LEVEN: A Large-Scale Chinese Legal Event Detection Dataset

Feng Yao[1*]**, Chaojun Xiao**[2,3*]**, Xiaozhi Wang**[2,3]**, Zhiyuan Liu**[2,3,4,5†]**,**
Lei Hou[2,3]**, Cunchao Tu**[6]**, Juanzi Li**[2,3]**, Yun Liu**[1]**, Weixing Shen**[1†]**, Maosong Sun**[2,3,4,5]

[1]School of Law, Institute for AI and Law, Tsinghua University, Beijing, China
[2]Dept. of Comp. Sci. & Tech., Institute for AI, Tsinghua University, Beijing, China
[3]Beijing National Research Center for Information Science and Technology, China
[4]International Innovation Center of Tsinghua University, Shanghai, China
[5]Beijing Academy of Artificial Intelligence, Beijing, China
[6]Beijing Powerlaw Intelligent Technology Co., Ltd., China
`{yaof20,xiaocj20}@mails.tsinghua.edu.cn`
`{liuzy,wxshen}@tsinghua.edu.cn`

Abstract

Recognizing facts is the most fundamental step in making judgments, hence detecting events in the legal documents is important to legal case analysis tasks. However, existing Legal Event Detection (LED) datasets only concern incomprehensive event types and have limited annotated data, which restricts the development of LED methods and their downstream applications. To alleviate these issues, we present LEVEN, a large-scale Chinese LEgal eVENt detection dataset, with $8,116$ legal documents and $150,977$ human-annotated event mentions in 108 event types. Not only charge-related events, LEVEN also covers general events, which are critical for legal case understanding but neglected in existing LED datasets. To our knowledge, LEVEN is the largest LED dataset and has dozens of times the data scale of others, which shall significantly promote the training and evaluation of LED methods. The results of extensive experiments indicate that LED is challenging and needs further effort. Moreover, we simply utilize legal events as side information to promote downstream applications. The method achieves improvements of average 2.2 points precision in low-resource judgment prediction, and 1.5 points mean average precision in unsupervised case retrieval, which suggests the fundamentality of LED. The source code and dataset can be obtained from https://github.com/thunlp/LEVEN.

1 Introduction

Finding out the occurred events and causal relations between them is fundamental to analyzing legal cases and making judgments. Legal event detection (LED) aims to automatically extract the

---

*Equal contribution. Listing order is random.
†Corresponding authors.

Fact Description
Alice drove a car at night and *crashed into* Bob, a pedestrian, on Green Avenue. To prevent being spotted, Alice took Bob away from the scene, *dumped* him under an isolated bridge and *drove off* in a panic. Two hours later, Bob *died* of excessive bleeding ...

Event Detection

Event Timeline

crashed into　　*dumped*　　*drove off*　　*died* / **Death**

Bodily_harm　　Desertion　　Escaping

dumped Trigger Word　　**Desertion** Event Type　　**Result**

Related Law Article
Traffic accident crime ... if the *hit-and-run* occurs, the crime should be sentenced to *imprisonment more than 3 years but less than 7 years* ... if the perpetrator *abandons* the victim, resulting in the *death*, he shall be convicted of Intentional homicide crime and sentenced to *death, life imprisonment or imprisonment of no less than 10 years* ...

Crime & Prison Term
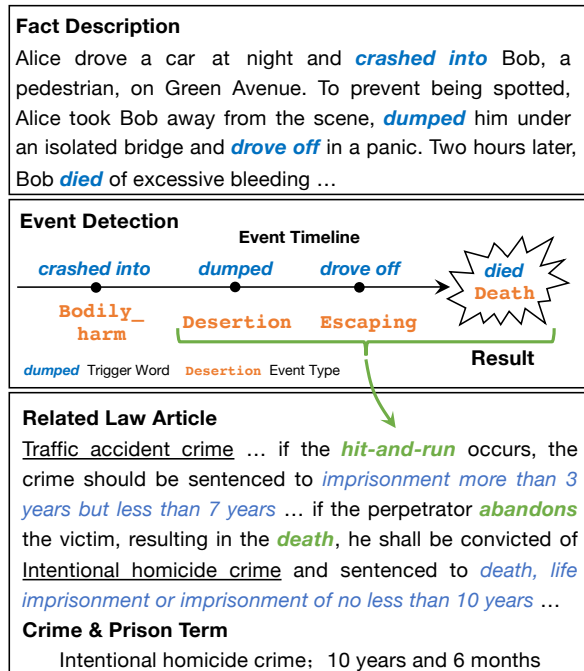Intentional homicide crime；10 years and 6 months

Figure 1: An example legal document describing the fact with the annotated event triggers, the corresponding event types, the related law article, and penalties.

event triggers from legal cases and then classify their corresponding event types, which will naturally benefit many legal artificial intelligence applications, such as Legal Judgment Prediction (LJP) and Similar Case Retrieval (SCR) (Zhong et al., 2020a). For instance, Figure 1 shows a case with the trigger words highlighted in the plain text and the corresponding event types. Based on the detected events, we can observe that Alice causes a traffic accident, and the subsequent `Desertion` and `Escaping` events jointly result in the `Death` event, which changes Alice's charge from traffic accident crime to intentional homicide crime and increases the expected penalties.

Inspired by the previous success for general-domain event detection (Ji and Grishman, 2008; Li et al., 2013; Chen et al., 2015; Nguyen et al., 2016; Feng et al., 2016; Yan et al., 2019; Wang et al., 2020), some works attempt to build LED systems with hand-crafted features (Lagos et al., 2010; Bertoldi et al., 2014), or neural networks (Li et al., 2019, 2020a; Shen et al., 2020). However, two major challenges of existing LED resources seriously restrict the development of LED methods:

(1) **Limited Data.** Existing LED datasets (Shen et al., 2020; Li et al., 2020a) only contain thousands of event mention annotations, which can not provide sufficient training signals and reliable evaluation results. To promote the progress of legal information extraction and legal document analysis, it is an urgent need to develop a large-scale and high-quality dataset for the LED task. (2) **Incomprehensive Event Schema.** Existing LED works merely concern a dozen of charge-oriented event types, which are either the judicial event types defined in general-domain datasets (Maxwell et al., 2009) or some newly-defined charge-oriented event types to meet specific downstream requirements (Li et al., 2019, 2020a; Shen et al., 2020). Their event schemata only cover a narrow scope of charges. Besides, existing datasets focus on charge-oriented events and ignore the general events in the cases, such as Desertion and Escaping in Figure 1, which are also critical for analyzing legal cases.

To alleviate the above issues and provide a solid foundation for LED, we present LEVEN, a large-scale Chinese legal event detection dataset, based on the cases published by the Chinese government[1]. We highlight LEVEN with the following merits:

(1) **Large scale**. LEVEN contains 8,116 legal documents covering 118 criminal charges and has 150,977 human-annotated event mentions, which is dozens of times larger than previous LED datasets. To the best of our knowledge, LEVEN is also the largest Chinese event detection dataset. Based on the scale, we believe LEVEN can well train and reliably benchmark data-driven LED methods, which shall promote this field. (2) **High coverage**. LEVEN contains 108 event types in total, including 64 charge-oriented events and 44 general events. The LEVEN event schema has a sophisticated hierarchical structure, which is shown in appendix E. To build the schema, we conduct a two-stage event schema construction process. We

first summarize the critical charge-oriented event types based on law articles and then simplify and supplement the event schema based on the events in sample cases. The two-stage process ensures the high coverage of LEVEN schema.

To explore the challenges of LEVEN, we implement some state-of-the-art models and evaluate them on our dataset. The results show that though existing models can achieve better performance on legal documents than in the general domain, it still needs future efforts to reach a practical level.

Moreover, we demonstrate the fundamentality of LED for downstream Legal AI applications. Specifically, we train an LED model on LEVEN and use it to detect events for unlabeled legal documents. We then use the auto-detected events as side information to handle LJP and SCR. Experiments show that the performance of these two tasks can be substantially improved in this simple way, indicating that LED can provide helpful fine-grained information and thus serve as a fundamental process in Legal AI. Hence we advocate more research attention to LED.

2 Related Work

2.1 Event Detection

Event detection (ED) is an important information extraction task and many efforts have been devoted to (Ji and Grishman, 2008; Li et al., 2013; Chen et al., 2015; Nguyen et al., 2016; Liu and Zhao, 2017; Zhao et al., 2018; Yan et al., 2019; Wang et al., 2021b). The majority of existing ED datasets are developed for the general domain (Christopher et al., 2006; Song et al., 2015; Wang et al., 2020) and mostly for English. Besides, some datasets are also developed for specific domains (Thompson et al., 2009; Kim et al., 2008; Ritter et al., 2012; Yang et al., 2018; Zheng et al., 2019) and Chinese (Li et al., 2020b). Considering the rapid growth of Chinese legal artificial intelligence (Zhong et al., 2020a), we believe constructing Chinese LED datasets is helpful and necessary. In the context of LED, some works define specific legal event types to analyze for legal documents (Maxwell et al., 2009; Lagos et al., 2010; Li et al., 2019; Shen et al., 2020; Li et al., 2020a), but these constructed datasets are typically small-scale and cannot well train and evaluate practical LED systems. Hence we construct LEVEN, which is the largest LED dataset and also the largest Chinese event detection dataset to our knowledge.

