAE
Sentence	ACE2005	General	English	599	33	35	4,090	✓	✓
	MAVEN	General	English	4,480	168	-	111,611	✓	✗
	DuEE	Financial	Chinese	11,224	65	121	19,640	✓	✓
	MNEE	Military	Chinese	13,000	8	10	6,997	✓	✓
Document	RAMS	General	English	3,993	139	65	9,124	✓	✓
	WikiEvents	General	English	246	50	59	3,951	✓	✓
	Duee-fin	Financial	Chinese	11,700	13	92	11,031	✓	✓
	ChfinAnn	Financial	Chinese	32,040	5	24	47,824	✗	✓
	DocEE	General	English	27,485	59	356	27,485	✓	✓
	<b>CMNEE</b>	Military	Chinese	17,000	8	11	29,223	✓	✓

Table 3: Statistics of CMNEE compared with existing widely-used EE datasets

ent roles in the same event (Sheng et al., 2021). For example, in the sentence “The submarine was delivered on January 9 after completing trials.”, the “submarine” is both the material delivered and the equipment for the trials. Currently, the FewFC dataset is the main dataset used for the study of overlapping events (Sheng et al., 2021; Cao et al., 2022), but the size of this dataset is relatively small, in which only 18% of the instances contain overlapping events, and its annotation is based on sentences. Overlapping events easily lead to confusion and omission of event information, which increases the difficulty of extracting event information completely.

### 3.3. Event Arguments Analysis

Since military text involves many proper nouns such as weapons and equipment, the proportion of long arguments is relatively high, and the proportion of arguments with more than 10 Chinese characters is nearly 17%, which increases the difficulty of argument extraction.

Considering when there are multiple events in a document, it is likely that the same argument belongs to more than one event at the same time, so we defined *shared arguments* as arguments that belong to different events at the same time with the same text or in the same co-reference arguments sub-list for ease of analysis. There are 19,353 shared arguments contained in CMNEE, which is about one-fifth of the common arguments. Shared arguments can reveal relationships between events. For example, if two events share a *Date* argument, it often indicates that these events occurred simultaneously. Accurately extracting all events which shared arguments based on CMNEE is difficult, but it is very important. So more attention should be paid into it. It is also an effective way to improve the performance.

Subset	Docs	Events	Args
Train	12,000	19,427	62,591
Dev	2,000	3,236	10,275
Test	3,000	6,560	20,842

Table 4: Statistics of splitting CMNEE

## 4. Experiments on CMNEE

### 4.1. Experimental setting

**Benchmark setting** Greedy algorithm is used to split CMNEE into training, development, and test sets to make event types of these subsets as identically distributed as possible. Table 4 lists basic statistics of the split result.

**Models** We selected several advanced baselines for evaluation and they can be categorized into three types. The first category includes models that are designed based on Chinese corpora and do not use triggers information, to experimentally verify the necessity of annotating trigger information for CMNEE. Category 1: 1) **DCFEE-O** & **DCFEE-M** (Yang et al., 2018) proposed a key-event detection method to complete event arguments extraction by filling event table according to key-event mention and surrounding sentences. DCFEE-O only produces one event record from one key-event sentence. DCFEE-M tries to get multiple possible argument combinations. 2) **GreedyDec** & **Doc2EDAG** (Zheng et al., 2019) transformed the event extraction task into a directed cyclic graph path extension task based on entity information. GreedyDec only fills one event table entry greedily. 3) **DEPPN** (Yang et al., 2021a) used a multi-granularity non-autoregressive decoder to generate events in parallel based on document-aware representations. The second category comprises

the current mainstream and advanced Event Extraction models to validate unique challenges of CMNEE compared to other datasets. Category 2: 4) **BERT+CRF** is a sequence labeling model. 5) **EEQA** (Du and Cardie, 2020) converted the event extraction task into the natural question answering task. 6) **TEXT2EVENT** (Lu et al., 2021) is a sequence-to-structure generation paradigm that can directly extract events from the text in an end-to-end manner. The third category consists of models that complete the Event Argument Extraction task based on golden triggers, in order to highlight the importance of resolving error propagation issues on CMNEE in Event Extraction tasks. Category 3: 7) **PAIE** (Ma et al., 2022) constructed prompts considering argument interaction to extract event arguments.

**Evaluation** Following the widely-used setting, we report the micro Precision, Recall, and F-1 scores for event detection and event argument extraction as our evaluation metrics (Wang et al., 2020; Tong et al., 2022).

