

Document-level Event Extraction with Efficient End-to-end Learning of Cross-event Dependencies

Kung-Hsiang Huang¹ Nanyun Peng^{1,2}

¹ Information Sciences Institute, University of Southern California

² Computer Science Department, University of California, Los Angeles

kunghsia@usc.edu, violetpeng@cs.ucla.edu

Abstract

Fully understanding narratives often requires identifying events in the context of whole documents and modeling the event relations. However, document-level event extraction is a challenging task as it requires the extraction of event and entity coreference, and capturing arguments that span across different sentences. Existing works on event extraction usually confine on extracting events from single sentences, which fail to capture the relationships between the event mentions at the scale of a document, as well as the event arguments that appear in a different sentence than the event trigger. In this paper, we propose an end-to-end model leveraging Deep Value Networks (DVN), a structured prediction algorithm, to efficiently capture cross-event dependencies for document-level event extraction. Experimental results show that our approach achieves comparable performance to CRF-based models on ACE05, while enjoys significantly higher computational efficiency.

1 Introduction

Narratives are account of a series of related events or experiences (Urdang, 1968). Extracting events in literature can help machines better understand the underlying narratives. A robust event extraction system is therefore crucial for fully understanding narratives.

Event extraction aims to identify events composed of a trigger of pre-defined types and the corresponding arguments from plain text (Grishman et al., 2005). To gain full information about the extracted events, entity coreference and event coreference are important, as demonstrated in Figure 1a. These two tasks require document-level modeling. The majority of the previous event extraction works focus on sentence level (Li and Ji, 2014; Huang et al., 2020; Lin et al., 2020). Some later works leverage document-level features, but still extract events at

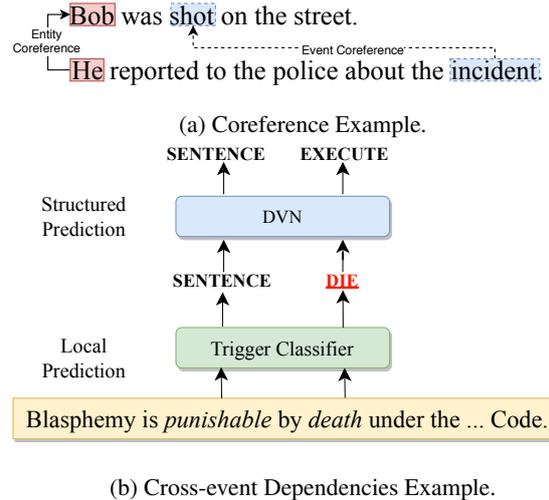


Figure 1: (a) demonstrates why coreference resolution is essential for event extraction. In the second sentence, without entity coreference, an event extraction system cannot identify which real-world entity does *He* refer to. Similarly, *incidence* and *shot* will be incorrectly linked to two different real-world events without event coreference. (b) shows the importance of cross-event dependencies. The local trigger classifier falsely classifies *death* as type DIE. Instead, it is an EXECUTE event as a person’s life is taken away by an authority. A structured prediction model that learns cross-event interactions can potentially infer the correct event type for *death* given the previous SENTENCE event is often carried out by authorities.

the scope of sentence (Yang and Mitchell, 2016; Zhao et al., 2018b; Wadden et al., 2019). More recently, Du and Cardie (2020) and Du et al. (2020) treat document-level event extraction as a template-filling task. Li et al. (2020a) performs event mention extraction and the two coreference tasks independently using a pipeline approach. However, none of the previous works learn entity and event coreference jointly with event mention extraction. We hypothesize that joint learning event mention extraction, event coreference, and entity coreference can result in richer representations and better performance.

Moreover, learning cross-event dependencies is crucial for event extraction. Figure 1b shows a real example from the ACE05 dataset on how learning dependencies among event mentions can help correct errors made by local trigger classifiers. However, efficiency is a challenge when modeling such dependencies at the scale of document. While some works attempted to capture such dependencies with conditional random field or other structured prediction algorithms on hand-crafted features (Li et al., 2013; Lin et al., 2020), these approaches subject to scalability issue and require certain level of human efforts. In this work, we study end-to-end learning methods of an efficient energy-based structured prediction algorithm, Deep Value Networks (DVN), for document-level event extraction.

The contribution of this work is two-fold. First, we propose a document-level event extraction model, DEED (**D**ocument-level **E**vent **E**xtraction with **D**VN). DEED utilizes DVN for capturing cross-event dependencies while simultaneously handling event mention extraction, event coreference, and entity coreference. Using gradient ascent to produce structured trigger prediction, DEED enjoys a significant advantage on efficiency for capturing inter-event dependencies. Second, to accommodate evaluation at the document level, we propose two evaluation metrics for document-level event extraction. Experimental results show that the proposed approach achieve comparable performance with much better training and inference efficiency than strong baselines on the ACE05 dataset.

2 Related Works

In this section, we summarize existing works on document-level information extraction and event extraction, and the application of structured prediction to event extraction tasks.

Document-level Information Extraction Information extraction (IE) is mostly studied at the scope of sentence by early works. (Ju et al., 2018; Qin et al., 2018; Stanovsky et al., 2018). Recently, there has been increasing interest in extracting information at the document-level. Jia et al. (2019) proposed a multiscale mechanism that aggregates mention-level representations into entity-level representations for document-level N -ary relation extraction. Jain et al. (2020) presented a dataset for salient entity identification and document-level N -ary relation extraction in scientific domain. Li et al.