---

[1] https://wenshu.court.gov.cn/

2.2 Legal Artificial Intelligence

Thanks to the rapid progress of natural language processing and the openness of legal documents, legal artificial intelligence (LegalAI) has drawn increasing attention from both AI researchers and legal professionals in recent years (Bommarito II et al., 2021; Ye et al., 2018; Chalkidis et al., 2021; Zhong et al., 2020a; Tsarapatsanis and Aletras, 2021; Wang et al., 2021a). LegalAI can not only provide handy references for people who are not familiar with legal knowledge, but also reduce the redundant paperwork for legal practitioners. Many efforts have been devoted to a variety of LegalAI tasks, including legal judgment prediction (Zhong et al., 2018; Chalkidis et al., 2019; Yang et al., 2019), legal question answering (Ravichander et al., 2019; Zhong et al., 2020b; Kien et al., 2020), contract review (Hendrycks et al., 2021; Zhang et al., 2021; Koreeda and Manning, 2021), legal case retrieval (Ma et al., 2021; Shao et al., 2021), and legal pre-trained models (Chalkidis et al., 2020; Xiao et al., 2021). Most existing works focus on the application in LegalAI while ignoring the basic key event information in the legal documents. Some works attempt to extract events from the legal documents (Li et al., 2019; Shen et al., 2020; Li et al., 2020a). But these works are limited to the event coverage and the number of annotation instances. We argue that our proposed large-scale and comprehensive dataset, LEVEN, can promote the development of legal event detection and thus benefit downstream legal artificial intelligence tasks.

3 Data Collection

Our ultimate goal is to construct a large-scale legal event detection dataset with a high-coverage event type schema and sufficient event instances, which is scarce in existing LED datasets. Therefore, we need to redefine a new event schema, select the trigger candidates, and annotate the corresponding event types. As criminal cases usually involve principal rights and complex facts, we focus on criminal legal events in this paper. In the following sections, we first introduce the construction of event schema and then describe the process of annotation of candidates and related event types.

3.1 Event Schema Construction

To construct an event schema with high coverage, we need to consider events for both judicial behaviors and general behaviors. Therefore, we follow a two-stage process to define our new event schema: 1) We first collect charge-oriented events based on the law articles and legal textbooks. 2) We then collect general events from the sampled case documents. The two-stage process enables LEVEN to cover essential events recorded in legal documents.

Inspired by previous works (Li et al., 2020a; Shen et al., 2020), in the first stage, we use law articles and a classical legal textbook, *Specific Theory of Criminal Law*, as our references to summarize the charge-oriented events. *Criminal Law* provides the definition of each criminal charge and a hierarchical structure for these charges. We first collect 459 criminal charges, which are then divided into 61 types based on the targets and measures of criminal behaviors. Considering that some criminal charges are too abstract to be specific event types (e.g., `dereliction`), we manually filter out them. Besides, as there are some similar charges involving the same event types (e.g., `intentional_homicide` and `involuntary_homicide`), we merge them. After the first stage, we obtain 198 event types highly correlating to the criminal charges.

As the charge-oriented event schema is constructed from legal professional references, there are two main issues: 1) The charge-oriented event schema mainly focuses on illegal behaviors, while ignoring important general behaviors. 2) There are some event types that infrequently or never occur in real-world cases. To address these issues, we further modify the event schema based on the summarization of real-world cases. Specifically, we sample 20 case documents for each criminal charge, which can ensure good coverage. And then we invite a legal expert to manually extract and summarize the event mentions occurring in sampled cases. Based on the extracted events, we further filter out the abstract event types and merge some detailed event types in the schema. We finally get 108 event types to annotate, with both charge-oriented events and general events.

According to the criminal theory, the key elements of the crime include the act, the harmful results, and the causal relation between them. Therefore, we organize the event types in a hierarchical structure, with three categories representing behavior and a category representing results. During the annotation process, the annotators are required to label the most fine-grained types. Please refer to Appendix E for details of the event schema.

Dataset	#Documents	#Tokens	#Sentences	#Event Types	#Event Mentions	Language	Domain
MAVEN	4,480	1,276k	49,873	168	118,732	English	General
ACE2005-zh	633	185k	7,955	33	4,090	Chinese	General
DuEE	11,224	530k	16,900	65	19,640	Chinese	General
DivorceEE*	3,100	–	–	13	–	Chinese	Legal
CLEE*	3,000	–	6,538	5	6,538	Chinese	Legal
DyHiLED*	–	–	–	11	2,380	Chinese	Legal
LEVEN	8,116	2,241k	63,616	108	150,977	Chinese	Legal

Table 1: The statistics of widely-adopted event detection datasets. For Chinese datasets, we adopt JIEBA toolkit to perform tokenization. Datasets denoted with * are not publicly available, and – means the value is not accessible.

3.2 Document Selection

To support the manual annotation, we adopt cases collected from the government website as our data source. Following Xiao et al. (2018), we only keep the criminal judgment documents for annotation.

We first extract the related charges with regular expression from the documents and divide each document into several parts based on the content, where only the fact description is maintained. Moreover, to ensure the dataset quality, we filter out the documents with less than 50 characters and more than 2, 500 characters in fact description. Notably, though we get 198 charges in the first stage of event schema construction, there are some charges where no cases are published due to the privacy and secrecy involved. Therefore, we get case documents with only 107 charges. We randomly sample 200 documents for charges with high frequency and maintain all cases for charges with low frequency. Finally, we select 8, 288 documents for annotation. After discarding the low-quality documents labeled by annotators, we finally retain 8, 166 documents.

3.3 Candidate Selection

The annotation of LED dataset requires the annotators to find the triggers from the documents and label the corresponding event types within 108 options. Following Wang et al. (2020), we adopt heuristic methods to automatically select the trigger candidates and narrow down the event type options for each trigger candidate.

Candidate trigger selection. Inspired by Chen et al. (2017), which utilizes the lexical unit in FrameNet (Baker et al., 1998) to select trigger words, we require a legal expert to collect semantic-related words for each event type in our schema. And we obtain a semantic vocabulary consisting of 1, 013 words with their corresponding event types. Then we apply tokenization and POS tagging with

JIEBA toolkit[2], and all the content words, including nouns and verbs, are selected as trigger candidates. Besides, the words in the collected vocabulary are also selected as trigger candidates.

Candidate event type selection. Further, we recommend 30 event types for each trigger candidate, which can provide references for annotators. We first calculate the cosine similarity between the representations of trigger candidates and event types. And then we rank the event types by the calculated similarity and retrieve the top 30 ones as the recommended candidates. Here, we utilize the representations calculated by SBERT (Reimers and Gurevych, 2019), which can generate semantic meaningful embeddings.

The automatic candidate selection mechanism aims to provide a good reference for the annotators. Notably, considering that not all triggers and event types can be automatically selected, we also require the annotators to label the words and event types that are not in the recommended list. The final annotation results show that 95.6% trigger words and 92.8% event types are recommended, and the rest are supplemented by annotators manually. The results demonstrate that the automatic candidate selection is helpful to improve the annotation efficiency, and the annotators can also label the trigger words and event types that are not recommended.

3.4 Human Annotation

The final process is to annotate the triggers from documents manually. We write a 59-page annotation guideline in Chinese to help the annotators better understand the annotation task. We also embed the guideline in the annotation platform so that the annotators can easily refer to it. A simplified version in English is provided in Appendix F.

Following previous works (Christopher et al., 2006; Wang et al., 2020), we adopt a two-stage

---

[2]https://github.com/fxsjy/jieba

Top-level Event Type	Category	#Type	#Mention	Percentage	Sub-type Examples
General_behaviors	Behavior	40	68,616	45.4%	Selling, Employing, Manufacturing
Prohibited_acts	Behavior	40	43,021	28.5%	Killing, Blackmail, Theft, Destroying
Judicature_related	Behavior	13	29,709	19.7%	Arrest, Surrendering
Consequences	Result	7	6,832	4.5%	Death, Injury, Being_trapped
Accident	Result	4	2,742	1.8%	Traffic_accident, Fire_accident
Natural_disaster	Majeure	4	57	0.03%	Drought, Flood_and_waterlogging

Table 2: Data distribution over the top-level event types and the corresponding categories and samples.

annotation process. In the first stage, we invite crowd-source annotators to choose the correct answers from case documents when given the candidate triggers and corresponding event types. Each document is annotated independently twice. The annotators first went through several hours of training for labeling, so as to ensure the annotation quality. Besides, for each labeled document, we discard the annotation results and require another two annotators to annotate it, if the inter-annotator agreement of the document is lower than 0.2. In the second stage, we invite experienced annotators to choose final event types given the results of the first-stage annotation. Only the results labeled differently in the first stage are required to be labeled again in the second stage.

We measure the data quality via inter-annotator agreements between two annotators, i.e., Cohen's Kappa (Cohen, 1960). The Kappa coefficient for the first stage is 0.609. To evaluate the data quality in the second stage, we randomly sample 5% documents to be labeled twice independently. The Kappa coefficient for the second stage is 0.875. The satisfactory Kappa coefficient demonstrates that LEVEN is a high-quality manually annotated LED dataset, and we hope the dataset can accelerate the development of LED and legal case analysis.

4 Data Analysis

In this section, we conduct analysis from various aspects to provide a deep understanding of LEVEN.

