

EXCEEDS: Extracting Complex Events via Nugget-based Grid Modeling in Scientific Domain

Anonymous ACL submission

Abstract

It is crucial to understand a specific domain by events. Extensive event extraction research has been conducted in many domains such as news, finance, and biology. However, event extraction in scientific domain is still insufficiently supported by comprehensive datasets and tailored methods. Compared with other domains, scientific domain has two characteristics: (1) denser nuggets and events, and (2) more complex information forms. To solve the above problem, considering these two characteristics, we first construct SciEvents, a large-scale multi-event document-level dataset with a schema tailored for scientific domain. It consists of 2,508 documents and 24,381 events under multi-stage manual annotation and quality control. Then, we propose EXCEEDS, an end-to-end scientific event extraction framework by encoding dense nuggets into a grid matrix and simplifying complex event extraction as a nugget-based grid modeling task. Experiments on SciEvents demonstrate state-of-the-art performances of EXCEEDS. Both the SciEvents dataset and the EXCEEDS framework will be released publicly to facilitate future research.

1 Introduction

Event extraction (EE) is a fundamental information extraction task aiming to extract structural event knowledge from plain texts (Peng et al., 2023). It is typically decomposed into two pipeline subtasks: event detection (ED) and event argument extraction (EAE). Specifically, ED identifies a word span (hereafter referred to as a **nugget**) that most clearly refers to the occurrence of an event, *i.e.*, event trigger, and also detects the event type evoked by the event trigger (Pouran Ben Veyseh et al., 2022). Given an event trigger and its event type, EAE further identifies nuggets as event arguments and classifies their roles in the event.

EE provides an effective abstraction for representing domain knowledge and supports down-

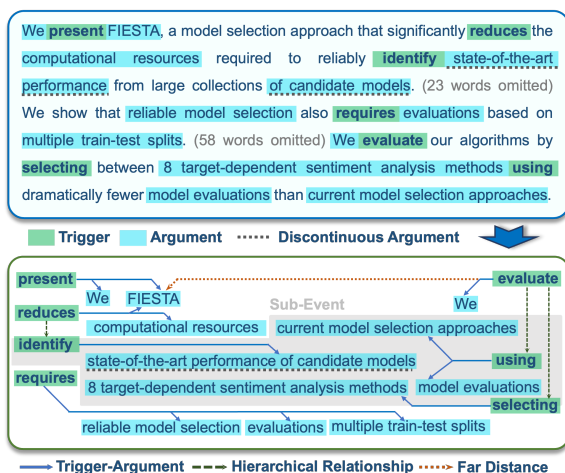


Figure 1: A real example from SciEvents. The upper panel displays a scientific paper abstract, and the lower panel shows the extracted events, highlighting the dense information and complex event structures characteristic of scientific texts.

stream tasks such as reasoning, summarization, and knowledge discovery. Consequently, extensive EE research has been conducted in various scenarios and domains such as internet news, radio conversations, internet blogs (Sundheim, 1992; Aguilar et al., 2014; Ebner et al., 2020), business (Yang et al., 2018; Liang et al., 2020), biology (Kim et al., 2011; Pyysalo et al., 2012), legislation (Shen et al., 2020; Yao et al., 2022), cybersecurity (Man Duc Trong et al., 2020; Satyapanich et al., 2020) and so on. Despite this progress, scientific literature has been growing rapidly in recent decades, with millions of new publications released every year. Such growth poses an urgent challenge for managing scientific domain knowledge, calling for effective EE solutions.

However, EE in scientific domain remains insufficiently characterized by existing datasets and methods. In particular, current resources and formulations struggle to capture two salient characteristics of scientific texts. **First**, compared with other

domains, scientific domain tends to contain **more complex information forms**. Although many EE methods are task-specialized and rely on domain-specific ontologies (Lu et al., 2022a), these ontologies typically adopt flat tabular schema, which (1) neglects the hierarchical structure of events, (2) restricts the continuity of arguments, and (3) complicates the coreference problem, while these complex information forms are common in scientific literature. For example in Figure 1, (1) the trigger *evaluate* and the trigger *using* form a hierarchical relationship; (2) the trigger *identity* connects a discontinuous argument *state-of-the-art performance of candidate models*; (3) the trigger *evaluate* connects an argument *FIESTA* with a far distance from itself. These three examples demonstrate complex nuggets and events in scientific domain. **Second**, scientific texts, especially literature abstracts, tend to contain **denser nuggets and events** (with statistical evidence presented in Section 3.2). Unlike many existing datasets that focus on sentence-level extraction or single-event documents, scientific domain EE requires modeling dense, document-level multi-event interactions (see Figure 1 and Section 3.2). Together, these two characteristics motivate the need for dedicated EE resources and methods for scientific domain.

To explore the unique characteristics of scientific domain EE, we first introduce SciEvents, a large-scale document-level multi-event dataset tailored for scientific literature. SciEvents contains 2,508 manually annotated abstracts with 24,381 events, and is constructed with a refined schema designed to capture dense and structurally complex event patterns. Dataset statistics show that SciEvents exhibits both **denser nugget distributions** and **more complex event structures** than existing domain-specific EE datasets, reflecting the information-intensive nature of scientific texts.

Building on these observations, we further propose **EXCEEDS**, an end-to-end framework to **EXtract Complex Events** via nuggEt-based grID modeling in Scientific domain. EXCEEDS represents pairwise token relations across the entire document in a word-word event grid, enabling unified modeling of dense multi-event contexts as well as complex nugget and event structures, including hierarchical relations, discontinuous arguments, and long-distance dependencies. This formulation allows EXCEEDS to effectively address the challenges posed by scientific texts under a single end-to-end framework.

We evaluate state-of-the-art and recent EE methods on SciEvents under an extended evaluation protocol that incorporates an event correlation metric for hierarchical EE. Experimental results show that EXCEEDS achieves consistently strong performance across tasks, while further analysis reveals that complex nugget structures, especially under dense scientific contexts, remain challenging for existing models.

In summary, our contributions are two-fold: (1) We introduce SciEvents, a large-scale document-level EE dataset for the scientific domain with a refined schema, providing a comprehensive benchmark for studying dense and complex scientific events. (2) We propose EXCEEDS, an end-to-end EE framework tailored to the challenges of density and structural complexity in scientific texts, which achieves state-of-the-art performance on SciEvents.

2 Related Works

Event Extraction Datasets EE datasets have been constructed across a wide range of domains. Early efforts mainly focus on the news domain, describing events in realistic scenarios (Walker et al., 2006; Aguilar et al., 2014; Song et al., 2015). General domain datasets are further built from diverse sources like Wikipedia (Wang et al., 2020; Li et al., 2021; Tong et al., 2022), Reddit (Ebner et al., 2020), Baidu news (Li et al., 2020c), FrameNet (Parekh et al., 2023) and multi-lingual candidate data (Pouran Ben Veyseh et al., 2022). In addition, domain-specific datasets have been developed for finance (Yang et al., 2018; Liu et al., 2019a), biomedicine and related fields (Kim et al., 2011; Sun et al., 2022; Ma et al., 2023), as well as other domains such as cybersecurity, law, and literature (Man Duc Trong et al., 2020; Shen et al., 2020; Sims et al., 2019). Despite these advances, event datasets for the scientific domain remain limited, and existing resources rarely analyze the characteristics of scientific texts. In this work, we systematically examine the characteristics of scientific abstracts and construct SciEvents, a document-level EE dataset tailored to the scientific domain.

Event Extraction Approaches EE has evolved from early sequence labeling methods to more advanced neural architectures. To jointly model heterogeneous elements in EE datasets (Peng et al., 2023), early work focused on joint extraction frameworks that capture dependencies within and across events (Nguyen et al., 2016; Liu et al., 2018;

Dataset	Domain	#Docs	#Tokens	#ETs	#Events	#ATs	#Arguments
ACE2005 (Doddington et al., 2004)	News	599	297,842	33	5,348	22	8,097
Genia2011 (Kim et al., 2011)	Biomedical	1,375	345,371	9	13,537	10	11,865
Genia2013 (Kim et al., 2013)	Biomedical	20	149,856	13	6,001	7	5,660
RAMS (Ebner et al., 2020)	News	9,124	1,218,622	139	9,124	65	21,237
CASIE (Satyapanich et al., 2020)	Cybersecurity	999	387,275	5	8,479	26	22,679
M2E2 (Li et al., 2020b)	Multimedia	6,013	169,990	8	1,105	15	1,659
WikiEvents (Li et al., 2021)	News	246	189,718	50	3,951	59	5,536
PHEE (Sun et al., 2022)	Drug Safety	4,827	106,447	2	5,019	16	25,760
Maccrobat-EE (Ma et al., 2023)	Clinical	200	107,130	14	13,128	22	8,599
SciEvents (Ours)	Science	2,508	439,890	10	24,381	20	56,411

Table 1: Basic statistics of widely-used domain-specific event datasets. This table only presents publicly available event datasets that include argument annotations. #ETs: number of event types. #ATs: number of argument types.

Yang et al., 2019; Nguyen et al., 2021; Tang et al., 2022; Lin et al., 2020). Subsequent studies reformulated EE as machine reading comprehension, enabling more flexible trigger and argument extraction via question answering (Chen et al., 2019; Du and Cardie, 2020a; Li et al., 2020a; Zhou et al., 2021; Wei et al., 2021). More recent approaches adopt sequence-to-structure generation with Transformer-based models, unifying ED and EAE within a single framework (Lu et al., 2021, 2022b; Lou et al., 2023; Wang et al., 2023a; Liu et al., 2022; Yang et al., 2024). With the emergence of large language models (LLMs), EE has further benefited from strong generalization and zero-shot capabilities (Wei et al., 2023; Gao et al., 2023a; Wang et al., 2023b; Sainz et al., 2023; Gao et al., 2023b; Li et al., 2024). However, many existing methods still struggle to model the compositional diversity of nuggets, particularly under complex scenarios involving discontinuous, overlapping, or reverse-ordered nuggets. In this paper, we introduce an end-to-end EE framework, EXCEEDS, tailored for complex nugget and event structures.

3 The SciEvents Dataset

To support systematic research on scientific EE, we construct SciEvents, a large-scale document-level EE dataset tailored for scientific literature. In this section, we will introduce the dataset construction process in Section 3.1, and present a comprehensive statistical analysis in Section 3.2.

3.1 Dataset Construction Process

Scientific abstracts are typically organized around four rhetorical components: *background*, *related work*, *methodology*, and *results*. Motivated by this regularity, we design an event schema comprising 10 event types that cover these components. For in-

stance, abstracts often summarize prior approaches and highlight their limitations; we capture such information using the *RelatedWorkStep* and *RelatedWorkFault* event types, respectively.

For each event type, we define a set of argument roles to encode the information that readers typically seek in scientific abstracts. To ensure both coverage and annotation feasibility, we develop the schema iteratively with domain experts, two professors and three senior Ph.D. students in computer science annotate a set of seed documents and revise the schema over four rounds. The final schema and detailed definitions of all event types and argument types are provided in Appendix E. We collect papers from the recent 4 years (2019-2022) ACL main conference paper abstracts as candidate data.

Annotation and Quality Control During the pre-annotation stage, we train three supervisors and screen candidates, resulting in seven qualified annotators for the official annotation. Quality inspection is conducted by three directors and two well-performing annotators, with annotators and inspectors strictly separated. If a document contains more than two annotation conflicts (including missing annotations), it is returned to the original annotator together with detailed revision comments provided by a quality inspector. Otherwise, minor conflicts are corrected by the inspection team, and corresponding feedback is provided to annotators to facilitate continuous improvement. A fully annotated example document can be found in Appendix F.

3.2 Dataset Statistics Analysis

This section will provide a comprehensive statistical analysis of SciEvents, with particular emphasis on information density and complexity.