(2020b) utilized a sequence labeling model with feature extractors at different level for document-level relation extraction in biomedical domain. Hu et al. (2020) leveraged contextual information of multi-token entities for document-level named entity recognition. A few studies which tackled document-level event extraction will be reviewed in Section 2.

Document-level Event Extraction Similar to other IE tasks, most event extraction methods make predictions within sentences. Initial attempts on event extraction relied on hand-crafted features and a pipeline architecture (Ahn, 2006; Gupta and Ji, 2009; Li et al., 2013). Later studies gained significant improvement from neural approaches, especially large pre-trained language models (Wadden et al., 2019; Nguyen et al., 2016; Liu et al., 2018; Lin et al., 2020; Balali et al., 2020). Recently, event extraction at the document level gains more attention. Yang et al. (2018) proposed a two-stage framework for Chinese financial event extraction: 1) sentence-level sequence tagging, and 2) document-level key event detection and heuristic-based argument completion. Zheng et al. (2019) transforms tabular event data into entity-based directed acyclic graphs to tackle the *argument scattering* challenge. Du and Cardie (2020) employed a multi-granularity reader to aggregate representations from different levels of granularity. However, none of these approaches handle entity coreference and event coreference jointly. Our work focus on extracting events at the scope of document, while jointly resolving both event and entity coreference.

Structured Prediction on Event Extraction Existing event extraction systems integrating structured prediction typically uses conditional random fields (CRFs) to capture dependencies between predicted events (Xu et al., 2019; Wang et al., 2018). However, CRF is only applicable to modeling linear dependencies, and has scalability issue as the computation cost at least grows quadratically in the size of label. Another line of solutions incorporated beam search with structured prediction algorithms. Li et al. (2013) leveraged structured perceptron to learn from hand-crafted global features. Lin et al. (2020) adopted hand-crafted global features with a global scoring function and uses beam search for inference. While these structured prediction methods can model beyond linear dependencies and alleviate the scalability issue, it requires pre-defined

orders for running beam search. In contrast, our method addresses the above two issues by adopting an efficient structured prediction algorithm, Deep Value Networks, which runs linear in the size of label and does not require pre-defined order for decoding.

3 Document-level Event Extraction

3.1 Task Definition

The input to the document-level event extraction task is a document of tokens $\mathcal{D} = \{d_0, d_1, \dots, d_m\}$, with spans $\mathcal{S} = \{s_0, s_1, \dots, s_n\}$ generated by iterating k-grams in each sentence (Wadden et al., 2019). Our model aims to jointly solve event mention extraction, event coreference, and entity coreference.

Event Mention Extraction refers to the subtask of 1) identifying event triggers in \mathcal{D} by predicting the event type for each token d_i . 2) Then, given each trigger, corresponding arguments in \mathcal{S} and argument roles are extracted. This task is similar to the sentence-level event extraction task addressed by previous studies (Wadden et al., 2019; Lin et al., 2020). The difference is that we require extracting *full spans* of all name, nominal, and pronoun arguments, while these works focus on extracting *head spans* of *name* arguments. **Entity Coreference** aims to find which entity mentions refer to the same entity. Our model predicts the most likely antecedent span s_j for each span s_i . **Event Coreference** is to recognize event mentions that are co-referent to each other. Similar to entity coreference, we predict the most likely antecedent trigger d_j for each predicted trigger d_i . **Entity Extraction** is performed as an auxiliary subtask for richer representations. Each entity mention corresponds to a span s_i in \mathcal{S} .

3.2 Task Evaluation

Evaluation metrics used by previous sentence-level event extraction studies (Wadden et al., 2019; Zheng et al., 2019; Lin et al., 2020) are not suitable for our task as event coreference and entity coreference are not considered. Du and Cardie (2020) evaluates entity coreference using bipartite matching. However, it does not consider event coreference and less informative arguments (nominal and pronoun). As a solution, we propose two metrics: DOCTRIGGER and DOCARGUMENT, to properly evaluate event extraction at the document level. The purpose is to conduct evaluation on

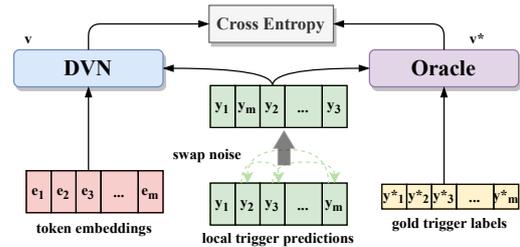


Figure 2: Use swap noise to enable DVN to continue learning from the oracle value function even when the local trigger classifier overfits on the training set.

event coreference clusters and argument coreference clusters. **DOCTRIGGER** considers trigger span, event type, and event coreference. Triggers in the same event coreference chain are clustered together. The metric first aligns gold and predicted trigger clusters, and computes a matching score between each gold-predicted trigger cluster pair. A predicted trigger cluster gets full score if all the associated triggers are correctly identified. To enforce the constraint that one gold trigger cluster can only be mapped to at most one predicted trigger cluster, Kuhn–Munkres algorithm (Kuhn, 1955) is adopted. **DOCARGUMENT** considers argument span, argument role, and entity coreference. We define an argument cluster as an argument with its co-referent entity mentions. Similar to DOCTRIGGER, DOCARGUMENT uses Kuhn–Munkres algorithm to align gold and predicted argument clusters, and compute a matching score between each argument cluster pair. An event extraction system should get full credits in DOCARGUMENT as long as it identifies the most informative co-referent entity mentions and does not predict false positive co-referent entity mentions.¹ Details of the evaluation metric are included in Appendix C.