4.1 Data Size

The detailed statistics of LEVEN and some widely-used event detection datasets are shown in Table 1. We compare LEVEN with two types of datasets: (1) **General-domain ED datasets.** ACE2005 (Christopher et al., 2006) is the most popular multi-lingual event extraction dataset and here we compare with its Chinese subset (denoted as ACE2005-zh). MAVEN (Wang et al., 2020) is the largest general-domain event detection dataset, with 168 event types and hundreds of thousands

of event instances. DuEE (Li et al., 2020b) is the largest Chinese ED dataset, which is collected from Chinese news articles. (2) **LED datasets.** DivorceEE (Li et al., 2019) focuses on event extraction in divorce cases. CLEE (Li et al., 2020a) is for larceny cases. DyHiLED (Shen et al., 2020) is a LED dataset with a hierarchical event schema.

From the comparisons, we can observe that LEVEN is the largest LED dataset with dozens of the scale of previous LED datasets and is also the largest Chinese event detection dataset. LEVEN's scale is even comparable to the previous largest general-domain event detection dataset MAVEN. Moreover, LEVEN contains the most event types among the Chinese event detection datasets. These suggest that LEVEN may help LED, Chinese ED, and general ED at the same time.

	#Doc.	#Sentences	#Event	#Negative.
Training	5,301	41,238	98,410	297,252
Validation	1,230	9,788	22,885	69,645
Test	1,585	12,590	29,682	90,512

Table 3: The detailed statistics of subsets of LEVEN.

4.2 Data Distribution

Event Types. As mentioned before, our event schema contains three event categories representing behavior, two event categories representing results, and one event category representing force majeure. The instance distribution over these top-level event types is shown in Table 2. There are 45.4% events belonging to general behavior, which is ignored in the previous LED dataset. It demonstrates that modeling the general events in LED is necessary. Besides, LEVEN meets long-tail distribution, which raises a challenge for future research.

Number of Instances. LEVEN is a large-scale dataset, where 89.6% event types contains more than 100 event mentions, and 43.4% event types contains more than $1,000$ event mentions. Therefore, LEVEN can provide sufficient training signals and reliable evaluation results for LED.

Model	Micro			Macro		
	Precision	Recall	F1	Precision	Recall	F1
DMCNN	**85.88** ± 0.70	79.70 ± 0.59	82.67 ± 0.08	80.55 ±0.49	73.31 ± 3.88	75.03 ± 0.40
BiLSTM	83.09 ± 0.89	85.16 ± 0.95	84.11 ± 0.24	78.70 ± 0.92	76.67 ± 2.23	76.65 ± 1.42
BiLSTM+CRF	84.74 ± 0.55	83.33 ± 0.49	84.03 ± 0.05	78.56 ± 1.31	72.60 ± 1.11	74.49 ± 0.77
BERT	84.19 ± 0.39	84.31 ± 0.34	84.25 ± 0.18	79.61 ± 0.91	76.76 ± 1.79	77.33 ± 1.30
BERT+CRF	83.82 ± 0.48	84.56 ± 0.52	84.19 ± 0.09	79.77 ± 1.10	77.65 ± 2.20	77.84 ± 1.58
DMBERT	84.77 ± 0.91	**86.22** ± 0.77	**85.48** ± 0.18	**81.57** ± 1.04	**80.90** ± 1.38	**80.34** ± 0.74

Table 4: The test performances of ED baselines on LEVEN. Refer to Appendix A.1 for validation performances.

5 Experiments

5.1 Benchmark Settings

We randomly split the dataset into training set, validation set, and test set according to the ratio, $0.65 : 0.15 : 0.2$. Following Wang et al. (2020), we provide official negative samples for a fair comparison between different methods. As stated in Section 3.3, we first employ Chinese word segmentation and POS-tagging to the documents, and then select the content words (verbs and nouns) or words in the human-collected semantic vocabulary as the trigger candidates. The detailed statistics of the data splits are listed in Table 3. As the dataset is unbalanced, we adopt both the micro-averaged and macro-averaged precision, recall, and F1 score as the evaluation metrics for the experiments.

5.2 Baseline Models

Event detection has been explored for decades. In this section, we evaluate several competitive baseline models, which are widely used in the general domain event detection task, on LEVEN, including (1) Token classification. We first encode the given sentences with deep neural networks, including **BiLSTM** (Hochreiter and Schmidhuber, 1997) and **BERT** (Devlin et al., 2019), and then use the hidden representations of the candidate triggers to classify their corresponding event types. (2) Dynamic max-pooling. These models adopt convolutional neural network (**DMCNN**, Chen et al. (2015)) or pre-trained language model, BERT (**DMBERT**, Wang et al. (2019)), to extract the sequence features, and employ dynamic pooling layers to obtain trigger-specific representation for each candidate. (3) Sequence labeling. Different from previous models, we adopt sequence labelling models (**BiLSTM+CRF**, **BERT+CRF**) to capture the correlations between different events. The implementation details can be found in Appendix A.1. We run each experiment 5 times, and the averages and standard deviations of the results are reported.

5.3 Overall Performance Comparison

The baseline results are shown in Table 4. And we can observe that (1) DMBERT can outperform other baselines significantly, with the micro-F1 score of 85.48%, which is still not satisfactory for real-world applications. (2) The standard deviations on the micro-metrics are relatively small, indicating that LEVEN contains sufficient data in the test set and can provide stable evaluation results. (3) From the comparison between BiLSTM-based and BERT-based models, we find that BERT-based models cannot achieve significant improvement on LEVEN. It suggests that designing event-oriented pre-trained models is necessary for LED, which we leave for future work. (4) CRF-based models perform slightly worse than their corresponding token classification models. We attempt to employ CRF to capture the dependencies of multiple events as suggested by Wang et al. (2020), while the result is inconsistent with the expectation. Therefore, it still needs exploration to model the correlations between multiple events in a single sentence.

Notably, as the legal documents are well-written and the used language is more standardized than the general domain, the event detection models can achieve better performance on LEVEN than on the general domain dataset (DMBERT can only achieve 67.1% micro-F1 score on MAVEN (Wang et al., 2020), while 85.5% on LEVEN.) Therefore, we can apply existing LED models to promote the downstream tasks (Section 5.5). However, the performance is still unsatisfactory and needs future research (Section 5.4).

5.4 Error Analysis

To analyze the defect of existing models and point out the future directions for LED, we conduct error analysis on the prediction errors of the model with the best performance. We categorize the prediction errors into several types and find some challenges which require future efforts.
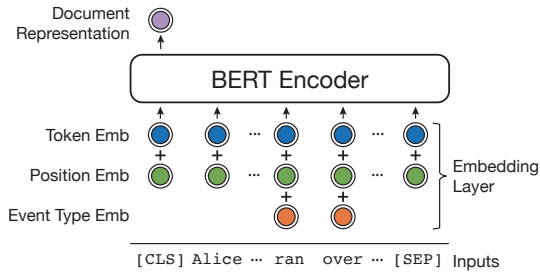
Figure 2: The framework for downstream tasks.

(1) **Long-tail Problem**. Though LEVEN contains hundreds of thousands of event mentions, there are some event types with limited instances inevitably. We compute the performance on low-frequency event types, where the micro-F1 score is 65.97% for event types with less than 50 instances and 72.24% for event types with less than 100 instances. There is still a large gap between the performance of the low-frequency types and the overall average performance. More discussion can be found in Appendix.

(2) **Context-aware Prediction**. Many triggers require the model to integrate the information of the complex context from argument entities or other sentences to predict corresponding event types. For instance, in the sentence *Bob rushed to call ENT to inform the situation*, if *ENT* is *the police* or *110*, the event type for trigger *call* is Reporting_to_police, while if *ENT* is other people, the event type is Reporting. Sometimes, the essential information comes from other sentences, which require the model to capture cross sentence dependency. We randomly sample 100 cases and ask another annotator to count the number of errors that need context-aware prediction. From the statistics, 36.98% errors are caused by incorrectly capturing contextual information, which still needs further effort.

(3) **Identification Mistakes**. Similar to Wang et al. (2020), the most common mistake is confusing the negative samples and positive samples, i.e., the false positive and false negative. The results show that 48.99% and 34.41% errors are false positive and false negative, respectively. Therefore, how to identify the event semantic is a challenge.

5.5 Applications of Legal Event Detection

Furthermore, in order to provide a perspective of how to use LEVEN for other Legal AI tasks and to verify the effectiveness of LED for legal documents analysis, we utilize legal events as side information

Model	Charge			Law			Term
	P	R	F1	P	R	F1	Dis ↓
50-shot							
BERT	76.6	**77.0**	76.8	73.6	**76.8**	75.2	2.398
+ event	**79.2**	76.2	**77.7**	**75.4**	75.6	**75.5**	**2.364**
full							
BERT	88.2	89.4	88.8	83.7	86.8	85.2	1.895
+ event	**88.2**	**89.7**	**88.9**	**83.8**	**87.7**	**85.7**	**1.878**

Table 5: The results for legal judgment prediction. Here P, R, and F1 indicate precision, recall, and F1 scores, respectively.

in two typical downstream tasks in legal artificial intelligence, including Legal Judgment Prediction (LJP) and Similar Case Retrieval (SCR).

In the following sections, we will first introduce the encoder architecture and applications of LED in legal judgment prediction and similar case retrieval. We compare the performance of the original BERT model and BERT model with event features to show the effectiveness of LED. Notably, the event features can either be used independently or fed into other models to further promote the performance. The details about model implementation and dataset statistics can be found in Appendix A.2.