Dataset	Domain	Density: Every 100 Tokens Contains			Complexity: Complex Forms			
		#Events	#Args	#Nugget Tokens	#D(%)	#O(%)	#R(%)	#S(%)
ACE2005	News	1.80	2.72	4.62	–	13.88	–	–
Genia2011	Biomedical	3.92	3.44	9.31	–	36.62	–	–
Genia2013	Biomedical	4.00	3.78	8.58	–	34.00	–	–
RAMS	News	0.75	1.74	4.06	–	10.51	–	–
CASIE	Cybersecurity	2.19	5.86	17.80	–	1.40	–	–
M2E2	Multimedia	0.65	0.98	1.89	–	6.08	–	–
WikiEvents	News	2.08	2.92	5.69	–	8.19	–	–
PHEE	Drug Safety	4.72	24.20	53.43	–	64.03	–	–
Maccrobat-EE	Clinical	12.25	8.03	38.65	–	3.71	–	–
SciEvents (Ours)	Science	5.54	12.82	39.49	3.08	33.70	1.01	25.63

Table 2: Statistics of density and complexity of widely-used domain-specific event datasets. This table only presents publicly available event datasets that include argument annotations. #D: Discontinuous nugget. #O: Overlapping nugget. #R: Reverse-order nugget. #S: Sub-event.

Basic Statistics Table 1 presents basic statistics of SciEvents and other widely-used domain-specific EE datasets covering diverse domains. Among these datasets, SciEvents is distinguished as a large-scale dataset specifically constructed for the scientific domain. In terms of annotation scale, SciEvents 24,381 event instances and 56,411 arguments, substantially exceeding most existing domain-specific event datasets. Notably, SciEvents achieves this scale with only 10 event types. This suggests that the large number of event instances in SciEvents primarily arises from frequent event occurrences within scientific documents, rather than an expanded or fine-grained schema, reflecting the information-intensive nature of scientific texts. Statistics can be found in Appendix G.

Information Density Statistics As shown in Table 2, SciEvents exhibits high information density under all token-normalized metrics, with 5.54 events, 12.82 arguments, and 39.49 nugget tokens per 100 tokens, indicating that a large proportion of tokens in scientific documents directly participate in event expressions.

Among other domain-specific datasets, similarly high density values are mainly observed in medical datasets such as PHEE and Maccrobat-EE, whose documents describe inherently information-intensive content (*e.g.*, drug safety reports and clinical records). By contrast, remaining datasets generally exhibit substantially lower densities. Overall, these results suggest that the elevated density of SciEvents reflects intrinsic properties of scientific texts under domain-specific settings.

Information Complexity Statistics The right part of Table 2 reports the proportions of com-

plex event forms. Unlike most existing domain-specific datasets that mainly annotate contiguous nuggets, SciEvents explicitly covers diverse complex structures, including overlapping, discontinuous, reverse-order nuggets, and sub-events. Specifically, SciEvents contains a substantial proportion of overlapping nuggets (33.70%) and sub-events (25.63%), together with non-negligible occurrences of discontinuous (3.08%) and reverse-order nuggets (1.01%). By contrast, other datasets only partially capture overlapping structures. These statistics reflect the structural complexity of scientific events and highlight the challenges they pose for EE models.

4 Event Extraction Problem Formulation

In domain-specific EE, a predefined schema is given as $S = \{T_E, T_A\}$, where T_E denotes an event type set and T_A denotes an argument type set. Each event type $t_e \in T_E$ is associated with a specific argument type set $T_A(t_e)$. Given a document D , EE aims to extract a set of events $E = \{e_1, e_2, \dots, e_M\}$ in D , where each event $e = \{t_e, t, A\}$ consists of an event type $t_e \in T_E$, a trigger t and a set of arguments $A = \{a_1, a_2, \dots, a_N\}$. Each argument $a = \{t_a, m\}$ consists of an argument type $t_a \in T_A(t_e)$ and a word span m . Both triggers and arguments are referred to as nuggets, whose word spans should be combinations of tokens in D . In SciEvents, we add an event correlation task to extract hierarchical event structures. Specifically, the trigger t_s of a sub-event $e_s = \{t_{se}, t_s, A_s\}$ will be regarded as an argument of a main-event $e_m = \{t_{me}, t_m, A_m\}$ with a certain argument type t_{sa} , *i.e.* $\{t_{sa}, t_s\} \in A_m$.

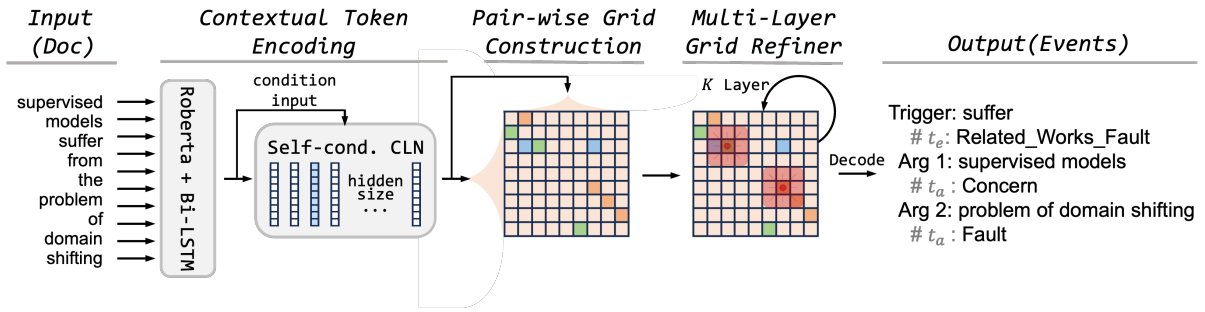


Figure 2: Overall architecture of EXCEEDS. The model encodes contextual token representations, constructs a word-word event grid to model pairwise relations, refines the grid, and decodes events from the refined grid.

5 The EXCEEDS Method

To address the challenges of high density and complexity in scientific EE, we propose EXCEEDS, an end-to-end framework that simplifies EE into a nugget-based grid modeling task. Section 5.1 introduces the word-word event grid construction, Section 5.2 presents the overall framework, and Section 5.3 describes the training objectives and inference procedure.

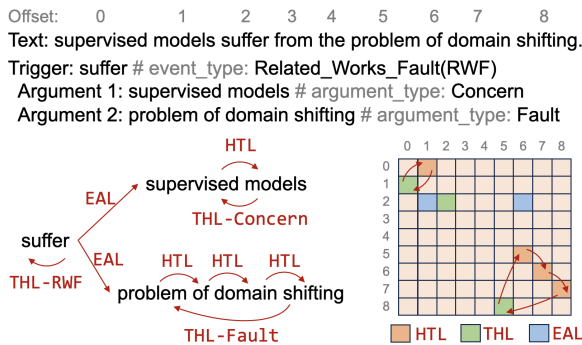


Figure 3: Illustration of the word-word event grid. HTL captures successive token order within a nugget; THL connects the last and first tokens to indicate nugget types; and EAL encodes relations across nuggets. The grid matrix (*right*) presents these relations.

5.1 Word-Word Event Grid

To effectively capture the complex structures of nuggets and events, we encode relations within a nugget and across different nuggets through a word-word grid. Formally, given a document $D = \{x_1, x_2, \dots, x_l\}$, we construct an $l \times l$ grid G , where each cell $G[i, j]$ stores the relation type $r \in R$ between token x_i (row) and token x_j (column). Specifically, within a nugget, we use head-tail-link (HTL) to represent the successive order between adjacent tokens and tail-head-link (THL) to connect the last token back to the first token,

which conveys nugget type information. Across different nuggets, we define event-argument-link (EAL) to represent the relation between a trigger and its argument, or between an event trigger and a sub-event trigger. For example, Figure 3 shows how HTL, THL and EAL are instantiated in Sci-Events, and how they are encoded into the corresponding cells of the grid, resulting in a unified representation for nuggets and events.

The benefits of this grid and relation design are threefold. First, it enables encoding of complex nugget structures within a document, including overlapping, discontinuous, and reverse-order nuggets. Second, it provides a unified formulation of event detection and event argument extraction in an end-to-end manner, allowing the framework to fully leverage contextual information without relying on separate pipeline modules. Third, it naturally captures hierarchical event relations by encoding relations between trigger pairs.

5.2 The Overall Framework

Figure 2 illustrates the overall architecture of our framework. Given a document, the model encodes contextual token representations and constructs a word-word grid to jointly model nugget structures and event relations in an end-to-end manner.

Contextual Token Encoding We encode the input document using a pretrained language model (Liu et al., 2019b) to obtain initial contextual representations, which are further refined by a bidirectional LSTM (Huang et al., 2015) to capture sequential dependencies. The resulting representations are normalized using Conditional Layer Normalization (CLN) (Liu et al., 2021) to enhance stability and contextual adaptability. Specifically, given the LSTM outputs $\mathbf{L} \in \mathbb{R}^{l \times d}$, CLN performs normalization with dynamically generated affine

parameters conditioned on the contextual representations themselves:

$$\mathbf{H} = \text{MLP}_\gamma(\mathbf{L}) \odot \frac{\mathbf{L} - \mu}{\sigma + \epsilon} + \text{MLP}_\beta(\mathbf{L}), \quad (1)$$

where μ and σ denote the mean and standard deviation computed along the feature dimension, ϵ is a smoothing parameter, and \odot denotes element-wise multiplication. The output $\mathbf{H} \in \mathbb{R}^{l \times d}$ serves as the contextualized word representations for subsequent pair-wise grid construction.

Pair-wise Grid Construction Given the contextualized word representations \mathbf{H} , we construct a word-word grid $\mathbf{G} \in \mathbb{R}^{l \times l \times C_g}$, where each cell corresponds to relation of a token pair (x_i, x_j) .

For each pair, we form a pair-wise representation by concatenating the token representations with a relative distance embedding:

$$\mathbf{z}_{i,j} = [\mathbf{h}_i; \mathbf{h}_j; \mathbf{d}_{i,j}], \quad (2)$$

which is projected into the grid feature space via a multilayer perceptron:

$$\mathbf{g}_{i,j} = \text{MLP}_{\text{pair}}(\mathbf{z}_{i,j}). \quad (3)$$

The resulting $\mathbf{g}_{i,j} \in \mathbb{R}^{C_g}$ constitutes the initial grid representation.

Grid Refiner The initial grid representations encode pair-wise token relations independently. To enable information propagation across related token pairs, we refine the grid with a stack of lightweight residual refinement blocks operating on the grid space.

Let $\mathbf{G}^{(0)} = \mathbf{G}$ and $\mathbf{G}^{(k)} \in \mathbb{R}^{l \times l \times C_g}$ denotes the grid features after the k -th refinement layer. Each block updates the grid by aggregating information from local neighborhoods and applying a residual transformation:

$$\mathbf{G}^{(k+1)} = \text{Norm}\left(\mathbf{G}^{(k)} + \mathcal{F}\left(\mathbf{G}^{(k)}\right)\right), \quad (4)$$

where $\mathcal{F}(\cdot)$ denotes a learnable local aggregation function on the grid, and is instantiated as stacked 2D convolutional refinement blocks in our implementation. After K refinement layers, we obtain the refined grid representations $\tilde{\mathbf{G}} = \mathbf{G}^{(K)}$.

5.3 Training and Inference

Loss Function Given $\tilde{\mathbf{G}} \in \mathbb{R}^{l \times l \times C_g}$, we project each grid cell to relation logits via a linear classifier, yielding $\mathbf{Y} \in \mathbb{R}^{l \times l \times |R|}$. Since multiple relation types may simultaneously hold for a token

pair, we formulate grid prediction as a multi-label classification problem.

We adopt a multi-label categorical cross-entropy loss (Su et al., 2022), which jointly optimizes positive and negative labels without requiring a predefined number of active labels. Formally, for each grid cell (i, j) , the loss is defined as

$$\mathcal{L}_{i,j} = \log\left(1 + \sum_{r \in \Omega^-} e^{y_{i,j}^r}\right) + \log\left(1 + \sum_{r \in \Omega^+} e^{-y_{i,j}^r}\right), \quad (5)$$

where Ω^+ and Ω^- denote the sets of positive and negative relation types for (x_i, x_j) , respectively.

Inference Following the multi-label formulation, we obtain the predicted relation set for each grid cell by a zero-threshold decision:

$$\hat{\mathbf{M}}_{i,j,r} = \mathbb{I}[y_{i,j}^r > 0], \quad (6)$$

where $\hat{\mathbf{M}} \in \{0, 1\}^{l \times l \times |R|}$ is the binary word-word relation grid. We then decode $\hat{\mathbf{M}}$ into a set of events by reconstructing nuggets and linking arguments to triggers. Specifically, we (1) recover nugget spans by traversing HTL and validating them with a closing THL-type, and (2) attach argument nuggets to trigger nuggets using EAL and schema constraints. Appendix A presents the detailed decoding algorithm.