4 Proposed Approach

We develop a base model that makes independent predictions for each subtask under a multi-task IE framework. The proposed end-to-end framework, DEED, then incorporates DVN into the base model to efficiently capture cross-event dependencies.

4.1 Base Model

Our BASE model is built on a span-based IE framework, DYGIE++ (Wadden et al., 2019). DYGIE++ learns entity classification, entity coreference, and event extraction jointly. The base model extends

¹We set the weights for name, nominal, and pronoun to be 1, 0.5, and 0.25, inspired by Chen and Ng (2013).

the entity coreference module of DYGIE++ to handle event coreference.

Encoding Ideally, we want to encode all tokens in a document $\mathcal{D} = \{d_1, d_2, \dots, d_m\}$ with embeddings that covers the context of the entire document. However, due to hardware limitation for long documents, each document is split into multi-sentences. Each multi-sentence corresponds to a chunk of consecutive sentences. We obtain rich contextualized embeddings for each multi-sentence of tokens $e = \{e_1, e_2, \dots, e_n\}$ using BERT-BASE (Devlin et al., 2019).

Span Enumeration Conventional event extraction systems use BIO tag scheme to identify the starting and ending position of each trigger and entity. Nevertheless, this method fails to handle nested entities. As a solution, we enumerate all possible spans to generate event mention and entity mention candidates from uni-gram to k -gram.² Each span s_i is represented by corresponding head token e_h , tail token e_t and the distance embeddings $c_{h,t}$, denoted as $\mathbf{x}_i = [e_h, e_t, c_{h,t}]$, following Wadden et al. (2019).

Classification We use task-specific feed-forward networks (FFN) to compute the label probabilities. Trigger extraction is performed on each token $\mathbf{y}_i^{trig} = \text{FFN}^{trig}(e_i)$, while entity extraction is done on each span $\mathbf{y}_i^{ent} = \text{FFN}^{ent}(\mathbf{x}_i)$. For argument extraction, event coreference, and entity coreference, we score each pair of candidate spans $\mathbf{y}_k^t = \text{FFN}^t([\mathbf{x}_i, \mathbf{x}_j])$, where t refers to a specific task. Cross-entropy loss is used to learn trigger extraction, argument extraction as follows

$$\mathcal{L}^t = \frac{1}{N^t} \sum_{i=1}^{N^t} \mathbf{y}_i^{t*} \log \mathbf{y}_i^t,$$

, where \mathbf{y}^{t*} denotes the ground truth labels, N^t denotes the number of instances, and t denotes different tasks.

For entity coreference and event coreference, BASE optimizes marginal log-likelihood for all correct coreferent spans given candidate spans.

$$\mathcal{L}^t = \log \prod_{i=1}^N \sum_{j \in \text{COREF}(i)} \mathbf{y}_{(i,j)}^t,$$

² k is empirically determined to be 12.

where $\text{COREF}(i)$ denotes the gold set of spans coreferent with candidate span i , and t denotes different tasks. The total loss function for BASE is the weighted sum of all tasks:

$$\mathcal{L}^{\text{BASE}} = \sum_t \beta^t \mathcal{L}^t,$$

β^t is the loss weight for task t .

4.2 Cross-event Dependencies

A main issue for document-level event extraction is the increased complexity for capturing event dependencies. Due to larger number of events at the scope of document, efficiency is a key challenge to modeling inter-event interactions. We incorporate DVN (Gygli et al., 2017) into BASE to solve this issue given its advantage in computation efficiency.

Deep Value Networks DVN is an energy-based *structured prediction* architecture $v(\mathbf{x}, \mathbf{y}; \theta)$ parameterized over θ that learns to evaluate the compatibility between a structured prediction \mathbf{y} and an input \mathbf{x} . The objective of $v(\mathbf{x}, \mathbf{y}; \theta)$ is to approximate an *oracle value function* $v^*(\mathbf{y}, \mathbf{y}^*)$, a function which measures the quality of the output \mathbf{y} in comparison to the groundtruth \mathbf{y}^* , *s.t.* $\forall \mathbf{y} \in \mathcal{Y}, v(\mathbf{x}, \mathbf{y}; \theta) \approx v^*(\mathbf{y}, \mathbf{y}^*)$. The final evaluation metrics are usually used as the *oracle value function* $v^*(\mathbf{y}, \mathbf{y}^*)$. For simplicity, we drop the parameter notion θ , and use $v(\mathbf{x}, \mathbf{y})$ to denote DVN instead.

The inference aims to find $\hat{\mathbf{y}} = \text{argmax}_{\mathbf{y}} v(\mathbf{x}, \mathbf{y})$ for every pair of input and output. A local optimum of $v(\mathbf{x}, \mathbf{y})$ can be efficiently found by performing gradient ascent that runs linear in the size of label. Given DVN’s higher scalability compared with other structured prediction algorithms, we leverage DVN to capture cross-event dependencies.