## 4.2. Overall Experimental Results

Overall experimental results can be seen in Tabel 5. From these results we could find that: 1) Almost all models have relatively low performance on CMNEE, which suggests that CMNEE is challenging and that event extraction for the military domain deserves further exploration. 2) Models not using trigger information are significantly less effective than those using trigger information. Due to the specificity of the military domain, military events are always centered around several specific types of entities. In the event schema, the overlap of argument roles is higher compared to other domains, and it is difficult to accurately distinguish different events. Therefore, incorporating trigger information for event types is conducive to better accomplishing the military event extraction task. Among them, PAIE used golden trigger information, and the event argument extraction effect is much better than other models. 3) It is worth noting that BERT+CRF performs well on CMNEE due to the advantage of CRF in modelling multi-event correlation, as can be seen from Figure 5, CMNEE has a high percentage of multi-events, and CRF is helpful in multi-event information extraction.

## 4.3. Analyzes on Data Size

In this section, we selected several models with better performance to analyze the benefits of a larger data scale. We randomly choose different proportions of documents from the CMNEE training set and compare the model performances (Trg-F1 for event detection and Arg-F1 for event argu-

ment extraction) trained with different data sizes. Results can be seen in Figure 6. We can observe that CMNEE can sufficiently train the models.

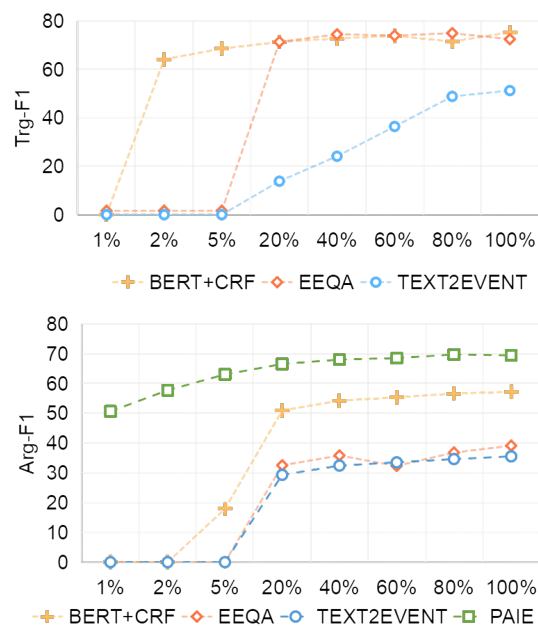


Figure 6: Model Performance (Trg-F1 & Arg-F1) change along with different training set data size.

## 4.4. Analyzes on Co-reference Arguments

Since the same argument contained in the event in the document may appear several times with the same text and different offsets, or in the form of abbreviation, code name, etc., i.e., there exists co-reference arguments. The extraction of co-reference arguments can also express the event information. For example, in the sentence “*The Seal nuclear submarine is a third-generation Russian submarine, construction of which began in 1991, and in 2008 the Seal underwent navigational tests in the Sea of Japan.*”, extracting “The Seal nuclear submarine” at the beginning of the sentence and “the seal” in the middle of the sentence both can be used as the Equipment argument of the Experiment event. Judging the extraction result only based on the events list does not meet the practical requirements. If co-reference argument information is taken into consideration, the performance of event argument extraction will be obviously improved as can be seen in Figure 7.

## 4.5. Error Analysis

In order to better understand the difficulties of CMNEE, We randomly selected 50 samples to analyze errors that occurred in event detection and

Models	Event Detection			Event Argument Extraction		
	Trg-P	Trg-R	Trg-F1	Arg-P	Arg-R	Arg-F1
DCFEE-O	-	-	-	30.3	22.3	25.7
DCFEE-M	-	-	-	26.4	22.0	24.0
GreedyDec	-	-	-	39.4	19.9	26.4
Doc2EDAG	-	-	-	54.3	23.9	33.2
DEPPN	-	-	-	38.2	35.0	36.5
BERT+CRF	73.1	77.7	75.3	63.1	52.3	57.2
EEQA	65.8	80.5	72.4	39.0	39.1	39.0
TEXT2EVENT	30.1	60.6	40.2	31.3	41.3	35.5
PAIE	-	-	-	72.0	67.0	69.4

Table 5: Overall event detection and event argument extraction performance on CMNEE