6 Experiment

6.1 Experiment Settings

Evaluation Metrics Four standard metrics are adopted: trigger identification (TI), trigger classification (TC), argument identification (AI), and argument classification (AC). In addition, we introduce event correlation (EC) to evaluate the extraction of hierarchical sub-event relations. Specifically, when the trigger of one event appears as an argument of another event, the two events are considered correlated through their triggers. An evaluation example is provided in Appendix B.

Baselines We conduct a comprehensive evaluation of state-of-the-art and recent EE models, which can be broadly categorized into three groups: (1) Global information extraction models that jointly model entities, relations, and events within a unified framework; (2) Discriminative EE models that formulate EE as token classification or sequential labeling problems; (3) Generative EE models that generate extractions via question answering or through a well-designed generative schema.

Model		TI	TC	AI	AC	EC
Global	OneIE (Lin et al., 2020)	75.72 ± 0.14	62.93 ± 0.17	30.30 ± 1.85	28.81 ± 1.77	37.41 ± 1.51
	ScentedEAE [†] (Yang et al., 2024)	73.27 ± 0.52	63.03 ± 0.20	36.70 ± 2.99	35.74 ± 2.41	37.88 ± 2.10
Discriminative	EEQA (Du and Cardie, 2020b)	74.85 ± 0.78	62.15 ± 0.73	37.75 ± 0.46	35.64 ± 0.59	44.81 ± 1.35
	PAIE [†] (Ma et al., 2022)	73.27 ± 0.52	63.03 ± 0.20	43.92 ± 0.22	42.06 ± 0.34	47.17 ± 0.92
	Tagprime (Hsu et al., 2023)	73.27 ± 0.52	63.03 ± 0.20	44.67 ± 0.13	42.69 ± 0.32	47.72 ± 0.32
	DEEIA [†] (Liu et al., 2024)	73.27 ± 0.52	63.03 ± 0.20	34.86 ± 0.81	33.30 ± 0.75	32.80 ± 1.67
Generative	BartGen [†] (Li et al., 2021)	73.27 ± 0.52	63.03 ± 0.20	39.85 ± 0.52	37.81 ± 0.39	42.75 ± 0.11
	DEGREE (Hsu et al., 2022)	65.72 ± 0.70	53.52 ± 0.75	28.40 ± 0.88	26.32 ± 0.79	29.53 ± 0.96
	KnowCoder (Li et al., 2024)	69.88 ± 0.61	52.02 ± 0.34	35.24 ± 0.11	33.43 ± 0.27	34.54 ± 0.82
EXCEEDS (Ours)		75.29 ± 0.32	63.74 ± 0.14	44.97 ± 0.28	43.20 ± 0.29	48.25 ± 0.10
– Contextual		75.29 ± 0.21	63.44 ± 0.67	44.07 ± 0.51	42.14 ± 0.39	47.64 ± 0.85
– Grid Refiner		75.36 ± 0.27	63.41 ± 0.67	44.30 ± 0.51	42.44 ± 0.59	48.04 ± 1.07

Table 3: Overall F1-score (%) on SciEvents. For [†]EAE-only models, trigger predictions are derived from Tagprime, which achieves the best ED performance among all baseline methods.

For a fair comparison, when evaluating EAE-only methods, we first apply a best-performing ED method to extract event triggers. The EAE-only methods then perform argument extraction conditioned on these predicted triggers.

Implementations We randomly split SciEvents into training, development, and test sets with a ratio of 80%/10%/10%. Each model is evaluated over three independent runs, and the average performance is reported. For models with the same architecture, we use the same pretrained backbone. Additional details are provided in Appendix C.

6.2 Experiment Results

Overall Results Table 3 presents the overall performance of different models on SciEvents. Results with precision and recall score can be found in Appendix D. Overall, EXCEEDS achieves the strongest performance on most evaluation metrics, particularly on AI, AC and EC. Notably, EXCEEDS outperforms the second-best model by 0.51% on AC and 0.53% on EC, indicating its advantage in extracting arguments and capturing hierarchical sub-event relations in scientific documents.

Among baseline families, global models show competitive performance on TI and TC, which can be attributed to their use of entity information during training. Discriminative models generally perform better on EAE than global and generative models, while generative approaches exhibit larger performance variance across metrics.

Ablation study shows that removing the contextual modeling module results in the most pronounced performance drop on EAE, with AC de-

creasing by 1.06%, indicating that grid-based modeling critically relies on high-quality contextual token representations. Excluding the grid refiner also leads to consistent degradation across metrics. These results suggest the effectiveness of contextual encoding and grid refinement in EXCEEDS. Despite these improvements, the overall performance on AC remains modest, indicating that SciEvents remains a challenging EE dataset.

Complex Scenarios Results Table 4 reports model performance on complex scenarios, where evaluation is conducted on subsets filtered by specific information forms. Overall, EXCEEDS consistently outperforms all baselines across complex scenarios, demonstrating its effectiveness in extracting complex nuggets and events. However, performance on discontinuous, overlapping, and reverse-order scenarios is substantially lower than the overall results, indicating that these forms remain particularly challenging for current EE models, even with specialized modeling designs.

Notably, generative models do not exhibit advantages under complex scenarios. In particular, they show pronounced performance degradation on overlapping nuggets, a trend also observed in EEQA, which adopts a generation-style formulation. This suggests that directly generating textual outputs is insufficient for accurately capturing complex nugget structures.

Error Analysis Figure 4 shows the identification error distributions for TI and AI, where errors are categorized into missed, predicted-long, predicted-short, and other overlap cases. Missed predictions dominate identification errors, accounting for 89.2% of TI errors and 84.6% of AI errors.

Model	Discontinuous	Overlapping		Reverse-order	TC	Sub-event	EC
	AC	TC	AC	AC		AC	
OneIE	–	62.20 \pm 1.07	17.33 \pm 0.98	–	51.19 \pm 0.69	37.41 \pm 1.51	41.39 \pm 2.72
ScentedEAE	–	55.77 \pm 2.40	11.05 \pm 1.23	–	43.23 \pm 2.35	38.00 \pm 2.19	33.04 \pm 1.45
Tagprime	–	55.03 \pm 1.59	18.11 \pm 1.13	–	53.84 \pm 0.81	47.89 \pm 0.39	48.11 \pm 2.45
EEQA	–	17.02 \pm 0.80	2.33 \pm 0.02	–	53.50 \pm 0.92	44.81 \pm 1.35	45.16 \pm 1.21
PAIE	–	49.62 \pm 1.71	13.18 \pm 0.67	–	53.66 \pm 1.75	47.34 \pm 1.01	49.08 \pm 1.75
DEEIA	–	51.31 \pm 2.84	13.13 \pm 1.21	–	42.73 \pm 1.20	32.80 \pm 1.67	31.30 \pm 3.73
BartGen	2.17 \pm 0.00	32.53 \pm 0.00	9.84 \pm 0.00	0.00 \pm 0.00	48.95 \pm 0.00	40.00 \pm 0.00	33.18 \pm 0.00
DEGREE	0.00 \pm 0.00	8.01 \pm 0.00	2.67 \pm 0.00	0.00 \pm 0.00	28.92 \pm 0.00	20.68 \pm 0.00	12.72 \pm 0.00
KnowCoder	0.00 \pm 0.00	18.56 \pm 0.96	6.45 \pm 0.33	0.00 \pm 0.00	42.36 \pm 1.32	34.81 \pm 0.89	40.33 \pm 2.41
EXCEEDS	13.86 \pm 2.29	62.46 \pm 3.86	22.46 \pm 2.01	7.27 \pm 5.14	55.13 \pm 0.32	48.32 \pm 0.18	51.15 \pm 0.56
– Contextual	13.41 \pm 1.34	59.72 \pm 7.56	20.43 \pm 3.20	9.55 \pm 4.41	54.72 \pm 1.13	47.64 \pm 0.85	50.18 \pm 1.30
– Grid Refiner	14.98 \pm 1.19	61.81 \pm 2.86	<u>22.07</u> \pm 0.80	11.99 \pm 2.12	55.28 \pm 0.91	<u>48.04</u> \pm 1.07	49.11 \pm 0.98

Table 4: F1-score (%) on complex nuggets and events. – indicates that the method cannot be evaluated.




Figure 4: Identification Error Distribution

This suggests that the primary challenge lies in failing to detect the presence of event mentions, rather than in inaccurately predicting their span boundaries. By contrast, boundary-related errors constitute a substantially smaller portion of the overall errors. This observation indicates nugget boundary imprecision is not the main bottleneck for identification performance. Overall, these results highlight that improving recall in dense scientific contexts remains a key challenge for event identification.

Figure 5 presents the misclassification patterns in EXCEEDS for TC and AC. Overall, errors in both TC and AC are dominated by confusions between semantically similar types. For TC, frequent confusions occur between closely related event types, notably *MDS* and *WKS*. For AC, the most common errors also arise from fine-grained role distinctions, such as *TriedC.* vs. *BaseC.* and *Subject* vs. *Object*. These patterns suggest that improving classification on SciEvents likely requires more fine-grained modeling of subtle semantic distinctions, e.g., specialized disambiguation components, or incorporating schema/template-level cues (Ma et al., 2022; Liu et al., 2024) to better differentiate similar types.

7 Conclusion

We address the challenge of understanding scientific domain with complex event structures through EE by introducing SciEvents, a large-scale document-level dataset with a refined schema tailored to the scientific domain, and EXCEEDS, a nugget-based grid modeling EE approach. Experiments show that SciEvents reflects the information-dense and structurally complex nature of scientific texts, while EXCEEDS achieves strong performance. Further analysis indicates that SciEvents remains a challenging dataset and suggests directions for future improvements.

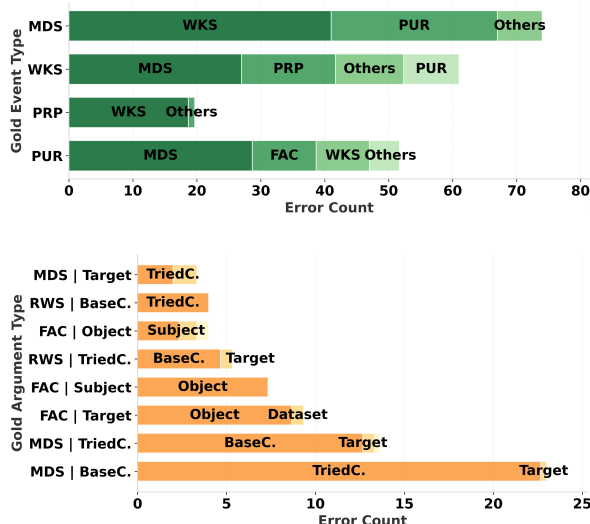


Figure 5: Classification Error Distribution. Y-axis denotes the misclassified gold type and each stacked bar indicates an incorrect predicted type and its frequency.

559 Limitations

560 **Abstract-Level Data Scope** SciEvents is con-
561 structed from paper abstracts, rather than full-
562 length scientific articles. While abstracts typically
563 provide concise and information-dense summaries
564 of scientific contributions, they do not capture all
565 event mentions that may appear in the main body
566 of a paper. In particular, important information con-
567 veyed through figures, tables, equations, or cross-
568 section references is not covered in our current
569 dataset. As a result, SciEvents does not account for
570 multimodal or long-range contextual signals that
571 may be essential for comprehensive scientific event
572 understanding. Extending the dataset to full-text
573 articles and incorporating multimodal information
574 remains an important direction for future work.

575 References

576 Jacqueline Aguilar, Charley Beller, Paul McNamee,
577 Benjamin Van Durme, Stephanie Strassel, Zhiyi
578 Song, and Joe Ellis. 2014. [A comparison of the
579 events and relations across ACE, ERE, TAC-KBP,
580 and FrameNet annotation standards](#). In *Proceedings
581 of the Second Workshop on EVENTS: Definition, De-
582 tection, Coreference, and Representation*, pages 45–
583 53, Baltimore, Maryland, USA. Association for Com-
584 putational Linguistics.

585 Yunmo Chen, Tongfei Chen, Seth Ebner, Aaron Steven
586 White, and Benjamin Van Durme. 2019. Reading the
587 manual: Event extraction as definition comprehen-
588 sion. *arXiv preprint arXiv:1912.01586*.

589 George Doddington, Alexis Mitchell, Mark Przybocki,
590 Lance Ramshaw, Stephanie Strassel, and Ralph
591 Weischedel. 2004. [The automatic content extrac-
592 tion \(ACE\) program – tasks, data, and evaluation](#). In
593 *Proceedings of the Fourth International Conference
594 on Language Resources and Evaluation (LREC’04)*,
595 Lisbon, Portugal. European Language Resources As-
596 sociation (ELRA).