Deep Value Networks Integration Local trigger classifier predicts the *event type scores* for each token independently. DVN takes in predictions from local trigger classifiers \mathbf{y}^{trig} and embeddings of all tokens e as inputs. Structured outputs $\hat{\mathbf{y}}^{trig}$ should correct errors made by the local trigger classifier due to uncaptured cross-event dependencies. $\hat{\mathbf{y}}^{trig}$ is obtained by performing h -iteration updates on local trigger predictions \mathbf{y}^{trig} using gradient ascent,³

³We set $h=20$ for best empirical performance.

$$\begin{aligned} \mathbf{y}^{t+1} &= \mathcal{P}_y(\mathbf{y}^t + \alpha \frac{d}{d\mathbf{y}} v(\mathbf{e}, \mathbf{y}^t)) \\ \hat{\mathbf{y}}^{trig} &= \mathbf{y}^h, \end{aligned} \quad (1)$$

where $\mathbf{y}^1 = \mathbf{y}^{trig}$, α denotes the inference learning rate, and \mathcal{P}_y denotes a function that clamps inputs into the range $(0, 1)$. The most likely event type for token i is determined by computing $\text{argmax}(\hat{\mathbf{y}}_i^{trig})$.

End-to-end DVN Learning We train DEED in an end-to-end fashion by directly feeding the local trigger predictions to both DVN and the oracle value function. The trigger classification F_1 metric adopted by previous works (Wadden et al., 2019; Lin et al., 2020) is used as the *oracle value function* $v^*(\mathbf{y}^{trig}, \mathbf{y}^{trig*})$. To accommodate continuous outputs, $v^*(\mathbf{y}^{trig}, \mathbf{y}^{trig*})$ needs to be relaxed. We relaxed the output label for each token from $[0, 1]$ to $(0, 1)$. Union and intersection set operations for computing the F_1 scores are replaced with element-wise minimum and maximum operations, respectively. The relaxed *oracle value function* is denoted as $\underline{v}^*(\mathbf{y}^{trig}, \mathbf{y}^{trig*})$. The loss function for the trigger DVN is the following:

$$\begin{aligned} \mathcal{L}^{DVN} &= \sum_{\mathbf{y}^{trig}} -\underline{v}^*(\mathbf{y}^{trig}, \mathbf{y}^{trig*}) \log v(\mathbf{e}, \mathbf{y}^{trig}) \\ &\quad - (1 - \underline{v}^*(\mathbf{y}^{trig}, \mathbf{y}^{trig*})) \log(1 - v(\mathbf{e}, \mathbf{y}^{trig})). \end{aligned} \quad (2)$$

The total loss function for training DEED end-to-end is the summation of BASE loss and DVN loss,

$$\mathcal{L}^{DEED} = \mathcal{L}^{BASE} + \mathcal{L}^{DVN}.$$

Noise Injection However, in this training setup, DVN observes a large portion of high scoring examples at the later stage of training process when the local trigger classifier starts to overfit on the training examples. A naive solution is feeding random noise to train DVN in addition to the outputs of local trigger classifier. Yet, the distribution of these noise are largely distinct from the output of trigger classifier, and therefore easily distinguishable by DVN. Thus, we incorporate swap noise into the local trigger predictions, where $s\%$ of the local trigger outputs \mathbf{y}^{trig} are swapped, as depicted

in Figure 2.⁴ This way, noisy local trigger predictions have similar distributions to the original trigger predictions. We also hypothesize that higher-confident predictions are often easier to identify, and swapping higher-confident trigger predictions may not help DVN learn. We experimented swapping only the lower-confident trigger predictions.

5 Experiments

5.1 Experimental Setup

Our models are evaluated on the ACE05 dataset, containing event, relation, entity, and coreference annotations. Experiments are conducted at the document level instead of sentence level as previous works (Wadden et al., 2019; Lin et al., 2020).

5.2 Baselines and Model Variations

We compare DEED with three baselines: (1) BASE, the base model described in Section 4.1; (2) BCRF extends BASE by adding a CRF layer on top of the trigger classifier; (3) OneIE⁺ is a pipeline composed of the joint model presented in Lin et al. (2020) and coreference modules adapted from BASE. Lin et al. (2020) is the state-of-the-art sentence-level event extraction model that utilizes beam search and CRF with global features to model cross sub-task dependencies. For fair comparison, all models are re-trained using BERT-BASE (Devlin et al., 2019) as the encoder.

In addition to the original DEED model, we consider three variations of it, as discussed in Section 4.2. **DEED w/RN** incorporates random noise while learning DVN, whereas **DEED w/SN** integrates swap noise. **DEED w/SNLC** is an extension of **DEED w/SN**, where swap noise is only applied to lower-confident trigger predictions.