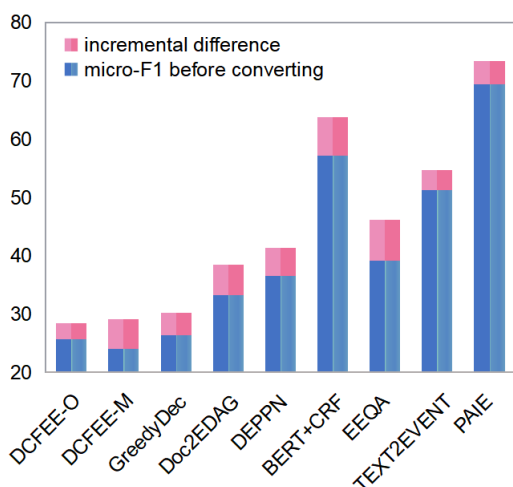


Figure 7: Converted metrics after considering co-reference arguments.

event argument extraction and categorized these errors into 3 main groups.

**Identification mistakes.** It is the most common error accounting for nearly 50% that negative instances are misclassified to positive types (FP) or positive instances are misclassified to N/A (FN). This suggests that understanding event semantic information based on complex text still requires a great deal of effort.

**Majority bias.** Due to the inherent unbalanced characteristics of CMNEE, when dealing with complex texts, instance tends to be classified into types of higher frequency of occurrence. This error type accounts for about 20% of all errors. Since we did not apply data augmentation or balancing during dataset construction and maintain the realworld distribution in CMNEE, it is expected that models that can cope with this problem and accomplish the extraction of fewer sample events more efficiently will be developed.

**Extraction boundary.** In this situation, extracted argument tokens are more or less than the actual labelled argument. For example, the correct equipment argument of Deploy is the *SAD anti-missile system*, but the extracted argument is the localized text *anti-missile system*, which omits key information. Another example is that location argument should be extracted is *South Korea* rather than *located in South Korea*, which introduces irrelevant information. How to better determine the extraction boundaries so that the extracted information is concise and effective deserves further research. This is one of the major reasons for the unsatisfactory performance of the event extraction model, with errors accounting for about 32%.

## 5. Related Work

Since event extraction is one of the important tasks in natural language processing and is significant in practical applications, various datasets are constructed. As shown in Table 3, the datasets constructed for different subtasks and different text granularity have their own characteristics. The most widely used dataset is ACE2005 (Christopher et al.), which defines 8 event types, 33 event subtypes and 30 argument roles, and contains 599 documents and 4,480 events. MAVEN (Wang et al., 2020) is the largest event detection dataset, which defines 168 event types and labels 19,640 events. For the document-level event extraction task, the commonly used datasets are RAMS (Ebner et al., 2020) and WikiEvents (Li et al., 2021a), with a smaller number of documents, 9124 and 246, respectively. ChfinAnn (Zheng et al., 2019) is constructed using distant supervision to assist in the construction, with a sizable scale, but it does not contain event trigger information, and can only be used for the event argument extraction task.

Most of the current relevant datasets are oriented to general domains or financial domains. There are also some datasets oriented



to other specific domains, such as biomedical domain (Pyysalo et al., 2013), literary domain (Sims et al., 2019), terrorist attack events (Grishman and Sundheim, 1996), breaking news (Fu et al., 2010), etc. To the best of our knowledge, there are currently no publicly available English datasets specifically dedicated to event extraction tasks in the military domain. However, some existing datasets do include certain military events. For instance, the ACE2005 dataset contains the “Conflict.Attack” category, while the MAVEN dataset includes categories like “Attack” and “Defending”. Additionally, the MUC-4 dataset encompasses information on terrorist attack events. The volume of military event-related data in these datasets is relatively limited and not easily distinguishable for specialized research purposes.

Recently, there have been some efforts using ChatGPT for data annotation in areas like Named Entity Recognition and Relation Extraction, as indicated in references (Ding et al., 2022; Goel et al., 2023; Zhang et al., 2023). However, due to the complex data structure involved, event extraction tasks prove more challenging for annotation. The potential reduction in manpower, time and costs were also limited. Considering CMNEE’s status as a pioneering dataset in the field, we aimed for utmost authenticity, reliability, and quality, ultimately opting for manual annotation. Nevertheless, we admit that the use of new techniques remains a promising avenue for future research and deserves more attention.