597 Xinya Du and Claire Cardie. 2020a. Event extraction by
598 answering (almost) natural questions. *arXiv preprint
599 arXiv:2004.13625*.

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

605 Seth Ebner, Patrick Xia, Ryan Culkin, Kyle Rawlins,
606 and Benjamin Van Durme. 2020. [Multi-sentence ar-
607 gument linking](#). In *Proceedings of the 58th Annual
608 Meeting of the Association for Computational Lin-
609 guistics*, pages 8057–8077, Online. Association for
610 Computational Linguistics.

Jun Gao, Huan Zhao, Changlong Yu, and Ruifeng Xu. 2023a. Exploring the feasibility of chatgpt for event extraction. *arXiv preprint arXiv:2303.03836*. 611 612 613

Jun Gao, Huan Zhao, Yice Zhang, Wei Wang, Chang- long Yu, and Ruifeng Xu. 2023b. Benchmark- ing large language models with augmented instruc- tions for fine-grained information extraction. *arXiv preprint arXiv:2310.05092*. 614 615 616 617 618

I-Hung Hsu, Kuan-Hao Huang, Elizabeth Boschee, Scott Miller, Prem Natarajan, Kai-Wei Chang, and Nanyun Peng. 2022. [DEGREE: A data-efficient generation-based event extraction model](#). In *Pro- ceedings of the 2022 Conference of the North Amer- ican Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1890–1908, Seattle, United States. Association for Computational Linguistics. 619 620 621 622 623 624 625 626 627

I-Hung Hsu, Kuan-Hao Huang, Shuning Zhang, Wenxin Cheng, Prem Natarajan, Kai-Wei Chang, and Nanyun Peng. 2023. [TAGPRIME: A unified framework for relational structure extraction](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12917–12932, Toronto, Canada. Association for Computational Linguistics. 628 629 630 631 632 633 634 635

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *Preprint*, arXiv:2106.09685. 636 637 638 639

Kuan-Hao Huang, I-Hung Hsu, Tanmay Parekh, Zhiyu Xie, Zixuan Zhang, Prem Natarajan, Kai-Wei Chang, Nanyun Peng, and Heng Ji. 2024. [TextEE: Bench- mark, reevaluation, reflections, and future challenges in event extraction](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 12804–12825, Bangkok, Thailand. Association for Computational Linguistics. 640 641 642 643 644 645 646 647

Zhiheng Huang, Wei Xu, and Kai Yu. 2015. [Bidirec- tional lstm-crf models for sequence tagging](#). *Preprint*, arXiv:1508.01991. 648 649 650

Jin-Dong Kim, Yue Wang, Toshihisa Takagi, and Aki- nori Yonezawa. 2011. [Overview of Genia event task in BioNLP shared task 2011](#). In *Proceedings of BioNLP Shared Task 2011 Workshop*, pages 7–15, Portland, Oregon, USA. Association for Computa- tional Linguistics. 651 652 653 654 655 656

Jin-Dong Kim, Yue Wang, and Yamamoto Yasunori. 2013. [The Genia event extraction shared task, 2013 edition - overview](#). In *Proceedings of the BioNLP Shared Task 2013 Workshop*, pages 8–15, Sofia, Bul- garia. Association for Computational Linguistics. 657 658 659 660 661

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. [Bart: De- noising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). *Preprint*, arXiv:1910.13461. 662 663 664 665 666 667

668	Fayuan Li, Weihua Peng, Yuguang Chen, Quan Wang,	Wanlong Liu, Li Zhou, Dingyi Zeng, Yichen Xiao,	726
669	Lu Pan, Yajuan Lyu, and Yong Zhu. 2020a. Event	Shaohuan Cheng, Chen Zhang, Grandee Lee, Malu	727
670	extraction as multi-turn question answering. In <i>Find-</i>	Zhang, and Wenyu Chen. 2024. Beyond single-event	728
671	<i>ings of the Association for Computational Linguistics:</i>	extraction: Towards efficient document-level multi-	729
672	<i>EMNLP 2020</i> , pages 829–838.	event argument extraction . In <i>Findings of the Asso-</i>	730
		<i>ciation for Computational Linguistics: ACL 2024</i> ,	731
673	Manling Li, Alireza Zareian, Qi Zeng, Spencer White-	pages 9470–9487, Bangkok, Thailand. Association	732
674	head, Di Lu, Heng Ji, and Shih-Fu Chang. 2020b.	for Computational Linguistics.	733
675	Cross-media structured common space for multime-		
676	dia event extraction . In <i>Proceedings of the 58th Annual</i>	Xiao Liu, Heyan Huang, Ge Shi, and Bo Wang. 2022.	734
677	<i>Meeting of the Association for Computational</i>	Dynamic prefix-tuning for generative template-based	735
678	<i>Linguistics</i> , pages 2557–2568, Online. Association	event extraction . In <i>Proceedings of the 60th Annual</i>	736
679	for Computational Linguistics.	<i>Meeting of the Association for Computational Lin-</i>	737
		<i>guistics (Volume 1: Long Papers)</i> , pages 5216–5228,	738
680	Sha Li, Heng Ji, and Jiawei Han. 2021. Document-level	Dublin, Ireland. Association for Computational Lin-	739
681	event argument extraction by conditional generation .	guistics.	740
682	In <i>Proceedings of the 2021 Conference of the North</i>		
683	<i>American Chapter of the Association for Computa-</i>	Xiao Liu, Heyan Huang, and Yue Zhang. 2019a. Open	741
684	<i>tional Linguistics: Human Language Technologies</i> ,	domain event extraction using neural latent variable	742
685	pages 894–908, Online. Association for Computa-	models . In <i>Proceedings of the 57th Annual Meet-</i>	743
686	tional Linguistics.	<i>ing of the Association for Computational Linguistics</i> ,	744
		pages 2860–2871, Florence, Italy. Association for	745
687	Xinyu Li, Fayuan Li, Lu Pan, Yuguang Chen, Weihua	Computational Linguistics.	746
688	Peng, Quan Wang, Yajuan Lyu, and Yong Zhu. 2020c.		
689	Duee: a large-scale dataset for chinese event extrac-	Xiao Liu, Zhunchen Luo, and Heyan Huang. 2018.	747
690	tion in real-world scenarios. In <i>Natural Language</i>	Jointly multiple events extraction via attention-based	748
691	<i>Processing and Chinese Computing: 9th CCF Inter-</i>	graph information aggregation. <i>arXiv preprint</i>	749
692	<i>national Conference, NLPCC 2020, Zhengzhou,</i>	<i>arXiv:1809.09078</i> .	750
693	<i>China, October 14–18, 2020, Proceedings, Part II 9</i> ,		
694	pages 534–545. Springer.	Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Man-	751
		dar Joshi, Danqi Chen, Omer Levy, Mike Lewis,	752
695	Zixuan Li, Yutao Zeng, Yuxin Zuo, Weicheng Ren,	Luke Zettlemoyer, and Veselin Stoyanov. 2019b.	753
696	Wenxuan Liu, Miao Su, Yucan Guo, Yantao Liu, Xi-	Roberta: A robustly optimized bert pretraining ap-	754
697	ang Li, Zhilei Hu, Long Bai, Wei Li, Yidan Liu, Pan	proach . <i>Preprint</i> , arXiv:1907.11692.	755
698	Yang, Xiaolong Jin, Jiafeng Guo, and Xueqi Cheng.		
699	2024. KnowCoder: Coding structured knowledge	Jie Lou, Yaojie Lu, Dai Dai, Wei Jia, Hongyu Lin, Xi-	756
700	into LLMs for universal information extraction . In	anpei Han, Le Sun, and Hua Wu. 2023. Universal	757
701	<i>Proceedings of the 62nd Annual Meeting of the As-</i>	information extraction as unified semantic matching.	758
702	<i>sociation for Computational Linguistics (Volume 1:</i>	<i>arXiv preprint arXiv:2301.03282</i> .	759
703	<i>Long Papers)</i> , pages 8758–8779, Bangkok, Thailand.		
704	Association for Computational Linguistics.	Yaojie Lu, Hongyu Lin, Jin Xu, Xianpei Han, Jialong	760
		Tang, Annan Li, Le Sun, Meng Liao, and Shaoyi	761
705	Xin Liang, Dawei Cheng, Fangzhou Yang, Yifeng Luo,	Chen. 2021. Text2event: Controllable sequence-to-	762
706	Weining Qian, and Aoying Zhou. 2020. F-hmtc:	structure generation for end-to-end event extraction.	763
707	Detecting financial events for investment decisions	<i>arXiv preprint arXiv:2106.09232</i> .	764
708	based on neural hierarchical multi-label text classi-		
709	fication . In <i>Proceedings of the Twenty-Ninth Inter-</i>	Yaojie Lu, Qing Liu, Dai Dai, Xinyan Xiao, Hongyu	765
710	<i>national Joint Conference on Artificial Intelligence,</i>	Lin, Xianpei Han, Le Sun, and Hua Wu. 2022a. Uni-	766
711	<i>IJCAI-20</i> , pages 4490–4496. International Joint Con-	fied structure generation for universal information	767
712	ferences on Artificial Intelligence Organization. Spe-	extraction . In <i>Proceedings of the 60th Annual Meet-</i>	768
713	cial Track on AI in FinTech.	<i>ing of the Association for Computational Linguistics</i>	769
		<i>(Volume 1: Long Papers)</i> , pages 5755–5772, Dublin,	770
714	Ying Lin, Heng Ji, Fei Huang, and Lingfei Wu. 2020.	Ireland. Association for Computational Linguistics.	771
715	A joint neural model for information extraction with		
716	global features . In <i>Proceedings of the 58th Annual</i>	Yaojie Lu, Qing Liu, Dai Dai, Xinyan Xiao, Hongyu	772
717	<i>Meeting of the Association for Computational Lin-</i>	Lin, Xianpei Han, Le Sun, and Hua Wu. 2022b. Uni-	773
718	<i>guistics</i> , pages 7999–8009, Online. Association for	fied structure generation for universal information	774
719	Computational Linguistics.	extraction. <i>arXiv preprint arXiv:2203.12277</i> .	775
720	Ruibo Liu, Jason Wei, Chenyan Jia, and Soroush	Mingyu Derek Ma, Alexander Taylor, Wei Wang, and	776
721	Vosoughi. 2021. Modulating language models with	Nanyun Peng. 2023. DICE: Data-efficient clinical	777
722	emotions . In <i>Findings of the Association for Com-</i>	event extraction with generative models . In <i>Proceed-</i>	778
723	<i>putational Linguistics: ACL-IJCNLP 2021</i> , pages	<i>ings of the 61st Annual Meeting of the Association for</i>	779
724	4332–4339, Online. Association for Computational	<i>Computational Linguistics (Volume 1: Long Papers)</i> ,	780
725	Linguistics.	pages 15898–15917, Toronto, Canada. Association	781
		for Computational Linguistics.	782