5.3 Overall Results

The overall results are summarized in Table 1. To measure the overall performance, a combined score (*Comb.*) is computed by multiplying DOCTRIGGER F_1 and DOCARGUMENT F_1 . DEED and BCRF achieve huge improvement on all metrics over BASE, suggesting the importance of cross-event dependency modeling for our task. Adding random noise or swap noise to train DVN both improve upon the vanilla training method. OneIE⁺ achieves the best DOCARGUMENT performance,

⁴ s is empirically set to 20

Model	DOCTRIGGER			DOCARGUMENT			Comb.
	Prec.	Rec.	F1	Prec.	Rec.	F1	
BASE	71.25	60.94	65.69	43.75	48.65	46.07	17.13
BCRF	71.87	65.18	68.36	49.84	52.16	50.97	34.84
OneIE ⁺	71.96	62.04	66.63	49.64	56.58	52.88	35.23
DEED	70.97	62.90	66.70	46.13	51.34	48.60	32.42
w/ RN	71.69	65.76	68.59	48.52	52.53	50.44	34.60
w/ SN	70.87	64.02	67.28	43.76	55.15	48.80	32.83
w/ SNLC	73.89	64.98	69.14	48.00	55.27	51.38	35.52

Table 1: Experimental results on ACE05 using document-level evaluation metrics. *RN*: random noise; *SN*: swap noise; *SNLC*: swap noise applying to lower-confident predicted triggers.

Model	Trig-I	Trig-C	Arg-I	Arg-C	Evt-Co	Ent-Co
BCRF	73.92	70.57	51.77	48.31	54.02	74.23
BASE	71.97	68.17	47.95	44.57	43.95	71.88
OneIE ⁺	73.91	71.01	57.19	53.89	42.75	77.00
DEED	73.68	69.62	52.35	48.24	53.85	75.77
w/ RN	72.33	68.20	51.33	48.66	49.86	74.39
w/ SN	74.19	69.54	51.27	48.10	48.94	75.60
w/ SNLC	75.06	71.73	55.12	52.09	50.11	76.98

Table 2: A breakdown of evaluation for each component in F1 evaluated on ACE05. *Trig*: trigger; *Arg*: argument; *I*: identification; *C*: classification; *Evt-Co*: event coreference; *Ent-Co*: entity coreference.

Model	Training (sec/ multi-sent)	Inference (sec/ doc)
BASE	0.52	1.50
BCRF	2.55	9.10
OneIE ⁺	1.21	15.89
DEED	0.71	1.52

Table 3: Comparison of training and inference time, evaluated on the training set and the dev set.

while **DEED w/SNLC** achieves the highest DOC-TRIGGER score and combined score.

6 Analysis

6.1 Performance of Each Component

To understand the capabilities of each module, we show an evaluation breakdown on each component following previous works (Wadden et al., 2019; Lin et al., 2020) in Table 2.⁵ Both BCRF and DEED obtain significant performance gain over BASE across all tasks. In terms of trigger-related tasks, *Trig-I* and *Trig-C*, **DEED w/SNLC** achieves the highest scores. Yet, BCRF performs the best on *Evt-Co*. This explains the close performance of **DEED w/SNLC** and BCRF on DOCTRIGGER,

⁵These studies focus on extracting head span of name argument, while we extract full span of all types of arguments.

as shown in Table 1. In terms of argument-related tasks, OneIE⁺ achieves the best performance on *Arg-I* and *Arg-C*. This suggests that cross-subtask modeling can be important to improve argument extraction. *Arg-I* and *Arg-C* are much lower than the reported scores by previous studies (Wadden et al., 2019; Lin et al., 2020). This suggests the difficulty of extracting full span of pronoun and nominal arguments.

6.2 Computation Time

Table 3 describes the computation time of different models. DEED only requires slightly more computation time in both training and inference time than BASE. By contrast, compared to BCRF, DEED is $\sim 3.5x$ faster in training time and $\sim 6x$ faster in inference time. This demonstrates the efficiency of our approach given the little increase in computation time and the significant performance gain comparable to BCRF detailed in Tables 1 and 2. We also added experiments with OneIE⁺ as a reference, but the comparison focuses on end-to-end frameworks.

Training Method	Loss (Cross Entropy)
Original	0.3613
RN	0.7451
SN	0.2393
SNLC	0.2298

Table 4: The average DVN loss of different DEED training methods on the test set. The lower the loss, the closer between DVN and the *oracle value function*.

6.3 Value Function Approximation

To show that the performance gain of DEED is resulted from improved capabilities of DVN in judging the structure of predicted triggers, we investigate how close DVN approximates the oracle value function under different training settings. We use cross entropy loss as the distance function between the output of DVN and and output of the oracle value function on the test set. The lower the loss is, the closer between the output of DVN and the output of the oracle value function. Table 4 shows the approximation results. The SNLC variation (swap noise applying to lower-confident predicted triggers) yields the lowest loss comparing to the base model and other variations. Along with the results shown in Table 2, we show that lower DVN loss results in better trigger scores. This demonstrates that integrating noise into DVN training procedure is effective in learning better DVN and obtaining better overall performance.

6.4 Error Analysis

We manually compared gold and predicted labels of event mentions on the ACE05 test set and analyzed the mistakes made by our model. These errors are categorized as demonstrated in Figure 3.

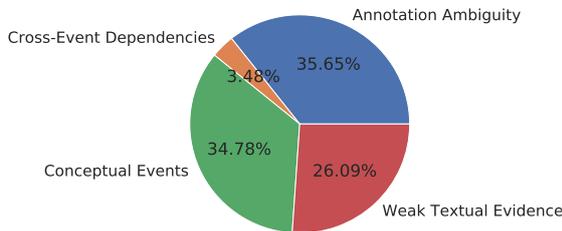


Figure 3: Distribution of errors made by DVN on the ACE05 test set.