## 6. Conclusion

To fill the gap between military event extraction and mainstream researches, we proposed a large-scale document-level event extraction dataset, CMNEE, which was developed based on open-source Chinese military news. CMNEE is an important complementary for existing event extraction datasets, and is helpful to advance the development of event extraction task research, particularly facilitate progress in the military domain. Compared to existing datasets, CMNEE is oriented to the military domain, featuring a higher proportion of overlapping events and longer arguments due to the characteristics of military texts, thereby increasing the difficulty of extraction. Additionally, it includes annotations for co-reference arguments, making the assessment of extraction results more rational. Experiments demonstrate that CMNEE is challenging and event extraction in military domain remains an open issue.

## 7. Limitations

We have to admit that CMNEE has some limitations. 1) Limited event types and roles. Due to the confidentiality and sensitivity of the military domain, there is limited publicly available information, and the development of event schema and annotation criteria relies heavily on domain expert knowledge, which is not easily available. We finally defined 8 event types that are more common and currently suitable for public research. In order to improve the diversity, applicability, and impact of CMNEE, we expected that more authoritative experts could be invited to increase and refine the event types. 2) Choose of language and annotation methodology. Due to the large size of CMNEE and it is annotated based on documents, the annotated information is also relatively complex, in the existing attempts to annotate with the help of large language models, the results are not satisfactory and the cost that can be saved is very limited. As the first dataset in the military domain, we want to ensure the quality of the dataset as much as possible, and finally choose pure manual annotation. Although we designed a two-stage multi-turns annotation approach to improve the efficiency while guaranteeing the quality, the cost of time, manpower, and expense for the data annotation is still very high, so we only annotate the Chinese corpus at the moment. Because Chinese event extraction is also a relatively large and active community. But in the future, expanding CMNEE to other languages and further try to annotate with new techniques are both necessary.

## 8. Acknowledgements

We thank the anonymous reviewers for their insightful comments and suggestions, researchers in the event extraction field for their excellent studies, data annotators for their hard work and all our team members for their kind support. This work was supported by National Science Foundation of China (Grant Nos.62376057) and the Start-up Research Fund of Southeast University (RF1028623234). All opinions are of the authors and do not reflect the view of sponsors. In addition, because CMNEE was used to support a competition<sup>6</sup> on the joint extraction of specific domain multi-events information held by *Laboratory for Big Data and Decision, National University of Defense Technology*, we also want to thank all experts of this competition for their professional opinions and participants for their entries which better validated the usefulness of CMNEE.

---

<sup>6</sup><https://www.datafountain.cn/competitions/987>

## 9. Bibliographical References

- David Ahn. 2006. [The stages of event extraction](#). In *Proceedings of the Workshop on Annotating and Reasoning about Time and Events*, pages 1–8, Sydney, Australia. Association for Computational Linguistics.
- Martin Bang. 2016. Pitfalls in military quantitative intelligence analysis: Incident reporting in a low intensity conflict. *Intelligence and National Security*, pages 49–73.
- Hu Cao, Jingye Li, Fangfang Su, Fei Li, Hao Fei, Shengqiong Wu, Bobo Li, Liang Zhao, and Donghong Ji. 2022. [OneEE: A one-stage framework for fast overlapping and nested event extraction](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1953–1964, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Walker Christopher, Strassel Stephanie, Medero Julie, and Maeda Kazuaki. Ace 2005 multilingual training corpus. <https://catalog.ldc.upenn.edu/LDC2006T06>.
- Kongbin Chu, Wenzhong Yang, Fuyuan Wei, and Jiangtao Shi. 2023. [Chinese document-level emergency event extraction dataset and corresponding methods](#).
- Bosheng Ding, Chengwei Qin, Linlin Liu, Lidong Bing, Shafiq R. Joty, and Boyang Li. 2022. [Is GPT-3 a good data annotator?](#) *CoRR*, abs/2212.10450.
- 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.
- 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.
- Lawrence Freedman. 1983. The soviet estimate: Us intelligence analysis and russian military strength. *International Affairs*, pages 261–262.
- Jian-feng Fu, Wei Liu, Zong-tian Liu, and Sha-sha Zhu. 2010. [A study of chinese event taggability](#). In *2010 Second International Conference on Communication Software and Networks*, pages 400–404.
- Hongbo Gao. 2021. *Military Event Extraction for Encyclopedia Data*. Ph.D. thesis, Dalian University of Technology.
- Akshay Goel, Almog Gueta, Omry Gilon, Chang Liu, Sofia Erell, Lan Huong Nguyen, Xiaohong Hao, Bolous Jaber, Shashir Reddy, Rupesh Kartha, Jean Steiner, Itay Laish, and Amir Feder. 2023. [Llms accelerate annotation for medical information extraction](#).
- 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*.
- Cuiyun Han, Jinchuan Zhang, Xinyu Li, Guojin Xu, Weihua Peng, and Zengfeng Zeng. 2022a. [Duee-fin: A large-scale dataset for document-level event extraction](#). In *Natural Language Processing and Chinese Computing*, pages 172–183, Cham. Springer International Publishing.
- Cuiyun Han, Jinchuan Zhang, Xinyu Li, Guojin Xu, Weihua Peng, and Zengfeng Zeng. 2022b. [Duee-fin: A large-scale dataset for document-level event extraction](#). In *Natural Language Processing and Chinese Computing*, pages 172–183, Cham. Springer International Publishing.
- I-Hung Hsu, Kuan-Hao Huang, Elizabeth Boschee, Scott Miller, Prem Natarajan, Kai-Wei Chang, and Nanyun Peng. 2021. [Event extraction as natural language generation](#). *CoRR*, abs/2108.12724.
- Heyan Huang, Xiao Liu, Ge Shi, and Qian Liu. 2023a. [Event extraction with dynamic prefix tuning and relevance retrieval](#). *IEEE Transactions on Knowledge and Data Engineering*, 35(10):9946–9958.
- Hongbin Huang, jiao Sun, Hui Wei, Kaiming Xiao, Mao Wang, and Xuan Li. 2023b. "a dataset of domain events based on open-source military news". *Science Data Bank*, pages 30(1–10).
- Lucian-Marius Ivanov. 2011. Management of intelligence activity improving the intelligence analysis process as a pre-requisite to increasing the efficiency of military intelligence services. *Revista Academiei Fortelor Terestre*, pages 20–28.
- Sha Li, Heng Ji, and Jiawei Han. 2021a. [Document-level event argument extraction by conditional generation](#). *CoRR*, abs/2104.05919.