783	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 <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 6759–6774, Dublin, Ireland. Association for Computational Linguistics.	841
784		842
785		843
786		844
787		845
788		
789		
790		
791	Hieu Man Duc Trong, Duc Trong Le, Amir Pouran Ben Veyseh, Thuat Nguyen, and Thien Huu Nguyen. 2020. Introducing a new dataset for event detection in cybersecurity texts . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 5381–5390, Online. Association for Computational Linguistics.	846
792		847
793		848
794		849
795		
796		
797		
798	Minh Van Nguyen, Viet Dac Lai, and Thien Huu Nguyen. 2021. Cross-task instance representation interactions and label dependencies for joint information extraction with graph convolutional networks . In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 27–38, Online. Association for Computational Linguistics.	850
799		851
800		852
801		853
802		854
803		855
804		856
805		
806		
807	Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. 2016. Joint event extraction via recurrent neural networks. In <i>Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies</i> , pages 300–309.	857
808		858
809		859
810		860
811		861
812		
813	Tanmay Parekh, I-Hung Hsu, Kuan-Hao Huang, Kai-Wei Chang, and Nanyun Peng. 2023. GENEVA: Benchmarking generalizability for event argument extraction with hundreds of event types and argument roles . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 3664–3686, Toronto, Canada. Association for Computational Linguistics.	862
814		863
815		864
816		865
817		866
818		867
819		868
820		869
821	Hao Peng, Xiaozhi Wang, Feng Yao, Kaisheng Zeng, Lei Hou, Juanzi Li, Zhiyuan Liu, and Weixing Shen. 2023. The devil is in the details: On the pitfalls of event extraction evaluation . In <i>Findings of the Association for Computational Linguistics: ACL 2023</i> , pages 9206–9227, Toronto, Canada. Association for Computational Linguistics.	870
822		871
823		872
824		873
825		
826		
827		
828	Amir Pouran Ben Veyseh, Minh Van Nguyen, Franck Dernoncourt, and Thien Nguyen. 2022. MINION: a large-scale and diverse dataset for multilingual event detection . In <i>Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 2286–2299, Seattle, United States. Association for Computational Linguistics.	874
829		875
830		876
831		877
832		878
833		879
834		880
835		881
836	Sampo Pyysalo, Tomoko Ohta, Makoto Miwa, Hanchchol Cho, Jun’ichi Tsujii, and Sophia Ananiadou. 2012. Event extraction across multiple levels of biological organization. <i>Bioinformatics</i> , 28(18):i575–i581.	882
837		883
838		884
839		885
840		886
	Oscar Sainz, Iker García-Ferrero, Rodrigo Agerri, Oier Lopez de Lacalle, German Rigau, and Eneko Agirre. 2023. Gollie: Annotation guidelines improve zero-shot information-extraction . <i>arXiv preprint arXiv:2310.03668</i> .	887
		888
		889
		890
		891
		892
		893
		894
	Taneeya Satyapanich, Francis Ferraro, and Tim Finin. 2020. Casie: Extracting cybersecurity event information from text . <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , 34(05):8749–8757.	895
		896
	Shirong Shen, Guilin Qi, Zhen Li, Sheng Bi, and Lusheng Wang. 2020. Hierarchical Chinese legal event extraction via pedal attention mechanism . In <i>Proceedings of the 28th International Conference on Computational Linguistics</i> , pages 100–113, Barcelona, Spain (Online). International Committee on Computational Linguistics.	
	Matthew Sims, Jong Ho Park, and David Bamman. 2019. Literary event detection . In <i>Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics</i> , pages 3623–3634, Florence, Italy. Association for 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 <i>Proceedings of the The 3rd Workshop on EVENTS: Definition, Detection, Coreference, and Representation</i> , pages 89–98, Denver, Colorado. Association for Computational Linguistics.	
	Jianlin Su, Mingren Zhu, Ahmed Murtadha, Shengfeng Pan, Bo Wen, and Yunfeng Liu. 2022. Zlpr: A novel loss for multi-label classification . <i>Preprint, arXiv:2208.02955</i> .	
	Zhaoyue Sun, Jiazheng Li, Gabriele Pergola, Byron Wallace, Bino John, Nigel Greene, Joseph Kim, and Yulan He. 2022. PHEE: A dataset for pharmacovigilance event extraction from text . In <i>Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing</i> , pages 5571–5587, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.	
	Beth M. Sundheim. 1992. Overview of the fourth Message Understanding Evaluation and Conference . In <i>Fourth Message Understanding Conference (MUC-4): Proceedings of a Conference Held in McLean, Virginia, June 16-18, 1992</i> .	
	Wei Tang, Benfeng Xu, Yuyue Zhao, Zhendong Mao, Yifeng Liu, Yong Liao, and Haiyong Xie. 2022. UniRel: Unified representation and interaction for joint relational triple extraction . In <i>Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing</i> , pages 7087–7099, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.	
	MeiHan Tong, Bin Xu, Shuai Wang, Meihuan Han, Yixin Cao, Jiangqi Zhu, Siyu Chen, Lei Hou, and	

897	Juanzi Li. 2022. DocEE: A large-scale and fine-grained benchmark for document-level event extraction . In <i>Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 3970–3982, Seattle, United States. Association for Computational Linguistics.	956
898		957
899		958
900		
901		959
902		960
903		961
904	Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and finetuned chat models . <i>Preprint</i> , arXiv:2307.09288.	962
905		963
906		964
907		965
908		966
909		967
910		968
911		969
912		970
913		
914		971
915		972
916		973
917		974
918		975
919		
920		976
921		977
922		978
923		979
924		980
925		981
926		
927	Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. Ace 2005 multilingual training corpus. <i>Techbical report, Linguistic Data Consortium</i> .	
928		
929		
930		
931	Bo Wang, Heyan Huang, Xiaochi Wei, Ge Shi, Xiao Liu, Chong Feng, Tong Zhou, Shuaiqiang Wang, and Dawei Yin. 2023a. Boosting event extraction with denoised structure-to-text augmentation . In <i>Findings of the Association for Computational Linguistics: ACL 2023</i> , pages 11267–11281, Toronto, Canada. Association for Computational Linguistics.	
932		
933		
934		
935		
936		
937		
938	Xiao Wang, Weikang Zhou, Can Zu, Han Xia, Tianze Chen, Yuansen Zhang, Rui Zheng, Junjie Ye, Qi Zhang, Tao Gui, et al. 2023b. Instructuie: Multi-task instruction tuning for unified information extraction. <i>arXiv preprint arXiv:2304.08085</i> .	
939		
940		
941		
942		
943	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 <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 1652–1671, Online. Association for Computational Linguistics.	
944		
945		
946		
947		
948		
949		
950		
951	Kaiwen Wei, Xian Sun, Zequn Zhang, Jingyuan Zhang, Guo Zhi, and Li Jin. 2021. Trigger is not sufficient: Exploiting frame-aware knowledge for implicit event argument extraction. In <i>Proceedings of the 59th Annual Meeting of the Association for Computational</i>	
952		
953		
954		
955		
	<i>Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)</i> , pages 4672–4682.	956
		957
		958
	Xiang Wei, Xingyu Cui, Ning Cheng, Xiaobin Wang, Xin Zhang, Shen Huang, Pengjun Xie, Jinan Xu, Yufeng Chen, Meishan Zhang, et al. 2023. Zero-shot information extraction via chatting with chatgpt. <i>arXiv preprint arXiv:2302.10205</i> .	959
		960
		961
		962
		963
	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 <i>Proceedings of ACL 2018, System Demonstrations</i> , pages 50–55, Melbourne, Australia. Association for Computational Linguistics.	964
		965
		966
		967
		968
		969
		970
	Sen Yang, Dawei Feng, Linbo Qiao, Zhigang Kan, and Dongsheng Li. 2019. Exploring pre-trained language models for event extraction and generation. In <i>Proceedings of the 57th annual meeting of the association for computational linguistics</i> , pages 5284–5294.	971
		972
		973
		974
		975
	Yu Yang, Jinyu Guo, Kai Shuang, and Chenrui Mao. 2024. Scented-EAE: Stage-customized entity type embedding for event argument extraction . In <i>Findings of the Association for Computational Linguistics: ACL 2024</i> , pages 5222–5235, Bangkok, Thailand. Association for Computational Linguistics.	976
		977
		978
		979
		980
		981
	Feng Yao, Chaojun Xiao, Xiaozhi Wang, Zhiyuan Liu, Lei Hou, Cunchao Tu, Juanzi Li, Yun Liu, Weixing Shen, and Maosong Sun. 2022. LEVEN: A large-scale Chinese legal event detection dataset . In <i>Findings of the Association for Computational Linguistics: ACL 2022</i> , pages 183–201, Dublin, Ireland. Association for Computational Linguistics.	982
		983
		984
		985
		986
		987
		988
	Yang Zhou, Yubo Chen, Jun Zhao, Yin Wu, Jiexin Xu, and Jinlong Li. 2021. What the role is vs. what plays the role: Semi-supervised event argument extraction via dual question answering. In <i>Proceedings of the AAAI conference on artificial intelligence</i> , volume 35, pages 14638–14646.	989
		990
		991
		992
		993
		994
	A Word-word Event Grid Decoding	995
	To ensure the structural validity of decoded nuggets and events, we apply two pruning heuristics: (1) an HTL chain is kept only if it can be closed by a THL-type edge; (2) an argument nugget is kept only if it can be linked to at least one trigger nugget via EAL and passes the ontology constraint.	996
		997
		998
		999
		1000
		1001
	With these two pruning heuristics, Algorithm 1 summarizes the detailed decoding process.	1002
		1003
	Notably, to prevent potential excessively long decoding time during early training stages, we adopt a conservative training strategy. In early epochs, model predictions may be unstable and could assign a large number of labels to grid cells, which	1004
		1005
		1006
		1007
		1008

in the worst case may lead to an exponential number of candidate HTL chains during DFS-based decoding and significantly slow down validation. To avoid such degenerate cases, we skip validation in the initial training phase (typically the first few epochs) and enable regular validation after this phase.

	Event 1	Event 2
trigger	A (Type M)	C (Type N)
argument	B (Type BT) C (Type CT)	D (Type DT)

Table 5: An example with two prediction events. A, B, C and D are token spans. M and N are event types. BT, CT and DT are argument types.

Metric	What elements should match ground-truth
TI	A C
TC	A+M C+N
AI	A+M+B A+M+C C+N+D
AC	A+M+B+BT A+M+C+CT C+N+D+DT
EC	A+M+C+CT+N

Table 6: An evaluation example for two prediction events. A, B, C and D are token spans. M and N are event types. BT, CT and DT are argument types.

B Evaluation Metrics Demonstration

In SciEvents, there are 3 tasks and 5 kinds of metrics: (1) Trigger Extraction, also known as Event Detection, includes Trigger Identification (TI) and Trigger Classification (TC). (2) Event Argument Extraction includes Argument Identification (AI) and Argument Classification (AC). (3) Sub-Event Extraction includes Event Correlation (EC).

In SciEvents, a nugget serves as the basic unit for evaluation, and an exact-match criterion is applied when assessing nugget token spans.

Table 5 shows an example with two prediction events. Table 6 shows how to evaluate these two prediction events.

Hyperparameter	Value
Roberta-large Learning Rate	1e-5
Warm Up Ratio	0.1
Other Learning Rate	1e-3
Batch Size	2
Epoch	20
Distance Embedding Size	20
Bi-LSTM Hidden Size	1024
Grid Channels (C_g)	256
Grid Refiner Dropout Rate	0.1
Other Dropout Rate	0.5
Grid Refiner Layers (K)	2
Grid Refiner Kernel	3

Table 7: Hyperparameter used in EXCEEDS.

C Implementation Details

Pretrained Backbone. For all evaluated models, we adopt either BART-large (Lewis et al., 2019) or RoBERTa-large (Liu et al., 2019b) as the pretrained backbone to ensure a fair and controlled comparison across architectures. KnowCoder (Li et al., 2024), which is based on large language models, employs LLaMA 2-7B (Touvron et al., 2023) as its pretrained backbone.

Hyperparameter Settings. For EXCEEDS, the hyperparameter settings used in our implementation are reported in Table 7. For the remaining models, we primarily apply a unified event extraction framework, TextEE (Huang et al., 2024), and adopt the hyperparameters recommended by TextEE. For models that are not supported by TextEE, we adapt SciEvents to their official code, and use the hyperparameter settings suggested in their implementations.

Architecture-Specific Implementations. For global information extraction models, we leverage the nugget type information provided in SciEvents (described in Appendix E.1) to facilitate entity training, following their original modeling assumptions.

For discriminative models, as most models rely on explicit start and end offsets for training and inference, discontinuous and reverse-order nuggets are not directly applicable under their modeling assumptions. Accordingly, these models are evaluated only on nugget instances with contiguous span representations.

For generative models, input-output interfaces and evaluation scripts are modified to support raw text, without relying on explicit span offsets, during training, inference, and evaluation. This adaptation

ensures that generative approaches can be fairly evaluated on SciEvents under an offset-free setting.

Training and Inference Cost Most models are trained and evaluated on NVIDIA RTX 3090 GPUs. Experiments involving large language models are conducted on NVIDIA A800 80GB PCIe GPUs, where parameter-efficient fine-tuning with LoRA (Hu et al., 2021) is adopted. Table 8 reports the average training time **per epoch** and inference time for each model, measured in GPU hours, providing a reference for their computational cost.