Encoder Architecture

As pre-trained language models have achieved promising results in many legal tasks (Chalkidis et al., 2020; Xiao et al., 2021), we adopt the BERT as our basic encoder. To verify the effectiveness of event detection in LegalAI, we only make minor changes in the embedding layer to integrate the event information. Figure 2 illustrates the encoder architecture. Given the input document, to highlight the event information, we first employ the **BERT+CRF** model to detect the trigger words and their event types from the case documents[3]. And then we utilize the event information in the BERT model by adding the event type embedding in the input embedding layer. The event type embedding is randomly initialized and updated during the training process. Specifically, for non-trigger tokens, we feed the sum of the token embeddings and position embeddings into the encoder. For trigger tokens, we also define an event type embedding for each event type and add the corresponding event type embeddings to the inputs.

---

[3]Notably, due to the high computational complexity of DMBERT, we use BERT+CRF here, which also achieve comparable results with DMBERT.

Model	MAP	NDCG@10	NDCG@20	NDCG@30	P@5	P@10
BM25	48.40	73.10	79.70	88.80	40.60	38.10
TFIDF	45.70	79.50	83.20	84.80	30.40	26.10
LMIR	49.50	76.90	81.80	90.00	43.60	40.60
Bag-of-Event	50.94	78.37	83.66	90.32	44.11	42.62
Bag-of-Event$_w$	**51.02**	**79.90**	**84.42**	**90.97**	**45.23**	**43.36**
BERT	51.92	79.23	84.12	91.28	44.49	40.10
+ event	**51.99**	**80.10**	**84.92**	**91.73**	**44.63**	**41.22**

Table 6: The experiment results under both unsupervised and supervised settings for similar case retrieval.

Legal Judgment Prediction

Legal judgment prediction (LJP) aims to predict the judgment results, including related laws, charges, and prison terms, based on the textual fact description, and LJP is an essential task for LegalAI (Zhong et al., 2018; Chalkidis et al., 2019; Yang et al., 2019). LJP requires the model to capture the key event information and mine the causal relationship between behaviors and consequences.

Therefore, in this section, we attempt to investigate the effect of LED for judgment prediction. We adopt the CAIL2018 (Xiao et al., 2018) as the evaluation benchmark, which is the largest LJP dataset. Following Zhong et al. (2020a), we formalize LJP as a multi-task learning problem. Specifically, we formalize law article prediction and charge prediction as multi-label classification tasks, and adopt binary cross-entropy function as the loss. We formalize prison term prediction as a regression task and adopt the log distance function as the loss. As for the output layer, we feed the document representation into three linear layers for the prediction of three tasks, respectively. In addition to training the model with the full dataset, we also explore the effectiveness of legal events under a low-resource setting. We only sample 50 cases for each charge and law article to train the model.

The results are shown in Table 5. We can observe that LED can promote the performance of LJP, especially under low-resource settings, which proves the effectiveness of LED. Besides, LED can only achieve slight improvement on charge prediction and law prediction under with full training dataset, while can achieve consistent improvement on prison term prediction. That is because prison term prediction is more complex and requires the model to capture both the criminal behaviors and severity-level of the consequences. Legal events can provide fine-grained information for predicting prison terms, and thus promote the performance under both low-resource and full dataset settings.

Similar Case Retrieval

Similar case retrieval (SCR) aims to retrieve supporting cases given a query case, which is a widely-applied task with high practical value (Kano et al., 2018; Shao et al., 2021). SCR requires the leverage of fine-grained information from multiple perspectives, including element-level, event-level, and law-level. In this paper, we adopt LeCaRD (Ma et al., 2021) as the evaluation benchmark, which contains 107 queries and 43,000 candidates. We adopt 5-fold cross-validation for evaluation, and employ the top-k Normalized Discounted Cumulative Gain (NDCG@k), Precision (P@k), and Mean Average Precision (MAP) as evaluation metrics.

We verify the effectiveness of utilizing event features for similar case retrieval task under both unsupervised and supervised settings. In the unsupervised setting, we adopt "Bag-of-Event", i.e., the frequency of each event, as the representation of each legal document, and use cosine similarity to compute the similarity scores between two different legal documents. Further, considering the fact that the events occurring in the legal cases are not equally important, we compute the inverse document frequencies for different event types in the TF-IDF fashion, which are used as the weights of different event types. We denote the weighted representation as Bag-of-Event$_w$. In the supervised setting, we train the BERT model in a sentence-pair classification paradigm. We concatenate the query case and candidate case as the input sequence, and require the model to classify whether the two cases are relevant or not.

The results are shown in Table 6. From the results, we can observe that both Bag-of-Event and Bag-of-Event$_w$ are powerful representation methods for similar case retrieval and can achieve superior performance than other unsupervised models. Besides, the event information can facilitate the performance of BERT model, which further proves that event information is crucial for case retrieval.

6 Conclusion and Future Work

In this paper, we construct the largest legal event detection dataset, LEVEN, which contains a comprehensive legal event schema and hundreds of thousands of event mentions. We evaluate several competitive baseline models and conduct error analysis for these models on LEVEN. The experimental results prove that it still needs future efforts to promote the development of LED. Furthermore, we employ LED for downstream legal document analysis, including legal judgment prediction and similar case retrieval. It indicates that LED can provide fine-grained information and serve as a fundamental process for legal artificial intelligence. In the future, we will explore to conduct more analysis on large-scale legal documents based on the event information, and annotate LEVEN with event relations and event arguments.

Ethical Considerations

LEVEN focuses on detecting events from the fact and does not involve any value judgment. LED aims to transform the unstructured legal text into structured event information, which is helpful to further processing. Therefore, our work can help reduce the workload for legal professionals and improve their work efficiency. Considering the fact that, like any other legal AI application, LED models would inevitably make mistakes and have negative influences, we argue that LED can only serve as an auxiliary tool for legal work and the final decision on a specific legal issue has to be ensured by legal professionals. In such case, we could exploit the advantage of legal AI and avoid the potential risk.

The corpus we use is released by the Chinese government and has been anonymized wherever necessary. Therefore, our dataset does not involve any personal privacy. In terms of human annotation, we first annotate a few examples on our own to approximate the workload and then determine the wages for annotators according to local standards.

Acknowledgement

References

Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet project. In *COLING 1998 Volume 1: The 17th International Conference on Computational Linguistics*.

Anderson Bertoldi, Rove Chishman, Sandro José Rigo, and Thaís Domênica Minghelli. 2014. Cognitive linguistic representation of legal events. *Proceedings of COGNITIVE*.

Michael J Bommarito II, Daniel Martin Katz, and Eric M Detterman. 2021. Lexnlp: Natural language processing and information extraction for legal and regulatory texts. In *Research Handbook on Big Data Law*. Edward Elgar Publishing.

Ilias Chalkidis, Ion Androutsopoulos, and Nikolaos Aletras. 2019. Neural legal judgment prediction in English. In *Proceedings of ACL*, pages 4317–4323, Florence, Italy. Association for Computational Linguistics.

Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. Legal-bert:"preparing the muppets for court'". In *Proceedings of EMNLP: Findings*, pages 2898–2904.

Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael James Bommarito, Ion Androutsopoulos, Daniel Martin Katz, and Nikolaos Aletras. 2021. Lexglue: A benchmark dataset for legal language understanding in english. *Available at SSRN 3936759*.

Xu Liheng Chen, Yubo, Kang Liu, Daojian Zeng, and Jun Zhao. 2015. Event extraction via dynamic multi-pooling convolutional neural networks. In *Proceedings of the ACL-IJNLP (Volume 1: Long Papers)*, pages 167–176, Beijing, China.

Yubo Chen, Shulin Liu, Xiang Zhang, Kang Liu, and Jun Zhao. 2017. Automatically labeled data generation for large scale event extraction. In *Proceedings of ACL (Volume 1: Long Papers)*, pages 409–419, Vancouver, Canada. Association for Computational Linguistics.

Walker Christopher, Strassel Stephanie, Medero Julie, and Maeda Kazuaki. 2006. Ace 2005 multilingual training corpus. *Linguistic Data Consortium, Philadelphia*, page 57.

Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.

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 NAACL-HLT 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Xiaocheng Feng, Lifu Huang, Duyu Tang, Heng Ji, Bing Qin, and Ting Liu. 2016. A language-independent neural network for event detection. In *Proceedings of ACL (Volume 2: Short Papers)*, pages 66–71, Berlin, Germany. Association for Computational Linguistics.

Dan Hendrycks, Collin Burns, Anya Chen, and Spencer Ball. 2021. Cuad: An expert-annotated nlp dataset for legal contract review. In *Proceedings of NeurIPS, Dataset*.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.

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

Yoshinobu Kano, Mi-Young Kim, Masaharu Yoshioka, Yao Lu, Juliano Rabelo, Naoki Kiyota, Randy Goebel, and Ken Satoh. 2018. Coliee-2018: Evaluation of the competition on legal information extraction and entailment. In *Proceedings of JSAI*, pages 177–192. Springer.

Phi Manh Kien, Ha-Thanh Nguyen, Ngo Xuan Bach, Vu Tran, Minh Le Nguyen, and Tu Minh Phuong. 2020. Answering legal questions by learning neural attentive text representation. In *Proceedings of ACL*, pages 988–998, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Jin-Dong Kim, Tomoko Ohta, Kanae Oda, and Jun'ichi Tsujii. 2008. From text to pathway: corpus annotation for knowledge acquisition from biomedical literature. In *Proceedings of APBC*, pages 165–175. World Scientific.

Yuta Koreeda and Christopher Manning. 2021. ContractNLI: A dataset for document-level natural language inference for contracts. In *Findings of EMNLP 2021*, pages 1907–1919, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Nikolaos Lagos, Frederique Segond, Stefania Castellani, and Jacki O'Neill. 2010. Event extraction for legal case building and reasoning. In *Proceedings of ICIIP*, pages 92–101. Springer.