- Sha Li, Heng Ji, and Jiawei Han. 2021b. [Document-level event argument extraction by conditional generation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 894–908, Online. Association for Computational Linguistics.
- Xinyu Li, Fayuan Li, Lu Pan, Yuguang Chen, Weihua Peng, Quan Wang, Yajuan Lyu, and Yong Zhu. 2020. [Duee: A large-scale dataset for chinese event extraction in real-world scenarios](#). In *Natural Language Processing and Chinese Computing*, pages 534–545, Cham. Springer International Publishing.
- 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.
- Di Lu, Shihao Ran, Joel Tetreault, and Alejandro Jaimes. 2023. [Event extraction as question generation and answering](#).
- Yaojie Lu, Hongyu Lin, Jin Xu, Xianpei Han, Jialong Tang, Annan Li, Le Sun, Meng Liao, and Shaoyi Chen. 2021. [Text2event: Controllable sequence-to-structure generation for end-to-end event extraction](#). *CoRR*, abs/2106.09232.
- Hui Li;Guoyu Yan;Xin Zhao;Jie Zhang;Ming Lyu. 2022. [Tactical mission event logic graph construction for network-centric warfare](#). *Alexandria Engineering Journal*, pages 9161–9173.
- 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.
- Hao Peng, Xiaozhi Wang, Feng Yao, Zimu Wang, Chuzhao Zhu, Kaisheng Zeng, Lei Hou, and Juanzi Li. 2023. [Omnievent: A comprehensive, fair, and easy-to-use toolkit for event understanding](#).
- Sampo Pyysalo, Tomoko Ohta, and Sophia Ananiadou. 2013. [Overview of the cancer genetics \(CG\) task of BioNLP shared task 2013](#). In *Proceedings of the BioNLP Shared Task 2013 Workshop*, pages 58–66, Sofia, Bulgaria. Association for Computational Linguistics.
- R Vallikannu;V Kanpur Rani;Bc Kavitha;P Sankar. 2023. [An analysis of situational intelligence for first responders in military](#). In *2023 International Conference on Artificial Intelligence and Applications (ICAIA) Alliance Technology Conference (ATCON-1)*.
- Julia Santucci. 2022. [After afghanistan: Intelligence analysis and us military missions](#). *Survival*, pages 157–178.
- Philip A Schrodt. 2012. [Precedents, progress, and prospects in political event data](#). *International Interactions*, 38(4):546–569.
- Jiawei Sheng, Shu Guo, Bowen Yu, Qian Li, Yiming Hei, Lihong Wang, Tingwen Liu, and Hongbo Xu. 2021. [CasEE: A joint learning framework with cascade decoding for overlapping event extraction](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 164–174, Online. Association for Computational Linguistics.
- Ge Shi, Yunyue Su, Yongliang Ma, and Ming Zhou. 2023. [A hybrid detection and generation framework with separate encoders for event extraction](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3163–3180, Dubrovnik, Croatia. Association for Computational Linguistics.
- Matthew Sims, Jong Ho Park, and David Bamman. 2019. [Literary event detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3623–3634, Florence, Italy. Association for Computational Linguistics.
- Elena A Skryabina, Naomi Betts, Gabriel Reedy, Paul Riley, and Richard Amlôt. 2020. [The role of emergency preparedness exercises in the response to a mass casualty terrorist incident: a mixed methods study](#). *International journal of disaster risk reduction*, 46:101503.
- 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.
- Qizhi Wan, Changxuan Wan, Keli Xiao, Dexi Liu, Chenliang Li, Bolong Zheng, Xiping Liu, and