Model		Training	Inference
Global	OneIE	0.1816	0.0036
	ScentedEAE	0.5249	0.0440
ED & EAE	EEQA-ED	0.0259	0.0056
	EEQA-EAE	5.2138	3.7292
	Tagprime-ED	0.2510	0.0014
	Tagprime-EAE	0.3910	0.0083
	DEGREE	0.3639	0.0067
	KnowCoder-ED [†]	0.2381	0.2150
	KnowCoder-EAE	1.5869	0.2897
	EXCEEDS (Ours)	0.0609	0.1692
EAE-only	PAIE	0.7536	0.0033
	DEEIA	0.1237	0.0165
	BartGen	0.2950	0.0236

Table 8: Average GPU hours required per training epoch and for inference across different models. [†] is based on a large language model and is fine-tuned using parameter-efficient adaptation.

D Full Experiment Results

Table 19, 20 and 21 presents the full experiment results, including precision and recall scores.

E Schema of SciEvents

The schema of SciEvents consists of nugget types and event types. In this section, we will introduce the description of each nugget type and the template of each event type.

E.1 Nugget Types

There are 10 nugget types in SciEvents as follows:

Research Organization / Group (OG) refers to a research team composed of people. Typical examples include: *We*; *Li et al., 2013*; *They*.

Approach (APP) refers to nouns, pronouns, and corresponding phrases that denote a complete method or algorithm with concrete inputs and outputs. Typical examples include: *... work*; *... model*;

... method; *... framework*; *... network*; *... algorithm*; *baselines*; *state-of-the-art*.

Module (MOD) refers to nouns, pronouns, and corresponding phrases that denote components of a method or architecture, such as modules or algorithmic elements. A single MOD item is usually not detailed enough to constitute a full APP. Typical examples include: *... encoders*; *... decoders*; *... module*; *... process*; *message propagation process*; *beam search*.

Feature (FEA) refers to nouns, pronouns, and corresponding phrases that denote features. Typical examples include: *... information*; *the first-order adjacency information*; *the relationships between labeled edges*.

Note the difference among APP, MOD, and FEA: **APP** refers to a complete method with concrete inputs (*e.g.*, a task) and outputs (*e.g.*, the desired results of the task). **MOD** refers to a sub-process or component within the overall APP framework, such as a module or algorithmic element. **FEA** refers to features utilized during the execution of an APP or a MOD, such as *positional information*, *vector representations of part-of-speech tags*, or *sentence length*.

Task (TAK) refers to phrases that denote the intention or objective of a task optimization, *i.e.*, the research focus or target point. These expressions are neutral. Typical examples include: *graph-to-sequence modeling*; *performance unimodal*; *performance multimodal*; *accuracy*; *F1 score*; *robustness*; *reproducibility*; *zero-shot translation quality*.

Dataset (DST) refers to nouns, pronouns, and corresponding phrases that denote datasets used or relied upon when describing an artifact, research objective, or experimental conclusion. Typical examples include: *TAC-KBP 2017 datasets*; *Chinese multimodal NER dataset*; *CNERTA*; *training data*.

Limit (LIM) refers to phrases that denote conditional or environmental limitations, often introduced with prepositions. Typical examples include: *for a small number of confusing type pairs*; *in existing verb metaphor detection benchmarks*; *of the dynamic self-attention*.

Strength (STR) refers to phrases that describe the advantages or strengths of an artifact, often with an evaluative or positive connotation. Typical examples include: *state-of-the-art performance*.

1142	Weakness (WEA) refers to phrases that describe	WorkStatement (WKS). As the schema shown	1188
1143	the disadvantages, shortcomings, or weaknesses	in Table 14, the WorkStatement event describes:	1189
1144	of an artifact, often with an evaluative or negative	<i><Researcher:arg1> report <Content:arg2></i>	1190
1145	connotation. Typical examples include: <i>most of</i>	<i>under <Condition:arg3> circumstance on</i>	1191
1146	<i>the mislabeling; biases and failure cases of beam</i>	<i><Dataset:arg4> datasets for <Target:arg5>.</i>	1192
1147	<i>search.</i>		
1148	Degree (DEG) refers to adjectives, adverbs, nu-	MethodStep (MDS). As the schema shown in	1193
1149	merals, or other expressions that describe the de-	Table 15, the MethodStep event describes: <i>Our</i>	1194
1150	gree or quantity of an event. Typical examples	<i>approach adopt <BaseComponent:arg1> with</i>	1195
1151	include: <i>only; not fully; 1.5%.</i>	<i><TriedComponent:arg2> under <Condition:arg3></i>	1196
1152		<i>circumstance on <Dataset:arg4> datasets for</i>	1197
1153	E.2 Event Types	<i><Target:arg5>.</i>	1198
1154	There are 10 event types classified by four different	(5) <i>Results</i> includes three kinds of event type:	1199
1155	rhetorical components as follows:	Finding (FIN). As the schema shown in Ta-	1200
1156	(1) <i>General.</i> <i>General</i> events occur in all four	ble 16, the Finding event describes: <i>In experi-</i>	1201
1157	components.	<i>ments, <Finder:arg1> find or demonstrate findings</i>	1202
1158		<i>that <Content:arg2>.</i>	1203
1159	Purpose (PUR). As the schema shown in Ta-	ExperimentCompare (CMP). As the schema	1204
1160	ble 9, the Purpose event describes: <i>In order to</i>	shown in Table 17, the ExperimentCompare event	1205
1161	<i>deal with <Aim:arg1> under <Condition:arg2></i>	describes: <i>Experimental results show that the</i>	1206
1162	<i>circumstance on <Dataset:arg3> datasets.</i>	<i><Metrics:arg6> of <Arg1:arg1> is <Extent:arg2></i>	1207
1163		<i><Result:arg3> than <Arg2:arg4> under <Con-</i>	1208
1164	(2) <i>Background</i> includes one kind of event type:	<i>dition:arg5> circumstance on <Dataset:arg7></i>	1209
1165		<i>datasets.</i>	1210
1166	IntroduceTarget (ITT). As the schema shown	OutcomeFact (FAC). As the schema shown	1211
1167	in Table 10, the IntroduceTarget event de-	in Table 18, the OutcomeFact event describes:	1212
1168	scribes: <i><Target:arg1> is the abstract research</i>	<i>Experimental results show that <Subject:arg1></i>	1213
1169	<i>target under <Condition:arg2> circumstance on</i>	<i>can <Extent:arg2> provide <Object:arg3> for</i>	1214
1170	<i><Dataset:arg3> datasets in this paper.</i>	<i><Target:arg4> under <Condition:arg5> circum-</i>	1215
1171		<i>stance on <Dataset:arg6> datasets because <Rea-</i>	1216
1172	(3) <i>Related Work</i> includes two kinds of event	<i>son:arg7> reasons.</i>	1217
1173	type:		
1174	RelatedWorkStep (RWS). As the schema	F Example of Document-Level Event	1218
1175	shown in Table 11, the RelatedWorkStep event	Annotation	1219
1176	describes: <i>Previously <Subject:arg1> on</i>	Table 22 presents a fully annotated example doc-	1220
1177	<i><Target:arg2> are mostly based on <BaseC-</i>	ument from SciEvents, illustrating all event in-	1221
1178	<i>omponent:arg3> with <TriedComponent:arg4></i>	stances annotated within a single document. For	1222
1179	<i>under <Condition:arg5> circumstance on</i>	each event, we show the corresponding trigger	1223
1180	<i><Dataset:arg6> datasets.</i>	nugget, argument nuggets, and their semantic roles,	1224
1181		as well as hierarchical sub-event relations when ap-	1225
1182	RelatedWorkFault (RWF). As the schema	licable. This example is provided to demonstrate	1226
1183	shown in Table 12, the RelatedWorkFault event	the density and structural complexity of event an-	1227
1184	describes: <i>Aiming to <Target:arg1>, to <Ex-</i>	notations in scientific documents, and to facilitate	1228
1185	<i>tent:arg5> degree, <Concern:arg2> has some</i>	a clearer understanding of the annotation schema	1229
1186	<i><Fault:arg6> faults under <Condition:arg3> cir-</i>	and evaluation setup.	1230
1187	<i>cumstance on <Dataset:arg4> datasets.</i>		
	(4) <i>Methodology</i> includes three kinds of event	G Distributions in SciEvents	1231
	type:	In this section, we will present comprehensive dis-	1232
	Propose (PRP). As the schema shown in Ta-	tributions in SciEvents.	1233
	ble 13, the Propose event describes: <i>In this pa-</i>		
	<i>per, <Proposer:arg1> propose <Content:arg2></i>		
	<i>for <Target:arg3>.</i>		

Nugget Type Distribution. Table 23 and Figure 6 present distribution of nugget types across train, develop and test splits.

Event Type Distribution. Table 24 and Figure 7 present distribution of event types across train, develop and test splits.

Argument Type Distribution. Table 25 and Figure 8 present distribution of argument types across train, develop and test splits.

Document Length-Event Instance Distribution. Figure 9 present the distribution of document length versus event instance.

Discontinuous Nugget Distribution. Figures 10a, 10b, and 10c present the distribution of discontinuous nuggets over nugget types, event types and argument types, respectively.

Overlapping Nugget Distribution. Figures 10d, 10e, and 10f present the distribution of overlapping nuggets over nugget types, event types and argument types, respectively.

Reverse-order Nugget Distribution. Figures 10g, 10h, and 10i present the distribution of reverse-order nuggets over nugget types, event types and argument types, respectively.

Sub-event Distribution. Figures 10j, 10k, and 10l present the distribution of sub-events over event types, argument types and sub-event types, respectively.

H Dataset Annotation Remunerations

During the official annotation stage, annotators spend approximately 364 minutes to annotate 19 consecutive documents, corresponding to an average of 19.1 minutes per document. Annotators are compensated at approximately \$4.5 per document, which is aligned with local wage standards and ensures fair remuneration for the annotation work.

Argument Type	Constrained Types
Aim	APP / MOD / FEA / DST / STR / WEA / TAK
Condition	LIM
Dataset	DST

Table 9: Schema of Purpose (PUR) Event

Argument Type	Constrained Types
Target	APP / MOD / FEA / DST / STR / WEA / TAK
Condition	LIM
Dataset	DST

Table 10: Schema of IntroduceTarget (ITT) Event

Argument Type	Constrained Types
Subject	APP / MOD / FEA / DST
BaseComponent	APP / MOD / FEA / DST
TriedComponent	APP / MOD / FEA / DST
Condition	LIM / E-RWS
Dataset	DST
Target	E-PUR / TAK / STR / WEA / APP / FEA / MOD

Table 11: Schema of RelatedWorkStep (RWS) Event

Argument Type	Constrained Types
Concern	APP / FEA / STR / WEA / MOD / DST
Fault	APP / FEA / STR / WEA / MOD / DST
Condition	LIM / E-RWF / E-RWS
Dataset	DST
Target	E-PUR / TAK / STR / WEA
Extent	DEG

Table 12: Schema of RelatedWorkFault (RWF) Event

Argument Type	Constrained Types
Proposer	OG
Content	APP / FEA / MOD / DST / TAK
Target	E-PUR / TAK / FEA / WEA

Table 13: Schema of Propose (PRP) Event

Argument Type	Constrained Types
Researcher	OG
Content	APP / MOD / FEA / DST / STR / WEA / TAK
Condition	LIM
Dataset	DST
Target	E-PUR / TAK / STR / WEA / APP / FEA / MOD

Table 14: Schema of WorkStatement (WKS) Event

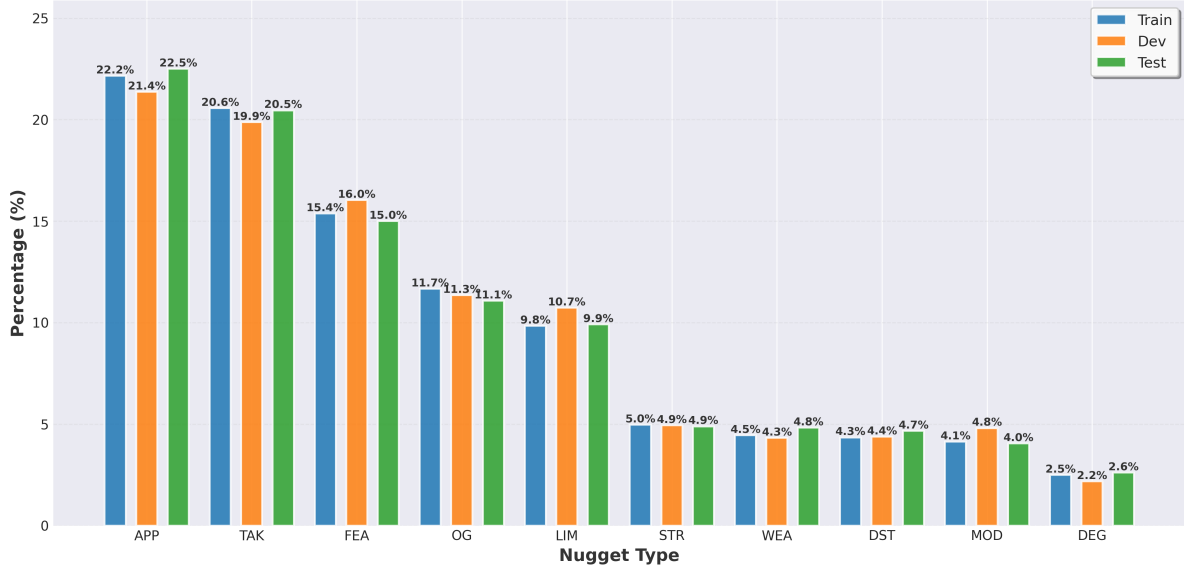


Figure 6: Percentage Distribution of Nugget Types across Train, Develop and Test Splits