Chuanyi Li, Yu Sheng, Jidong Ge, and Bin Luo. 2019. Apply event extraction techniques to the judicial field. In *Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers*, pages 492–497.

Qi Li, Heng Ji, and Liang Huang. 2013. Joint event extraction via structured prediction with global features. In *Proceedings of ACL (Volume 1: Long Papers)*, pages 73–82, Sofia, Bulgaria. Association for Computational Linguistics.

Qingquan Li, Qifan Zhang, Junjie Yao, and Yingjie Zhang. 2020a. Event extraction for criminal legal text. In *Proceedings of ICKG*, pages 573–580. IEEE.

Shen Li, Zhe Zhao, Renfen Hu, Wensi Li, Tao Liu, and Xiaoyong Du. 2018. Analogical reasoning on Chinese morphological and semantic relations. In *Proceedings of ACL(Volume 2: Short Papers)*, pages 138–143, Melbourne, Australia. Association for Computational Linguistics.

Xinyu Li, Fayuan Li, Lu Pan, Yuguang Chen, Weihua Peng, Quan Wang, Yajuan Lyu, and Yong Zhu. 2020b. Duee: A large-scale dataset for chinese event extraction in real-world scenarios. In *Proceedings of NLPCC*, pages 534–545. Springer.

Yubo Liu Kang Liu, Shulin Chen and Jun Zhao. 2017. Exploiting argument information to improve event detection via supervised attention mechanisms. In *Proceedings of ACL (Volume 1: Long Papers)*, pages 1789–1798, Vancouver, Canada.

Yixiao Ma, Yunqiu Shao, Yueyue Wu, Yiqun Liu, Ruizhe Zhang, Min Zhang, and Shaoping Ma. 2021. Lecard: A legal case retrieval dataset for chinese law system. In *Proceedings of SIGIR*.

K Tamsin Maxwell, Jon Oberlander, and Victor Lavrenko. 2009. Evaluation of semantic events for legal case retrieval. In *Proceedings of the WSDM: Workshop on Exploiting Semantic Annotations in Information Retrieval*, pages 39–41.

Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. 2016. Joint event extraction via recurrent neural networks. In *Proceedings of NAACL-HLT 2016*, pages 300–309, San Diego, California. Association for Computational Linguistics.

Abhilasha Ravichander, Alan W Black, Shomir Wilson, Thomas Norton, and Norman Sadeh. 2019. Question answering for privacy policies: Combining computational and legal perspectives. In *Proceedings of EMNLP-IJCNLP*, pages 4947–4958, Hong Kong, China. Association for Computational Linguistics.

Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of EMNLP-IJCNLP*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.

Alan Ritter, Mausam, Oren Etzioni, and Sam Clark. 2012. Open domain event extraction from twitter. In *The 18th ACM SIGKDD, Beijing, China, August 12-16, 2012*, pages 1104–1112. ACM.

Yunqiu Shao, Yueyue Wu, Yiqun Liu, Jiaxin Mao, Min Zhang, and Shaoping Ma. 2021. Investigating user behavior in legal case retrieval. In *Proceedings of SIGIR*.

Shirong Shen, Guilin Qi, Zhen Li, Sheng Bi, and Lusheng Wang. 2020. Hierarchical Chinese legal event extraction via pedal attention mechanism. In *Proceedings of COLING*, pages 100–113, Barcelona, Spain. International Committee on Computational Linguistics.

Zhiyi Song, Ann Bies, Stephanie Strassel, Tom Riese, Justin Mott, Joe Ellis, Jonathan Wright, Seth Kulick, Neville Ryant, and Xiaoyi Ma. 2015. From light to rich ERE: Annotation of entities, relations, and events. In *Proceedings of the The 3rd Workshop on EVENTS*, pages 89–98, Denver, Colorado. Association for Computational Linguistics.

Paul Thompson, Syed A Iqbal, John McNaught, and Sophia Ananiadou. 2009. Construction of an annotated corpus to support biomedical information extraction. *BMC bioinformatics*, 10(1):1–19.

Dimitrios Tsarapatsanis and Nikolaos Aletras. 2021. On the ethical limits of natural language processing on legal text. In *Findings of ACL 2021*, pages 3590–3599, Online. Association for Computational Linguistics.

Xiaozhi Wang, Xu Han, Zhiyuan Liu, Maosong Sun, and Peng Li. 2019. Adversarial training for weakly supervised event detection. In *Proceedings of NAACL-HLT*, pages 998–1008, Minneapolis, Minnesota. Association for Computational Linguistics.

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 EMNLP*, pages 1652–1671, Online. Association for Computational Linguistics.

Yuzhong Wang, Chaojun Xiao, Shirong Ma, Haoxi Zhong, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2021a. Equality before the law: Legal judgment consistency analysis for fairness. *ArXiv preprint*, abs/2103.13868.

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 ACL-IJNLP (Volume 1: Long Papers)*, pages 6283–6297, Online. Association for Computational Linguistics.

Chaojun Xiao, Xueyu Hu, Zhiyuan Liu, Cunchao Tu, and Maosong Sun. 2021. Lawformer: A pre-trained language model for chinese legal long documents. *AI Open*.

Chaojun Xiao, Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Zhiyuan Liu, Maosong Sun, Yansong Feng, Xianpei Han, Zhen Hu, Heng Wang, et al. 2018. Cail2018: A large-scale legal dataset for judgment prediction. *ArXiv preprint*, abs/1807.02478.

Haoran Yan, Xiaolong Jin, Xiangbin Meng, Jiafeng Guo, and Xueqi Cheng. 2019. Event detection with multi-order graph convolution and aggregated attention. In *Proceedings of EMNLP-IJCNLP*, pages 5766–5770, Hong Kong, China. 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.

Wenmian Yang, Weijia Jia, Xiaojie Zhou, and Yutao Luo. 2019. Legal judgment prediction via multi-perspective bi-feedback network. In *Proceedings of IJCAI 2019, Macao, China, August 10-16, 2019*, pages 4085–4091. ijcai.org.

Hai Ye, Xin Jiang, Zhunchen Luo, and Wenhan Chao. 2018. Interpretable charge predictions for criminal cases: Learning to generate court views from fact descriptions. In *Proceedings of NAACL-HLT, Volume 1 (Long Papers)*, pages 1854–1864, New Orleans, Louisiana. Association for Computational Linguistics.

Shuo Zhang, Junzhou Zhao, Pinghui Wang, Nuo Xu, Yang Yang, Yiting Liu, Yi Huang, and Junlan Feng. 2021. Learning to check contract inconsistencies. In *Proceedings of the AAAI*, volume 35, pages 14446–14453.

Yue Zhao, Xiaolong Jin, Yuanzhuo Wang, and Xueqi Cheng. 2018. Document embedding enhanced event detection with hierarchical and supervised attention. In *Proceedings of ACL (Volume 2: Short Papers)*, pages 414–419, Melbourne, Australia. Association for Computational Linguistics.

Shun Zheng, Wei Cao, Wei Xu, and Jiang Bian. 2019. Doc2EDAG: An end-to-end document-level framework for Chinese financial event extraction. In *Proceedings of EMNLP-IJCNLP*, pages 337–346, Hong Kong, China. Association for Computational Linguistics.

Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Chaojun Xiao, Zhiyuan Liu, and Maosong Sun. 2018. Legal judgment prediction via topological learning. In *Proceedings of EMNLP*, pages 3540–3549, Brussels, Belgium. Association for Computational Linguistics.

Haoxi Zhong, Chaojun Xiao, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2020a. How does NLP benefit legal system: A summary of legal artificial intelligence. In *Proceedings of ACL*, pages 5218–5230, Online. Association for Computational Linguistics.

Haoxi Zhong, Chaojun Xiao, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. 2020b. JEC-QA: A legal-domain question answering dataset. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020*, pages 9701–9708. AAAI Press.

A Implementation Details

In this section, we introduce the hyper-parameters of the baseline models, models for legal judgment prediction, and models for similar case retrieval.

A.1 Baseline Models

For all baseline models, we run the model 5 times to get stable results, and the average performance is reported. For each model, we choose the checkpoint with the best performance on the validation set to evaluate on the test set. We train these models on GeForce RTX 2080Ti GPUs, and use Adam to optimize the models. The validation performances are shown in Table 7.

Model	Micro			Macro		
	P	R	F1	P	R	F1
DMCNN	86.15	79.27	82.57	79.42	69.77	73.00
BiLSTM	83.01	84.30	83.65	78.45	73.39	74.27
BiLSTM+CRF	84.63	83.10	83.86	80.99	73.39	75.73
BERT	84.35	83.80	84.07	80.21	76.08	77.38
BERT+CRF	83.72	84.13	83.93	78.38	75.39	76.01
DMBERT	83.40	86.76	85.05	79.18	79.28	78.42

Table 7: Validation Performances of ED baselines.

For **DMCNN**, the hyper-parameters are the same as Chen et al. (2015), excluding the unmentioned dimension of word embedding and learning rate. We use JIEBA toolkit[4] to perform the Chinese word segmentation, and use the pre-trained word vectors released in Li et al. (2018). Specifically, the word embedding is the one trained by the Wikipedia-zh corpus with word, character, and N-gram context. The training parameters are shown in Table 8.

For **BiLSTM-based models**, we also use JIEBA to perform word segmentation, and adopt the same pre-trained Chinese word vectors used in DMCNN. The detailed training hyper-parameters are shown in Table 9.