Annotation ambiguity A significant portion of the false positive errors are caused by the ambiguity of the task. Such ambiguity can result in disagreement between human annotators. For example,

Lebanese Prime Minister Rafiq Hariri submitted his *resignation* Tuesday and it was accepted

by President Emile Lahoud.

In the sentence above, the trigger label for token *resignation* should be END-POSITION, according to the annotation guideline. Yet, it is not annotated as a trigger in gold annotation. In other cases, two sentences with similar structures contain inconsistent gold annotation, such as:

Separately, *former* WorldCom CEO Bernard Ebbers failed on April 29 to make a first repayment of 25 million dollars ...

Former senior banker Callum McCarthy begins what is one of the most important jobs in London ’s financial world in September

The two examples above share similar context. However, the *former* in the first sentence is not involved with any event, whereas the *former* in the second sentence is annotated as an END-POSITION typed trigger.

Conceptual Events Another common source of false positive errors is extracting “conceptual” events, which did not happen or may happen in the future. For instance,

... former WorldCom CEO Bernard Ebbers failed on April 29 to make a first *repayment* of 25 million dollars ...

Our model predicts the word *repayment* as an TRANSFER-MONEY, which is true if it indeed happened, except it *failed*, as indicated in the beginning of the sentence. To handle this type of error, models need to be aware of the tense and whether there is a negative sentiment associated with the predicted events.

Weak Textual Evidence Our model commonly made false negative errors in cases where the textual information is vague.

But both men observed an uneasy truce over US concerns about Russian *aid* to the nuclear program of Iran ...

In the above sentence, DVN fails to identify the token *aid* as a trigger of type TRANSFER-MONEY. In fact, it is hard to determine whether the *aid* is monetary or military given the context of the whole document. In this case, models have to be aware of information from other sources, such as knowledge bases or other news articles.

Cross-event Dependencies Although our model is able to correct many mistakes made by BASE that requires modeling of cross-event dependencies, as

		Within sentence	Cross sentence
BASE	Correct	161	126
	Incorrect	71	45
DEED	Correct	166	136
	Incorrect	66	35

Table 5: Trigger predictions comparison between BASE and DEED. *Cross sentence* refers to triggers with co-referent triggers that lie in different sentences.

demonstrated in Table 5, there are still a few cases where our model fails.

... after the city 's bishop committed *suicide* over the 1985 blasphemy law . Faisalabad 's Catholic Bishop John Joseph , who had been campaigning against the law , *shot* himself in the head outside a court in Sahiwal district when the judge ... himself in the head outside a court

In the above example, DVN correctly predict *suicide* as a DIE typed trigger, but falsely predict *shot* as type ATTACK instead of type DIE. If our model could capture the interactions between *suicide* and *shot*, it would be able to process this situation. There is still room to improve in cross-event dependency modeling.

7 Conclusion

In this paper, we investigate document-level event extraction that requires joint modeling of event and entity coreference. We propose a document-level event extraction framework, DEED, which uses DVN to capture cross-event dependencies, and explore different end-to-end learning methods of DVN. Experimental results show that DEED achieves comparable performance to competitive baseline models, while DEED is much favorable in terms of computation efficiency. We also found that incorporating noise into end-to-end DVN training procedure can result in higher DVN quality and better overall performance.

8 Ethics

Biases have been studied in many information extraction tasks, such as relation extraction (Gaut et al., 2020), named entity recognition (Mehrabi et al., 2020), and coreference resolution (Zhao et al., 2018a). Nevertheless, not many works investigate biases in event extraction tasks, particularly ACE05.

We analyze the portion of male pronouns (he, him, and his) and female pronouns (she and her) in the

ACE05 dataset. In total, there are 2780 male pronouns, while only 970 female pronouns appear in the corpus. We would expect the trained model to perform better when extracting events where male arguments are involved, and make more mistakes for event involving female arguments due to the significant imbalance between male and female entity annotation. After analyzing the performance of DEED w/ SNLC on the test set, we found that it scores 54.90 and 73.80 on *Arg-C F₁* for male and female pronoun arguments, respectively. Surprisingly, our model is better at identifying female pronoun arguments than male pronoun arguments.

While our proposed framework may not subject to gender biases in ACE05, whether such issue can occur when our model is deployed for public use is unknown. Rigorous studies on out-of-domain corpus is needed to answer this question.

Acknowledgements

We appreciate insightful feedback from PLUSLab members and the anonymous reviewers. This research was sponsored by the Intelligence Advanced Research Projects Activity (IARPA), via Contract No. 2019-19051600007. The views and conclusions of this paper are those of the authors and do not reflect the official policy or position of IARPA or the US government.

References

- David Ahn. 2006. [The stages of event extraction](#). In *Proceedings of the Workshop on Annotating and Reasoning about Time and Events*, pages 1–8, Sydney, Australia. Association for Computational Linguistics.
- Ali Balali, Masoud Asadpour, Ricardo Campos, and Adam Jatowt. 2020. [Joint event extraction along shortest dependency paths using graph convolutional networks](#). *Knowledge-Based Systems*, 210:106492.
- Chen Chen and Vincent Ng. 2013. [Linguistically aware coreference evaluation metrics](#). In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 1366–1374, Nagoya, Japan. Asian Federation of Natural Language Processing.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–