Argument Type	Constrained Types
BaseComponent	APP / MOD / FEA / DST
TriedComponent	APP / MOD / FEA / DST
Condition	LIM / E-MDS
Dataset	DST
Target	E-PUR / TAK / STR / WEA / APP / FEA / MOD

Table 15: Schema of MethodStep (MDS) Event

Argument Type	Constrained Types
Finder	OG
Content	E-FAC / E-CMP

Table 16: Schema of Finding (FIN) Event

Argument Type	Constrained Types
Subject	APP / MOD / FEA / STR / WEA / TAK / DST
Object	APP / MOD / FEA / STR / WEA / TAK / DST
Condition	LIM / E-FAC
Reason	LIM / E-FAC
Dataset	DST
Target	E-PUR / TAK / STR / WEA
Extent	DEG

Table 18: Schema of OutcomeFact (FAC) Event

Argument Type	Constrained Types
Arg1	E-FAC / APP / MOD / FEA / DST
Arg2	E-FAC / APP / MOD / FEA / DST
Condition	LIM / E-FAC
Dataset	DST
Result	STR / WEA
Metrics	TAK
Extent	DEG

Table 17: Schema of ExperimentCompare (CMP) Event




Figure 7: Percentage Distribution of Event Types across Train, Develop and Test Splits




Figure 8: Percentage Distribution of Argument Types across Train, Develop and Test Splits

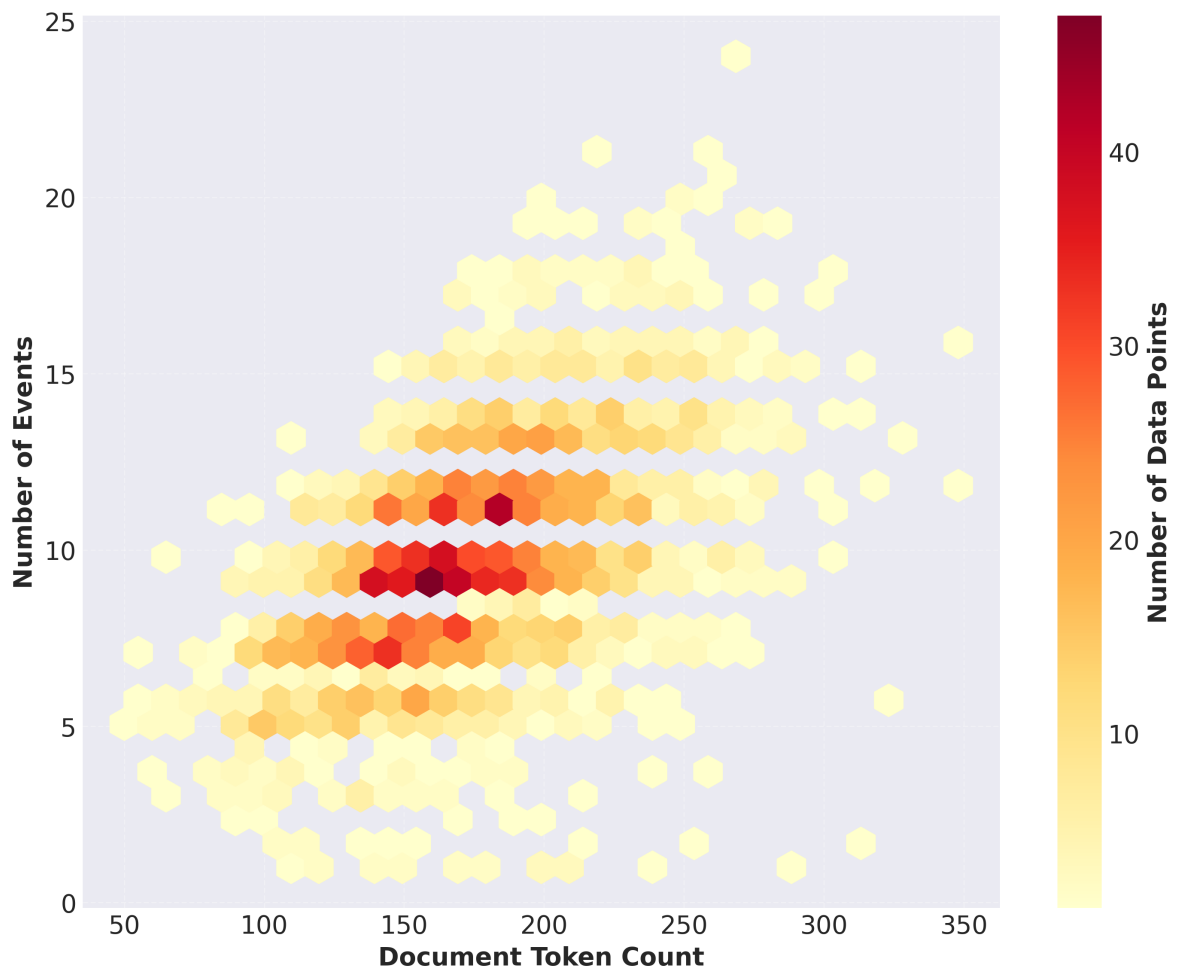





Figure 9: Document Length-Event Instance Distribution



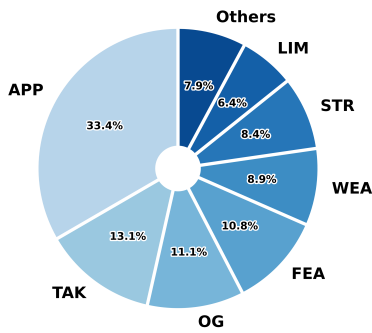
(a) Discontinuous Nugget: Nugget Type



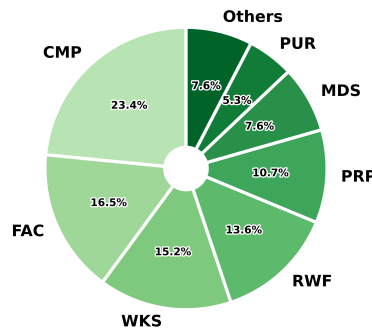
(b) Discontinuous Nugget: Event Type




(c) Discontinuous Nugget: Arg Type




(d) Overlap Nugget: Nugget Type




(e) Overlap Nugget: Event Type




(f) Overlap Nugget: Arg Type




(g) Reverse-Order Nugget: Nugget Type




(h) Reverse-Order Nugget: Event Type




(i) Reverse-Order Nugget: Arg Type



(j) Sub-Event: Event Type



(k) Sub-Event: Argument Type



(l) Sub-Event: Sub-Event Type

Figure 10: Distributions of Complex Nuggets and Events

Algorithm 1: Decoding grid into events

Input: Binary grid $\hat{\mathbf{M}} \in \{0, 1\}^{l \times l \times |R|}$; label vocabulary \mathcal{V} (including HTL, EAL, and THL-types); ontology checker $\text{VALID}(et, at)$.

Output: Event set E in the target format.

Initialize Forward \leftarrow empty map; // HTL: head \rightarrow next tokens

Initialize Tails \leftarrow empty map; // THL-type: head \rightarrow possible tails

Initialize Links \leftarrow empty set; // EAL: (trigger-head, argument-head)

for $i \leftarrow 1$ **to** l **do**

for $j \leftarrow 1$ **to** l **do**

$\mathcal{R}_{i,j} \leftarrow \{r \in R \mid \hat{\mathbf{M}}_{i,j,r} = 1\}$;

foreach $r \in \mathcal{R}_{i,j}$ **do**

if $r = \text{HTL}$ **and** $i \neq j$ **then**

 Forward[i] \leftarrow Forward[i] \cup { j };

else if $r = \text{EAL}$ **then**

 Links \leftarrow Links \cup {(i, j)};

else

 // r is a THL-type label indicating mention type

 Tails[j] \leftarrow Tails[j] \cup { i };

// Step 1: recover nuggets by DFS over HTL and close with THL-type

Initialize Mentions $\leftarrow \emptyset$;

foreach head h in Tails **do**

 Run DFS starting from h following Forward edges to enumerate paths $p = [h, \dots, t]$;

 Keep p only if $t \in \text{Tails}[h]$; // Heuristic (1): must be closed by THL-type

 Add each kept path as a mention span into Mentions;

// Step 2: assign mention types via closing THL-type and split triggers/arguments

Initialize Triggers $\leftarrow \emptyset$, Args $\leftarrow \emptyset$;

foreach mention span $p = [h, \dots, t]$ in Mentions **do**

$\mathcal{T} \leftarrow \{r \in R \mid \hat{\mathbf{M}}_{t,h,r} = 1 \wedge r \neq \text{HTL} \wedge r \neq \text{EAL}\}$;

foreach $\tau \in \mathcal{T}$ **do**

if $\tau \in \text{ET}$ **then**

 Triggers \leftarrow Triggers \cup {(p, τ)}

else

 Args \leftarrow Args \cup {(p, τ)}

// Step 3: build events by linking arguments to triggers via EAL + ontology constraints

Initialize $E \leftarrow \{e = (et, \text{trigger} = p, A = \emptyset) \mid (p, et) \in \text{Triggers}\}$;

foreach (p_a, at) \in Args **do**

 Let h_a be the head (first token) of p_a ;

 Find all triggers (p_t, et) such that (h_t, h_a) \in Links and $\text{VALID}(et, at)$;

if no such trigger exists **then**

continue; // Heuristic (2): drop arguments not attachable to any trigger

foreach matched trigger event e **do**

 Add (at, p_a) into $e.A$;

return E ;

Model		TI(%)			TC(%)		
		P	R	F1	P	R	F1
Global	OneIE	74.18 \pm 0.24	77.33 \pm 0.38	75.72 \pm 0.14	61.68 \pm 0.06	64.23 \pm 0.42	62.93 \pm 0.17
Discriminative	EEQA	77.04 \pm 2.13	72.96 \pm 3.19	74.85 \pm 0.78	64.03 \pm 1.87	60.53 \pm 2.64	62.15 \pm 0.73
	Tagprime	75.72 \pm 1.16	71.04 \pm 1.98	73.27 \pm 0.52	64.85 \pm 1.54	61.36 \pm 1.05	63.03 \pm 0.20
Generative	DEGREE	64.98 \pm 1.72	66.53 \pm 1.08	65.72 \pm 0.70	51.93 \pm 1.80	55.27 \pm 0.42	53.52 \pm 0.75
	KnowCoder	69.83 \pm 0.58	69.95 \pm 0.86	69.88 \pm 0.61	52.02 \pm 0.30	52.03 \pm 0.55	52.02 \pm 0.34
EXCEEDS (Ours)		73.90 \pm 0.78	76.76 \pm 1.44	75.29 \pm 0.32	62.28 \pm 0.78	65.31 \pm 1.14	63.74 \pm 0.14
– Contextual		73.02 \pm 1.38	77.75 \pm 1.38	75.29 \pm 0.21	61.34 \pm 1.46	65.74 \pm 1.26	63.44 \pm 0.67
– Grid Refiner		73.73 \pm 0.73	77.09 \pm 1.01	75.36 \pm 0.27	61.76 \pm 0.97	65.16 \pm 1.02	63.41 \pm 0.67

Table 19: Precision, recall and F1-score (%) of trigger identification (TI) and trigger classification (TC) on SciEvents. EAE-only models are not presented.