For **BERT-based models**, we adopt BERT-base as the basic encoder with the bert-base-chinese

[4]https://github.com/fxsjy/jieba

Batch Size	170
Dropout Rate	0.5
Learning Rate	1×10^{-3}
Kernel Size	3
Hidden Size	200
Dimension of PF	5
Dimension of Word Embedding	300

Table 8: Hyper-parameters for DMCNN.

Batch Size	200
Dropout Rate	0.5
Learning Rate	1×10^{-3}
Kernel Size	3
Hidden Size	256
Dimension of Word Embedding	300

Table 9: Hyper-parameters for BiLSTM-based models.

checkpoint[5]. For BERT, BERT-CRF and DMBERT, the training hyper-parameters are shared and the training hardwares are $4\times$ RTX 2080TI. The detailed hyper-parameters are shown in Table 10.

Batch Size	64
Dropout Rate	0.5
Adam ϵ	1×10^{-8}
Learning Rate	5×10^{-5}
Validation Steps During Training	500

Table 10: Hyper-parameters for BERT-based models.

For **CRF-based models**, we use BIO tagging schema for training and evaluation.

A.2 Downstream Applications

For all downstream application experiments, we first adopt the **BERT+CRF** model to detect the trigger words in the original downstream datasets. We add an extra **Event Type Embedding Layer** to incorporate the event information, where the embedding matrix with the shape of 109×768 is randomly initialized and updated during training.

For **Legal Judgment Prediction (LJP)** task, we adopt the dataset released in the first stage of CAIL2018[6] as the benchmark. We evaluate the LJP task with the detected events under both full data and low-resource settings. The data size is shown in Table 11. The training data for the low-resource setting is obtained by randomly selecting 50 samples for each label and the corresponding data we used is also released in the github repository. The experiments are based on the source code released in Zhong et al. (2020a).

[5]https://huggingface.co/bert-base-chinese
[6]https://github.com/china-ai-law-challenge/CAIL2018

Setting	Training	Validation	Test
Full-data	154,592	17,131	32,508
Low-resource	12,702		

Table 11: Statistics of Data for LJP Experiments.

For **Similar Case Retrieval (SCR)**, we use LeCaRD[7] as benchmark and implement the models based on the code released in Xiao et al. (2021). As the documents in LeCaRD are usually longer than 512, we truncate the text to feed into the encoder. Specifically, the maximum lengths we use for the query and candidates are 100 and 409, respectively. We also conduct experiments in the unsupervised setting, where we use a 108-dimension vector as the representation of a document and each entry of the vector is the number of the detected events.

B Performance on Different Top-level Types

Top-level Event Type	precision	recall	F1
General_behaviors	83.71	85.67	84.86
Prohibited_acts	83.01	82.93	82.97
Judicature_related	94.17	91.89	93.01
Consequences	84.54	82.92	83.73
Accident	86.04	84.40	85.21
Natural_disaster	77.78	63.64	70.00

Table 12: Performance of DMBERT on different top-level event types.

In this section, we further analyze the performance of DMBERT on different top-level types to explore the fine-grained performance on LEVEN. From the results, we can observe that as the Judicature_related events are usually described in legal terminologies, thus the models can easily identify trigger words correctly by memorizing a set of specific words. As the Natural_disaster only contains tens of event instances, the model cannot be well-trained for these event types.

C Performance on Long-tail Event Types

In this section, we further analyze the performance on long-tail event types in detail. From Table 13, we can observe that overall, the model performance decreases as the number of training instances decreases. But there are some exceptions. Some

event types only contain limited instances while the model can achieve high F1-scores on these types. For instances, Suicide only contains 55 event mentions, but the model can achieve 95.65% F1-score due to its non-diverse expression. Though some long-tail event types can be predicted accurately, there are 9 long-tail event types that can only reach F1-scores lower than 0.6. Therefore, we argue that detecting the event types accurately with limited instances needs future efforts.

F1-score	[0,0.4)	[0.4,0.6)	[0.6,0.8)	[0.8,0.9)	[0.9,1.0]	sum
#low-freq.	5	4	4	4	4	21
#mid-freq.	0	0	9	13	6	28
#high-freq.	0	0	14	23	22	59

Table 13: Distribution of event types by their performance on the test set. Here, low-freq and high-freq represent the number of event types that have less than 150 event mentions and more than 500 event mentions. And mid-freq denotes the number of event types containing between 150 and 500 event mentions.

D Data Distribution

To help the following researchers to better understand the features and details of LEVEN, we present more data analysis in this section regarding multiple events in one sentence and the sentence length distribution.

Number of Events in One Sentence. Legal cases usually involve complicated facts, and it is common in LEVEN that there are multiple events mentioned in one sentence. Table 14 shows the percentage of sentences containing different numbers of events.

#Event/Sent.	0	1	[2,5)	[5,10)	[10,100)
Percentage (%)	12.8	26.7	47.9	11.6	1.0

Table 14: The percentage of sentences containing different numbers of sentences.

Length&Number of Sentences. LEVEN is constructed based on real-world corpus, which makes it a perfect resource for developing practical applications. Figure 3 and 4 exhibit a comparison between the sentence length and number distributions of LEVEN and CAIL2018(Xiao et al., 2018), which is the largest legal judgment prediction dataset with over 1.7 million criminal judgment documents and can serve as a good real-world reference, indicating that LEVEN is consistent with the reality.
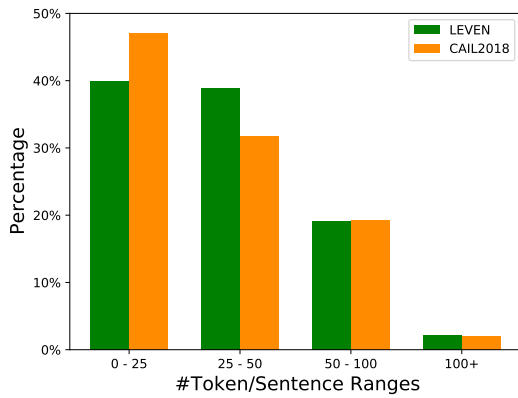
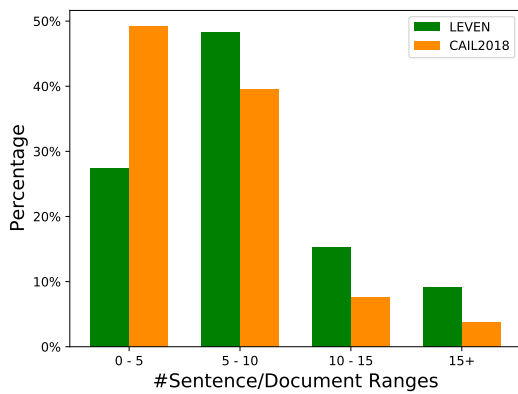Figure 3: Sentence length distributions of LEVEN and CAIL2018.

Figure 4: Sentence number distributions of LEVEN and CAIL2018.

E Event Type Schema and Description

To promote future research, we provide the hierarchical event schema in Figure 5, and the list of event types, including the event names and the corresponded descriptions, in Table 15, 16, 17, and 18.