- 4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xinya Du and Claire Cardie. 2020. [Document-level event role filler extraction using multi-granularity contextualized encoding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8010–8020, Online. Association for Computational Linguistics.
- Xinya Du, Alexander Rush, and Claire Cardie. 2020. Document-level event-based extraction using generative template-filling transformers. *arXiv preprint arXiv:2008.09249*.
- Andrew Gaut, Tony Sun, Shirlyn Tang, Yuxin Huang, Jing Qian, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2020. [Towards understanding gender bias in relation extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2943–2953, Online. Association for Computational Linguistics.
- R. Grishman, D. Westbrook, and Adam Meyers. 2005. Nyu’s english ace 2005 system description.
- Prashant Gupta and Heng Ji. 2009. [Predicting unknown time arguments based on cross-event propagation](#). In *Proceedings of the ACL-IJCNLP 2009 Conference Short Papers*, pages 369–372, Suntec, Singapore. Association for Computational Linguistics.
- Michael Gygli, Mohammad Norouzi, and A. Angelova. 2017. Deep value networks learn to evaluate and iteratively refine structured outputs. In *ICML*.
- Anwen Hu, Zhicheng Dou, Jian-Yun Nie, and Ji-Rong Wen. 2020. Leveraging multi-token entities in document-level named entity recognition. In *AAAI*, pages 7961–7968.
- Kung-Hsiang Huang, Mu Yang, and Nanyun Peng. 2020. [Biomedical event extraction with hierarchical knowledge graphs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1277–1285, Online. Association for Computational Linguistics.
- Sarthak Jain, Madeleine van Zuylen, Hannaneh Hajishirzi, and Iz Beltagy. 2020. [SciREX: A challenge dataset for document-level information extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7506–7516, Online. Association for Computational Linguistics.
- Robin Jia, Cliff Wong, and Hoifung Poon. 2019. [Document-level n-ary relation extraction with multi-scale representation learning](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3693–3704, Minneapolis, Minnesota. Association for Computational Linguistics.
- Meizhi Ju, Makoto Miwa, and Sophia Ananiadou. 2018. [A neural layered model for nested named entity recognition](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1446–1459, New Orleans, Louisiana. Association for Computational Linguistics.
- Harold W Kuhn. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97.
- Manling Li, Alireza Zareian, Ying Lin, Xiaoman Pan, Spencer Whitehead, Brian Chen, Bo Wu, Heng Ji, Shih-Fu Chang, Clare Voss, Daniel Napierski, and Marjorie Freedman. 2020a. [GAIA: A fine-grained multimedia knowledge extraction system](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 77–86, Online. Association for Computational Linguistics.
- Qi Li and Heng Ji. 2014. [Incremental joint extraction of entity mentions and relations](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 402–412, Baltimore, Maryland. Association for Computational Linguistics.
- Qi Li, Heng Ji, and Liang Huang. 2013. [Joint event extraction via structured prediction with global features](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 73–82, Sofia, Bulgaria. Association for Computational Linguistics.
- Zhiheng Li, Zhihao Yang, Yang Xiang, Ling Luo, Yuanyuan Sun, and Hongfei Lin. 2020b. Exploiting sequence labeling framework to extract document-level relations from biomedical texts. *BMC bioinformatics*, 21:1–14.
- Ying Lin, Heng Ji, Fei Huang, and Lingfei Wu. 2020. [A joint neural model for information extraction with global features](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7999–8009, Online. Association for Computational Linguistics.
- Xiao Liu, Zhunchen Luo, and Heyan Huang. 2018. Jointly multiple events extraction via attention-based graph information aggregation. In *EMNLP*.
- Ninareh Mehrabi, Thammie Gowda, Fred Morstatter, Nanyun Peng, and Aram Galstyan. 2020. [Man is to person as woman is to location: Measuring gender bias in named entity recognition](#). In *Proceedings of the 31st ACM Conference on Hypertext and Social Media, HT ’20*, page 231–232, New York, NY, USA. Association for Computing Machinery.
- Thien Huu Nguyen, Kyunghyun Cho, and Ralph Grishman. 2016. [Joint event extraction via recurrent neural networks](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 300–309, San Diego, California. Association for Computational Linguistics.
- Pengda Qin, Weiran Xu, and William Yang Wang.

2018. **Robust distant supervision relation extraction via deep reinforcement learning**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2137–2147, Melbourne, Australia. Association for Computational Linguistics.
- Gabriel Stanovsky, Julian Michael, Luke Zettlemoyer, and Ido Dagan. 2018. **Supervised open information extraction**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 885–895, New Orleans, Louisiana. Association for Computational Linguistics.
- Laurence Urdang. 1968. *The Random House dictionary of the English language*. New York : Random House.
- David Wadden, Ulme Wennberg, Yi Luan, and Hananeh Hajishirzi. 2019. Entity, relation, and event extraction with contextualized span representations. In *EMNLP/IJCNLP*.
- Yan Wang, Jian Wang, Hongfei Lin, Xiwei Tang, Shaowu Zhang, and Lishuang Li. 2018. Bidirectional long short-term memory with crf for detecting biomedical event trigger in fasttext semantic space. *BMC bioinformatics*, 19(20):507.
- Meng Xu, Xin Zhang, and Lixiang Guo. 2019. Jointly detecting and extracting social events from twitter using gated bilstm-crf. *IEEE Access*, 7:148462–148471.
- Bishan Yang and Tom M. Mitchell. 2016. **Joint extraction of events and entities within a document context**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 289–299, San Diego, California. 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.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018a. **Gender bias in coreference resolution: Evaluation and debiasing methods**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 15–20, New Orleans, Louisiana. Association for Computational Linguistics.
- Yue Zhao, Xiaolong Jin, Yuanzhuo Wang, and Xueqi Cheng. 2018b. **Document embedding enhanced event detection with hierarchical and supervised attention**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (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 the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 337–346, Hong Kong, China. Association for Computational Linguistics.