Model		AI(%)			AC(%)		
		P	R	F1	P	R	F1
Global	OneIE	28.71 \pm 1.80	32.09 \pm 1.90	30.30 \pm 1.85	27.29 \pm 1.73	30.49 \pm 1.82	28.81 \pm 1.77
	ScentedEAE	50.73 \pm 1.66	29.07 \pm 4.12	36.70 \pm 2.99	49.48 \pm 2.32	28.29 \pm 3.61	35.74 \pm 2.41
Discriminative	EEQA	32.65 \pm 0.97	44.93 \pm 2.45	37.75 \pm 0.46	30.75 \pm 0.78	42.55 \pm 2.53	35.64 \pm 0.59
	PAIE	46.86 \pm 0.26	41.34 \pm 0.59	43.92 \pm 0.22	44.82 \pm 0.31	39.63 \pm 0.65	42.06 \pm 0.34
	Tagprime	47.79 \pm 0.71	41.95 \pm 0.47	44.67 \pm 0.13	45.67 \pm 0.55	40.09 \pm 0.67	42.69 \pm 0.32
	DEEIA	43.06 \pm 2.41	29.33 \pm 0.37	34.86 \pm 0.81	41.05 \pm 2.23	28.06 \pm 0.36	33.30 \pm 0.75
Generative	BartGen	41.21 \pm 1.34	38.60 \pm 0.58	39.85 \pm 0.52	39.10 \pm 1.21	36.63 \pm 0.48	37.81 \pm 0.39
	DEGREE	29.03 \pm 1.33	27.81 \pm 0.50	28.40 \pm 0.88	26.87 \pm 1.21	25.79 \pm 0.43	26.32 \pm 0.79
	KnowCoder	37.82 \pm 0.17	33.00 \pm 0.17	35.24 \pm 0.11	35.89 \pm 0.35	31.29 \pm 0.24	33.43 \pm 0.27
EXCEEDS (Ours)		50.86 \pm 2.68	40.48 \pm 1.92	44.97 \pm 0.28	48.82 \pm 2.81	38.91 \pm 1.66	43.20 \pm 0.29
– Contextual		48.64 \pm 2.16	40.40 \pm 1.70	44.07 \pm 0.51	46.46 \pm 2.35	38.68 \pm 1.43	42.14 \pm 0.39
– Grid Refiner		49.64 \pm 1.24	40.03 \pm 0.94	44.30 \pm 0.51	47.51 \pm 1.37	38.39 \pm 0.87	42.44 \pm 0.59

Table 20: Precision, recall and F1-score (%) of argument identification (AI) and argument classification (AC) on SciEvents. For EAE-only models, trigger predictions are derived from Tagprime.

Model		EC(%)		
		P	R	F1
Global	OneIE	35.26 \pm 1.48	39.85 \pm 1.61	37.41 \pm 1.51
	ScentedEAE	51.21 \pm 2.24	30.12 \pm 2.32	37.88 \pm 2.10
Discriminative	EEQA	44.45 \pm 3.77	45.57 \pm 2.38	44.81 \pm 1.35
	PAIE	49.98 \pm 2.11	44.70 \pm 0.06	47.17 \pm 0.92
	Tagprime	50.22 \pm 1.42	45.53 \pm 1.46	47.72 \pm 0.32
	DEEIA	40.64 \pm 3.35	27.53 \pm 0.79	32.80 \pm 1.67
Generative	BartGen	46.93 \pm 1.37	39.30 \pm 0.87	42.75 \pm 0.11
	DEGREE	39.60 \pm 1.01	23.55 \pm 0.86	29.53 \pm 0.96
	KnowCoder	41.03 \pm 0.94	29.85 \pm 1.22	34.54 \pm 0.82
EXCEEDS (Ours)		48.28 \pm 1.05	48.29 \pm 1.22	48.25 \pm 0.10
– Contextual		47.05 \pm 1.10	48.32 \pm 2.13	47.64 \pm 0.85
– Grid Refiner		47.33 \pm 0.86	48.80 \pm 1.80	48.04 \pm 1.07

Table 21: Precision, recall and F1-score (%) of event correlation (EC) on SciEvents. For EAE-only models, trigger predictions are derived from Tagprime.

Adaptive Compression of Word Embeddings					
Document ID: f63de5c23cce0cc5bb67d42ab12e7bed					
<p>Abstract: Distributed representations of words have been an indispensable component_{E1} for natural language processing (NLP) tasks. However, the large memory footprint_{E2} of word embeddings makes it challenging to deploy NLP models to memory-constrained devices (e.g., self-driving cars, mobile devices). In this paper, we propose_{E3} a novel method to adaptively compress_{E4} word embeddings. We fundamentally follow_{E5} a code-book approach that represents_{E6} words as discrete codes such as (8, 5, 2, 4). However, unlike prior works that assign the same length of codes to all words, we adaptively assign_{E8} different lengths of codes to each word by learning_{E7} downstream tasks. The proposed method works in two steps. First, each word directly learns to select_{E10} its code length in an end-to-end manner by applying_{E9} the Gumbel-softmax tricks. After selecting the code length, each word learns_{E12} discrete codes through_{E11} a neural network with a binary constraint. To showcase_{E14} the general applicability of the proposed method, we evaluate_{E13} the performance on four different downstream tasks. Comprehensive evaluation results clearly show_{E15} that our method is effective_{E16} and makes_{E17} the highly compressed word embeddings without hurting the task accuracy. Moreover, we show_{E18} that our model assigns_{E20} word to each code-book by considering_{E19} the significance of tasks.</p>					
Event ID	Event Type	Trigger	Arg Type	Arg Text	Nugget Type
E1	ITT	component	Target	natural language processing	TAK
E2	RWF	large memory footprint	Concern	word embeddings	MOD
E3	PRP	propose	Proposer Content Target	we method adaptively compress	OG APP E-PUR [†]
E4	PUR*	adaptively compress	Aim	word embeddings	MOD
E5	WKS	follow	Researcher Content Target	we code - book approach represents	OG APP E-PUR [†]
E6	PUR*	represents	Aim Condition	words as discrete codes	FEA LIM
E7	WKS	learning	Content Researcher Target	downstream tasks we assign	TAK OG E-PUR [†]
E8	PUR*	assign	Aim Condition	different lengths of codes to each word	FEA LIM
E9	MDS	applying	Target TriedComponent BaseComponent	select gumbel - softmax tricks word	E-PUR [†] APP FEA
E10	PUR*	select	Aim Condition	code length in an end - to - end manner	TAK LIM
E11	MDS	through	Target BaseComponent TriedComponent Condition	learns word neural network with a binary constraint after selecting the code length	E-PUR [†] FEA APP LIM
E12	PUR*	learns	Aim	discrete codes	TAK
E13	WKS	evaluate	Researcher Content Condition Target	we performance on four different downstream tasks showcase	OG TAK LIM E-PUR [†]
E14	PUR*	showcase	Aim	general applicability	TAK
E15	FIN	show	Content Content	effective makes	E-FAC [†] E-FAC [†]
E16	FAC*	effective	Subject	method	APP
E17	FAC*	makes	Condition Subject Object	without hurting the task accuracy method highly compressed word embeddings	LIM APP STR
E18	FIN	show	Finder Content	we considering	OG E-FAC [†]
E19	FAC*	considering	Object Target Subject	significance of tasks assigns model	TAK E-PUR [†] APP
E20	PUR	assigns	Aim Condition	word to each code - book	FEA LIM

Table 22: Event extraction annotations for the paper *Adaptive Compression of Word Embeddings*. * indicates that the event is a sub-event; † indicates that the argument is a sub-event argument (nugget_type starts with E-).

Nugget Type	Train	Develop	Test	Total
APP	8,860 (22.16%)	1,068 (21.38%)	1,146 (22.50%)	11,074 (22.12%)
TAK	8,219 (20.56%)	993 (19.88%)	1,042 (20.46%)	10,254 (20.48%)
FEA	6,147 (15.38%)	801 (16.03%)	764 (15.00%)	7,712 (15.40%)
OG	4,663 (11.67%)	567 (11.35%)	564 (11.07%)	5,794 (11.57%)
LIM	3,933 (9.84%)	536 (10.73%)	505 (9.92%)	4,974 (9.94%)
STR	1,986 (4.97%)	247 (4.94%)	249 (4.89%)	2,482 (4.96%)
WEA	1,779 (4.45%)	216 (4.32%)	246 (4.83%)	2,241 (4.48%)
DST	1,736 (4.34%)	219 (4.38%)	238 (4.67%)	2,193 (4.38%)
MOD	1,652 (4.13%)	240 (4.80%)	206 (4.04%)	2,098 (4.19%)
DEG	998 (2.50%)	109 (2.18%)	133 (2.61%)	1,240 (2.48%)
Total	39,973	4,996	5,093	50,062

Table 23: Distribution of Nugget Types across Train, Develop and Test Splits

Event Type	Train	Develop	Test	Total
PUR	3,320 (17.04%)	395 (16.16%)	416 (16.99%)	4,131 (16.94%)
WKS	2,733 (14.02%)	354 (14.48%)	332 (13.56%)	3,419 (14.02%)
FAC	2,528 (12.97%)	313 (12.80%)	324 (13.24%)	3,165 (12.98%)
MDS	2,385 (12.24%)	303 (12.39%)	300 (12.25%)	2,988 (12.26%)
PRP	2,094 (10.75%)	235 (9.61%)	257 (10.50%)	2,586 (10.61%)
RWF	1,549 (7.95%)	201 (8.22%)	205 (8.37%)	1,955 (8.02%)
CMP	1,536 (7.88%)	192 (7.85%)	195 (7.97%)	1,923 (7.89%)
FIN	1,480 (7.59%)	192 (7.85%)	175 (7.15%)	1,847 (7.58%)
ITT	1,400 (7.18%)	190 (7.77%)	174 (7.11%)	1,764 (7.24%)
RWS	463 (2.38%)	70 (2.86%)	70 (2.86%)	603 (2.47%)
Total	19,488	2,445	2,448	24,381

Table 24: Distribution of Event Types across Train, Develop and Test Splits

Argument Type	Train	Develop	Test	Total
Content	7,010 (15.55%)	877 (15.61%)	853 (14.92%)	8,740 (15.49%)
Target	6,982 (15.49%)	844 (15.02%)	876 (15.32%)	8,702 (15.43%)
Condition	3,919 (8.69%)	534 (9.51%)	506 (8.85%)	4,959 (8.79%)
Aim	3,574 (7.93%)	424 (7.55%)	443 (7.75%)	4,441 (7.87%)
BaseComponent	2,858 (6.34%)	377 (6.71%)	366 (6.40%)	3,601 (6.38%)
Subject	2,662 (5.91%)	326 (5.80%)	365 (6.38%)	3,353 (5.94%)
TriedComponent	2,143 (4.75%)	274 (4.88%)	279 (4.88%)	2,696 (4.78%)
Researcher	2,067 (4.59%)	256 (4.56%)	250 (4.37%)	2,573 (4.56%)
Object	2,009 (4.46%)	250 (4.45%)	256 (4.48%)	2,515 (4.46%)
Proposer	1,901 (4.22%)	217 (3.86%)	240 (4.20%)	2,358 (4.18%)
Fault	1,577 (3.50%)	191 (3.40%)	231 (4.04%)	1,999 (3.54%)
Arg1	1,329 (2.95%)	164 (2.92%)	155 (2.71%)	1,648 (2.92%)
Result	1,302 (2.89%)	170 (3.03%)	166 (2.90%)	1,638 (2.90%)
Arg2	1,250 (2.77%)	155 (2.76%)	152 (2.66%)	1,557 (2.76%)
Concern	1,066 (2.36%)	149 (2.65%)	144 (2.52%)	1,359 (2.41%)
Extent	998 (2.21%)	109 (1.94%)	133 (2.33%)	1,240 (2.20%)
Dataset	914 (2.03%)	115 (2.05%)	136 (2.38%)	1,165 (2.07%)
Metrics	794 (1.76%)	90 (1.60%)	92 (1.61%)	976 (1.73%)
Finder	695 (1.54%)	94 (1.67%)	74 (1.29%)	863 (1.53%)
Reason	26 (0.06%)	2 (0.04%)	0 (0.00%)	28 (0.05%)
Total	45,076	5,618	5,717	56,411

Table 25: Distribution of Argument Types across Train, Develop and Test Splits