F Annotation Guidelines

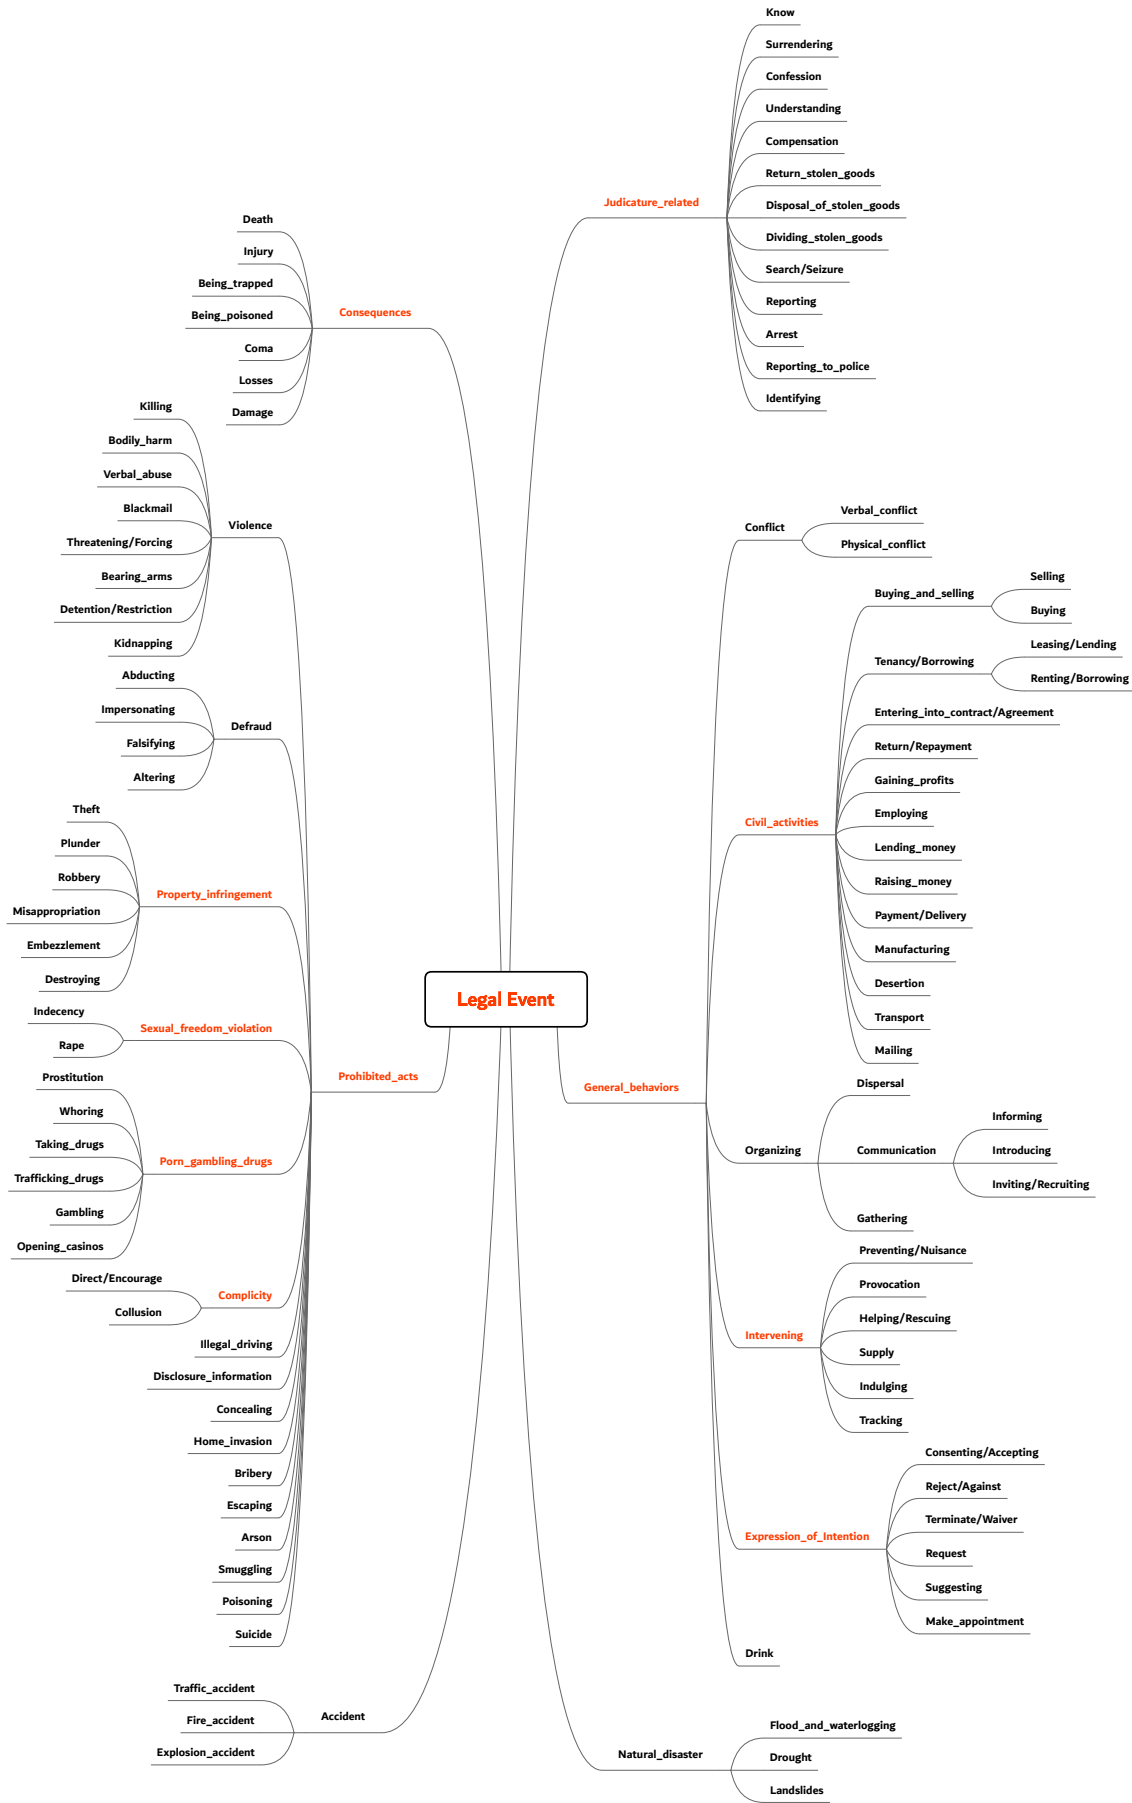The annotation guidelines can be obtained from https://github.com/thunlp/LEVEN.

Figure 5: The detailed event schema of our proposed LEVEN.

Event Type Name	Descriptions
	Judicature Related Events
Judicature_related	JUDICATURE_RELATED events mainly refer to the activities of judicial organs or some legal penalty circumstances.
Know	A KNOW event means the doer ought to know the fact or understand the fact clearly.
Surrendering	A SURRENDERING event refers to the doer voluntarily surrendering after committing a crime.
Confession	A CONFESSION event refers to the suspect or defendant telling the facts to the police.
Understanding	AN UNDERSTANDING event refers to the forgiveness from the victim or victim's families to the criminal.
Compensation	A COMPENSATION event refers to the act of compensating the victim for his loss, damage, or injury.
Return_stolen_goods	A TERURN_STOLEN_GOODS event refers to the act of returning the stolen money or stolen goods to the victim or government.
Disposal_of_stolen_goods	A DISPOSAL_OF_STOLEN_GOODS event refers to the act of destroying stolen goods, selling stolen goods, or squandering stolen money, that is, the stolen goods/money have been disposed of.
Dividing_stolen_goods	A DICIDING_STOLEN_GOODS event refers to the act of sharing stolen goods or money.
Search/Seizure	A SEARCH/SEIZURE event mainly refers to the search and inspection of the suspect's body, articles, residence, or other space by the reconnaissance personnel, or the seizure of contraband, including the seizure of real estate. However, illegal search or seizure by non-reconnaissance personnel can also mark this event.
Reporting	A REPORTING event refers to the act of reporting bad people or bad things to relevant units.
Arrest	AN ARREST event refers to the act of detaining or arresting suspects.
Reporting_to_police	A REPORTING_TO_POLICE event refers to the act of calling the police to ask for help or reporting a case to the police.
Identifying	AN IDENTIFYING event refers to a kind of behavior in which the investigation organ appoints or hires people with expertise to make a scientific judgment and draw professional conclusions on the specialized problems in criminal cases in order to solve the specialized problems in criminal cases.
	Accident Events
Accident	AN ACCIDENT event refers to accidental loss or disaster.
Traffic_accident	A TRAFFIC_ACCIDENT event occurs when a traffic accident happens, which usually causes personal injury, death or property loss.
Fire_accident	A FIRE_ACCIDENT event refers to the disaster caused by uncontrolled combustion.
Explosion_accident	AN EXPLOSION_ACCIDENT event refers to the disaster caused by a sudden release of a large amount of energy, which leads to property losses and personal casualties.
	Natural Disaster Events
Natural_disaster	NATURAL_DISASTER events refer to Natural phenomena or man-made influences that endanger human survival or damage the human living environment.
Flood_and_waterlogging	A FLOOD_AND_WATERLOGGING event occurs where a large amount of water covers an area that is usually dry.
Drought	A DROUGHT event occurs when there is little or no rain during a long period of time.
Landslides	A LANDSLIDES event refers to a geographic disaster caused by a mass of earth or rock falling down the slope of a mountain.
	Consequence Events
Consequence	CONSEQUENCE events contain the fact of damage to the object caused by harmful acts.
Death	A DEATH event refers to the state of a human being dead.
Injury	AN INJURY event refers to the fact of personal injury.
Being_trapped	A BEING_TRAPPED event means the state in which people are physically in trouble and can't get out.
Being_poisoned	A BEING_POISONED event refers to one's discomfort caused by toxic effects, emphasizing the state of one's being poisoned.
Coma	A COMA event refers to the state of one's unconsciousness.
Losses	A LOSSES event refers to the fact of property loss.
Damage	A DAMAGE event refers to the fact that the property has been damaged.
	General Behavior Events (I)
General_behavior	GENERAL_BEHAVIOR events contain common behaviors in daily life, which usually do not violate laws.
Conflict	A CONFLICT event refers to two or more parties having verbal, physical, or other conflicts, disputes, or contradictions.

Table 15: Event type list (I), including the event type names and the corresponded descriptions.