- Rong Hu. 2023. [Joint document-level event extraction via token-token bidirectional event completed graph](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10481–10492, Toronto, Canada. Association for Computational Linguistics.
- Lancheng Wang, Silong Zhang, and Hexu Xu. 2021a. Research on the construction of event evolutionary graph of network public opinion. *Chinese Journal of Medical Library and Information Science*, pages 17–23.
- Sijia Wang, Mo Yu, Shiyu Chang, Lichao Sun, and Lifu Huang. 2022. [Query and extract: Refining event extraction as type-oriented binary decoding](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 169–182, Dublin, Ireland. Association for Computational Linguistics.
- Xiaozhi Wang, Hao Peng, Yong Guan, Kaisheng Zeng, Jianhui Chen, Lei Hou, Xu Han, Yankai Lin, Zhiyuan Liu, Ruobing Xie, Jie Zhou, and Juanzi Li. 2023. [Maven-arg: Completing the puzzle of all-in-one event understanding dataset with event argument annotation](#).
- 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](#). *CoRR*, abs/2004.13590.
- Ziqi Wang, Xiaozhi Wang, Xu Han, Yankai Lin, Lei Hou, Zhiyuan Liu, Peng Li, Juanzi Li, and Jie Zhou. 2021b. [CLEVE: Contrastive Pre-training for Event 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 6283–6297, Online. Association for Computational Linguistics.
- Jinhang Xu, Biying Yan, Feng Chen, and Yaqi Cao. 2021. Dynamic fusion of emergency response plan based on domain knowledge graph. *Computer Systems and Applications*, pages 1–13.
- Xin Xu, Jian Xu, Guoqing Ruan, Hongyi Bao, and Jiadong Sun. 2022. A pipeline-based multimodal military event argument extraction framework. In *CCKS 2022 - Evaluation Track*, pages 21–29, Singapore. Springer Nature Singapore.
- Zhiyang Xu, Jay Yoon Lee, and Lifu Huang. 2023. [Learning from a friend: Improving event extraction via self-training with feedback from Abstract Meaning Representation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10421–10437, Toronto, Canada. Association for Computational Linguistics.
- Hang Yang, Yubo Chen, Kang Liu, Yang Xiao, and Jun Zhao. 2018. [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.
- Hang Yang, Dianbo Sui, Yubo Chen, Kang Liu, Jun Zhao, and Taifeng Wang. 2021a. [Document-level event extraction via parallel prediction networks](#). 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 6298–6308, Online. Association for Computational Linguistics.
- 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.
- Xuesheng Yang, Qian Huang, and Yong Guan. 2021b. Method of emergency response force planning. *Command Control and Simulation*, pages 114–118.
- 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.
- Haitao Zhang, Jiawei Li, Weili Li, and Yashu Liu. 2021. [Research on the construction of major emergencies event centric knowledge graph](#). *Library and Information Service*, 65(133-140).
- Ruoyu Zhang, Yanzeng Li, Yongliang Ma, Ming Zhou, and Lei Zou. 2023. [Llmaa: Making large language models as active annotators](#).
- Yuheng Zhang, Sheng Zheng, and Zhang Sheng. 2022. [Event extraction for military target motion in open-source military news](#). In *2022 International Conference on Artificial Intelligence and Computer Information Technology (AICIT)*, pages 1–4.



Shun Zheng, Wei Cao, Wei Xu, and Jiang Bian.  
2019. [Doc2edag: An end-to-end document-level framework for chinese financial event extraction](#). *CoRR*, abs/1904.07535.