A Data Statistics

The statistics of ACE05 are shown in Table 6. We observe that the event coreference annotation is very sparse.

Split	Docs	Events	Entities	Ent-C	Evt-C
Train	529	4202	47569	6814	482
Dev	28	450	3423	553	45
Test	40	403	3673	577	58

Table 6: Data statistics of ACE05. *Ent-C* and *Evt-C* denote the number of entity and event coreference clusters, respectively.

B Implementation Details

We adopted part of the pre-processing pipelines from Wadden et al. (2019) for data cleaning and dataset splitting.

BASE, BCRF, and DVN are optimized with BERTADAM for 250 epochs with batch size of 16. BERT-BASE is fine-tuned with learning rate of $1e-4$ and no decay, while the other components are trained with learning rate of $1e-3$ and weight decay of $1e-2$. Training is stopped if the dev set *Arg-C* F_1 score does not improve for 15 consecutive epochs. OneIE⁺ is trained with the default parameters described in Lin et al. (2020). All experiments are conducted on a 12-CPU machine running CentOS Linux 7 (Core) and NVIDIA RTX 2080 with CUDA 10.1.

C Document-level Evaluation Metrics

D Development Set Performance

Algorithm 1 Document-level Trigger Evaluation Metric

```
1: function DOCTRIGGER(gold events  $G$ , predicted events  $P$ )
2:   Let  $match = false\text{-}alarm = miss = hit = 0$ 
3:   Let  $M$  be a trigger matching matrix.
4:   for  $g$  in  $G$ .triggers do
5:     for  $p$  in  $P$ .triggers do
6:       if ! SAMEEVENTTYPE( $g, p$ ) then
7:          $match = 0$ 
8:       else
9:          $match = \text{Trig-I}(p, g)$ 
10:      end if
11:       $M[g.idx, p.idx] = match$ 
12:    end for
13:  end for
14:   $assignments = \text{KUHN-MUNKRES}(M)$ 
15:  for  $i, j$  in  $assignments$  do
16:    if  $G$ .triggers[ $i$ ] is null then
17:       $false\text{-}alarm += 1$ 
18:    else if  $P$ .triggers[ $j$ ] is null then
19:       $miss += 1$ 
20:    else
21:       $match += M[i][j]$ 
22:       $hit += 1$ 
23:    end if
24:  end for
25:  return ( $match, false\text{-}alarm, miss, hit$ )
26: end function
```

Model	Trig-I	Trig-C	Arg-I	Arg-C	Evt-Co	Ent-Co
BASE	74.63	70.49	56.82	52.41	30.64	67.31
BCRF	76.53	72.89	59.62	54.47	33.16	68.72
OneIE ⁺	76.78	73.56	63.12	59.32	35.81	70.78
DEED	77.11	72.31	62.42	55.80	31.90	69.57
w/ RN	75.74	70.94	61.45	55.18	34.88	68.56
w/ SN	77.81	74.53	61.90	55.52	38.55	69.48
w/ SNLC	76.76	72.13	62.78	57.45	31.32	70.78

Table 7: A breakdown of evaluation on the dev set for each model. The corresponding test set performance is shown in Table 2.

Algorithm 2 Document-level Argument Evaluation Metric

```

1: function DOCARGUMENT(gold events  $G$ , pre-
   predicted events  $P$ )
2:   Let  $match = false\text{-}alarm = miss = hit = 0$ 
3:   Let  $M$  be an argument matching matrix.
4:   for  $g$  in  $G.arguments$  do
5:     for  $p$  in  $P.arguments$  do
6:        $M[i, j] = ARGMATCH(g, p)$ 
7:     end for
8:   end for
9:    $assignments = KUHN\text{-}MUNKRES(M)$ 
10:  for  $i, j$  in  $assignments$  do
11:    if  $G.arguments[i]$  is null then
12:       $false\text{-}alarm += 1$ 
13:    else if  $P.arguments[j]$  is null then
14:       $miss += 1$ 
15:    else
16:       $match += M[i][j]$ 
17:       $hit += 1$ 
18:    end if
19:  end for
20:  return ( $match, false\text{-}alarm, miss, hit$ )
21: end function

```

Algorithm 3 Argument match called by Algorithm 2

```

1: function ARGMATCH(gold argument cluster
    $g$ , predicted argument cluster  $p$ )
2:   if not SAMEROLE( $g, p$ ) or not
3:     not SAMEEVENTTYPE( $g, p$ ) then
4:     return 0
5:   end if
6:    $BMA = BESTMATCHEDARGUMENT(p, g)$ 
7:    $w = GETWEIGHT(BMA)$   $\triangleright$  The weights
   for name, nominal, pronoun are 1, 0.5, 0.25.
8:    $false\text{-}alarm = |p - g|$   $\triangleright$  Set operation
9:   return  $w \times (1 - \frac{false\text{-}alarm}{|p|})$ 
10: end function

```