Event Type Name	Descriptions
	General Behavior Events (II)
Verbal_conflict	A VERBAL_CONFLICT event refers to oral conflicts happen between two or more people without physical contact.
Physical_conflict	A PHYSICAL_CONFLICT event refers to a physical clash that happens between two or more people, including fighting. This event emphasizes the mutual behavior of both parties, pay attention to distinguish this event from a BODILY_HARM event, which emphasizes that one hurts another.
Civil_activities	CIVIL_ACTIVITIES events contain typical activities in civil and commercial areas.
Buying_and_selling	A BUYING_AND_SELLING event refers to the act of transacting within or between groups, including the exchange of goods and online transactions.
Selling	A SELLING event refers to one's act of selling something for a profit.
Buying	A BUYING event refers to one's act of buying or consuming something.
Tenancy/Borrowing	A TENANCY/BORROWING event refers to the relationship between two groups/persons to lease or rent something.
Leasing/Lending	A LEASING/LENDING event refers to the act of renting or lending something to others.
Renting/Borrowing	A RENTING/BORROWING event refers to the act of renting or borrowing something from others.
Return/Repayment	A RETURN/REPAYMENT refers to the act of returning something to its original place or owner.
Gaining_profits	A GAINING_PROFITS event refers to one obtaining money or other benefits through a certain act or activity.
Employing	AN EMPLOYING event refers to the act of giving others a job to do for payment.
Lending_money	A LENDING_MONEY event refers to specialized institutions or people making loans to earn profits, including bank loans and individual loans.
Raising_money	A RAISING_MONEY event refers to the act of raising money from unspecified majority people.
Payment/Delivery	A PAYMENT/DELIVERY event refers to the act of giving money or other things to others.
Entering_into_contract/agreement	AN ENTERING_INTO_CONTRACT/AGREEMENT event refers to the act of two or more person/groups signing contracts, including written contracts, written agreements, oral agreements, etc.
Manufacturing	A MANUFACTURING event refers to producing, manufacturing, or making tangible objects, emphasizing from scratch, excluding "noise", "explosion" or other intangible objects.
Desertion	A DESERTION event refers to one's act of actively abandoning or discarding something or someone.
Transport	A TRANSPORT event refers to one's act of transporting someone or something from one place to another.
Mailing	A MAILING event refers to delivering documents or articles through the post office or third-party postal service.
Organizing	AN ORGANIZING event refers to the act of arranging scattered people or things to serve a common goal.
Dispersal	A DISPERSAL event refers to the act of spreading information, data, rumors to the unspecified majority of people on the Internet or in public.
Communication	A COMMUNICATION event generally refers to the connection between two or more people, such as making a phone call.
Informing	AN INFORMING event refers to one's act of telling others information or reminding others of certain information, or the notified one should not have known the information.
Introducing	AN INTRODUCING refers to one's behavior to make other people or groups know each other or have a connection, excluding product instructions (because the introduction here does not mean "intermediary", but just a kind of teaching).
Inviting/Recruiting	AN INVITING/RECRUITING event refers to the acts of recruiting, inviting others to a place, or inviting others to do something or participate in an activity.
Gathering	A GATHERING event refers to the act of gathering a group of people together.
Intervening	AN INTERVENING event refers to one's act of intervening in an ongoing event.
Preventing/Nuisance	A PREVENTING/NUISANCE event refers to one's act of preventing things from going smoothly or hindering others from doing something by words or actions.
Provocation	A PROVOCATION event refers to one attempting to trigger off conflicts with others, or trigger off conflicts between other two groups.

Table 16: Event type list (II), including the event type names and the corresponded descriptions.

Event Type Name	Descriptions
General Behavior Events (III)	
Helping/Rescuing	A HELPING/RESCUING event refers to one's act of helping others to do something in the process of life, work or crime, it is limited to behavioral help, excluding providing materials, suggestions, etc. A HELPING/RESCUING event also refers to one's act of saving, rescuing, or assisting others who are injured or in trouble.
Supply	A SUPPLY event refers to one providing materials, conditions, intelligence information, or other specific things to others, excluding abstract things such as "providing help" or "providing advice".
Indulging	AN INDULGING event refers to one's act of allowing bad things to develop without any interference.
Tracking	A TRACKING event refers to one's act of following others quietly without being detected.
Expression_of_Intention	EXPRESSION_OF_INTENTION events contain the acts of one expressing a certain intention in a verbal way.
Consenting/Accepting	A CONSENTING/ACCEPTING event refers to one agreeing with the opinions of others, accepting others' asks, or accepting the property given by others.
Reject/Against	A REJECT/AGAINST event refers to one rejecting others' asks or the property given by others.
Terminate/Waiver	A TERMINATE/WAIVER event refers to one stopping doing something, giving up the original persistence, or giving up a right.
Request	A REQUEST event refers to one putting forward specific matters or wishes, hoping or requiring others to realize them.
Suggesting	A SUGGESTING event refers to one putting forward a plan or idea to others.
Make_appointment	A MAKE_APPOINTMENT event refers to the act of two or more people discussing and determining something.
Drink	A DRINK event refers to one's act of drinking alcohol, usually accompanied by other behaviors, such as driving, etc.
Prohibited Acts Events (I)	
Prohibited_acts	PROHIBITED_ACTS events contain behaviors prohibited by law, including not only typical criminal behaviors, but also behaviors that are not up to the degree of crime but prohibited by law. Therefore, events in this part are events that should be given negative evaluation, which is opposite to general behaviors.
Violence	VIOLENCE events contain violent behaviors that are intended to hurt others' mental or physical health, including physical force as well as language.
Killing	A KILLING event refers to one's act of killing others in order to make others die.
Bodily_harm	A BODILY_HARM event refers to the act of harming the physical health of others, usually manifested in beating.
Verbal_abuse	A VERBAL_ABUSE event refers to the act of insulting, attacking or hurting others through language. Pay attention to distinguishing this event from a VERBAL_CONFLICT event, which emphasizes mutual abuse.
Blackmail	A BLACKMAIL event refers to the act of demanding money from others by threatening or deceiving them.
Threatening/Forcing	A THREATENING/FORCING event refers to the act of forcing others to do or not do something through violence or power, mostly referring to the use of force to make others obey.
Bearing_arms	A BEARING_ARMS event refers to one's holding or carrying sticks, props, guns, or other instruments.
Detention/restriction	A DETENTION/RESTRICTION event refers to the act of depriving or restricting the freedom of others, such as binding or detaining people in specific places.
Kidnapping	A KIDNAPPING event refers to the act of taking hostages by violent means in exchange for interests, emphasizing that the object must be people.
Defraud	A DEFRAUD event refers to the act of covering up the real situation with false words or actions to deceive others.
Abducting	AN ABDUCTING event refers to one's act of cheating someone away by luring, cheating, or other means.
Impersonating	AN IMPERSONATING event refers to the act of disguising a real thing with a false thing or one's act of pretending to be somebody in order to trick people.
Falsifying	A FALSIFYING event refers to the act of making fake goods or false news.
Altering	AN ALTERING event refers to the act of modifying real basis A without authorization to make it have another illusion B.
Property Infringement	PROPERTY INFRINGEMENT events contain acts of infringing upon others' property rights and interests of others.
Theft	A THEFT event refers to one's act of stealing others' property by secret.
Plunder	A PLUNDER event refers to one's act of seizing property blatantly in front of the victims and taking them away, including seizing guns or knives, excluding competing for customers or land rights. The object of robbery must be tangible things.

Table 17: Event type list (III), including the event type names and the corresponded descriptions.

Event Type Name	Descriptions
	Prohibited Acts Events (II)
Robbery	A ROBBERY event refers to one's act of using violent means to rob others' property, such as robbery with a knife. The establishment of this event is strict. If it is impossible to judge whether it is a ROBBERY event, then PLUNDER may be marked.
Misappropriation	A MISAPPROPRIATION event refers to the act of changing the original use of the property to another without authorization.
Embezzlement	AN EMBEZZLEMENT event refers to one's act of taking others' property illegally, including real estate, emphasizing the state of possession.
Destroying	A DESTROYING event refers to one's act of destroying property, this event has a subject, which is the main difference against A DAMAGE event.
Sexual_freedom_violation	SEXUAL_FREEDOM_VIOLATION events contain acts of making others unable to freely dispose of their sexual rights by means of inducement, deception, coercion, etc.
Indecency	AN INDECENCY event refers to one's act of forcibly sexually harassing others by touching private parts or other acts other than adultery.
Rape	A RAPE event refers to one's act of forcing women to have sex when they do not want to.
Porn_gambling_drugs	PORN_GAMBLING_DRUGS events contain illegal or criminal phenomena involving pornography, gambling, and drugs.
Prostitution	A PROSTITUTION event refers to women providing paid sexual services to others.
Whoring	A WHORING event refers to one purchasing sexual service with money.
Taking_drugs	A TAKING_DRUGS event refers to one's act of taking drugs.
Trafficking_drugs	A TRAFFICKING_DRUGS event refers to one's act of peddling drugs.
Gambling	A GAMBLING event refers to one's act of gambling.
Opening_casinos	AN OPENING_CASINOS event refers to one's act of opening casinos for multiple plays to gamble on.
Complicity	COMPLICITY events occur when intentional contacts happen between two or more criminals.
Direct/Encourage	A DIRECT/ENCOURAGE refer to one's act of letting others commit crimes by means of command, inspiration, or temptation. Specifically, a DIRECT event refers to the act of summoning others to commit criminal acts or other negative acts according to the instigator's intention. AN ENCOURAGE event refers to one's act of making people who do not have criminal intention have the intention of committing a crime.
Collusion	A COLLUSION event refers to the act of two or more people scheming a crime plan together.
Illegal_driving	AN ILLEGAL_DRIVING event refers to one's act of driving a car illegally.
Disclosure_information	A DISCLOSURE_INFORMATION event refers to one's act of disclosing information that should be kept secret.
Concealing	A CONCEALING event refers to one's act of hiding something from discovery.
Home Invasion	A HOME INVASION event refers to one's act of invading or sneaking into other people's private space without the permission of others. This event is usually the pre-act of another criminal act (such as theft or rape).
Bribery	A BRIBERY event refers to one's act of bribing others with property to seek illegitimate interests or accepting others' property to seek illegitimate interests for others.
Escaping	AN ESCAPING event means one's escaping and hiding in order to avoid capture.
Arson	AN ARSON event refers to one's act of setting on fire.
Smuggling	A SMUGGLING event refers to the act of one's illegally transporting goods into or out of the country in violation of customs regulations.
Poisoning	A POISONING event refers to one putting poison in containers or a specific environment in order to kill people, animals, or plants.
Suicide	A SUICIDE event refers to the act of one's killing himself.

Table 18: Event type list (IV), including the event type names and the corresponded descriptions.

201