

# Document-level Entity-based Extraction as Template Generation

Kung-Hsiang Huang<sup>1,2</sup> Sam Tang<sup>3</sup> Nanyun Peng<sup>1,3</sup>

<sup>1</sup> Information Sciences Institute, University of Southern California

<sup>2</sup> Department of Computer Science, University of Illinois Urbana-Champaign

<sup>3</sup> Computer Science Department, University of California, Los Angeles

khhuang3@illinois.edu samtang1430@gmail.com

violetpeng@cs.ucla.edu

## Abstract

Document-level entity-based extraction (EE), aiming at extracting entity-centric information such as entity roles and entity relations, is key to automatic knowledge acquisition from text corpora for various domains. Most document-level EE systems build extractive models, which struggle to model long-term dependencies among entities at the document level. To address this issue, we propose a generative framework for two document-level EE tasks: role-filler entity extraction (REE) and relation extraction (RE). We first formulate them as a template generation problem, allowing models to efficiently capture cross-entity dependencies, exploit label semantics, and avoid the exponential computation complexity of identifying N-ary relations. A novel cross-attention guided copy mechanism, TOPK COPY, is incorporated into a pre-trained sequence-to-sequence model to enhance the capabilities of identifying key information in the input document. Experiments done on the MUC-4 and SciREX dataset show new state-of-the-art results on REE (+3.26%), binary RE (+4.8%), and 4-ary RE (+2.7%) in F1 score <sup>1</sup>.

## 1 Introduction

Document-level entity-based extraction (EE) are tasks that extract entity-centric information, such as entities and their relations, from unstructured text across multiple sentences. With the rise of big data in recent years, document-level EE is growing in importance with applications such as understanding clinical reports (Nye et al., 2020), extracting document-level events (Huang and Peng, 2021), and building knowledge graphs from journals (Wu et al., 2020). In this work, we focus on two classic tasks of document-level EE: role-filler entity extraction (REE) and relation extraction (RE).

Recent works on document-level EE usually build

<sup>1</sup>The source code is publicly available at <https://github.com/PlusLabNLP/TempGen>

We conducted experiments on the Stanford Natural Language Inference (SNLI) dataset ... (39 sentences) ... Our best score on this task is 87.3 % accuracy. ... (20 sentences) ... we adopt MAP and MRR as the evaluation metrics for this task.

	Relations
Gold	SNLI-accuracy
Ours	SNLI-accuracy
SciREX-P	SNLI-MAP
	SNLI-MRR

Figure 1: A comparison between our approach and a competitive extractive system, SciREX-P (Jain et al., 2020), on a relation extraction example from SciREX. The task is to extract entities and identify which entities are related from the given scientific article. Due to the long distances between entities, SciREX-P struggles to extract the right entity pair that has a relation, while our approach correctly identifies them. This reflects our method’s advantage in modeling long-term cross-entity dependencies.

task-specific classifiers on top of large pre-trained language models. For example, Du and Cardie (2020a) builds a sequence tagging framework with multi-granularity representations based on BERT (Devlin et al., 2019) for role-filler entity extraction. Jain et al. (2020) builds a relation extraction pipeline upon SciBERT (Beltagy et al., 2019). However, there are a few drawbacks of this model architecture. First, as the size of the document increases, it becomes increasingly difficult for extractive methods to capture cross-entity dependencies among entity types due to long distances between entities, as shown in Figure 1. Additionally, discriminative models have no information regarding the semantics of the labels when classifying relations or entity types. Thus, it is unable to take advantage of the label semantics embedded in the pre-trained encoders.

Motivated by these challenges, we propose to formulate REE and RE tasks as template generation. Due to the autoregressive nature of generative setup, this formulation makes dependencies among

the output entities easier to capture compared to sequence tagging methods. Moreover, label names are incorporated into the decoder targets for exploiting label semantics not present in the extractive counterparts. Furthermore, for tasks that involve the identification of  $N$ -ary relations, this formulation significantly alleviates the computational complexity of comparing exponential combinations of entities. A generative framework, Cross-attention Guided Template Generation (TEMPGEN), that incorporates a novel copy mechanism into a pre-trained sequence-to-sequence model is proposed to solve the template generation problem effectively.

Our contributions can be summarized as follows:

- We propose to formulate document-level EE tasks as a template generation problem, which allows our generative framework to effectively capture cross-entity dependencies, better identify entities with label semantics, and avoid the exponential computation complexity of identifying  $N$ -ary relations.
- We devise a novel copy mechanism based on cross-attention to enable our model to better learn how to copy key information from the input document.
- Our approach achieves state-of-the-art results on MUC role-filler entity extraction task and SciREX relation extraction task, while being data efficient compared to previous systems.

## 2 Tasks

This section gives an overview of the two document-level EE tasks we tackled in this work: role-filler entity extraction (REE) and relation extraction (RE).

### 2.1 Role-filler Entity Extraction

The REE task aims to extract all entities involved in events from the input article (Du et al., 2021). It differs from the event template extraction task introduced by the MUC-4 dataset (muc, 1992) in that only one event template, as opposed to many, is outputted for each input document. For documents associating with multiple event templates (all events in MUC-4 are of ATTACK type), the event templates are collapsed as one — systems are required to identify all entities associate with different events for each role type. An event template consists of a set of pre-defined roles, and each

role is filled with zero to many entities, as shown in Figure 2. An entity is characterized by a group of mentions, which are spans of text in the input document.

### 2.2 Relation Extraction

We focus on end-to-end document-level relation extraction where systems first extract entities from the input document and then identify the  $N$ -ary non-typed relations among the extracted entities. The SciREX (Jain et al., 2020) dataset is the only dataset that supports such end-to-end configurations that we know of. Thus, we follow the definition of document-level RE in SciREX, which contains binary and 4-ary relation annotation. A binary relation contains two typed entities, and a 4-ary relation contains four typed entities. An entity is represented by a cluster of mentions, similar to the REE task. Systems should first extract salient entities of pre-defined types<sup>2</sup>. Then, binary and 4-ary relations among salient entities are identified. A binary relation example is shown in Figure 2.

## 3 Proposed Methods

In this section, we first illustrate how the REE and RE tasks can be framed as a template generation problem. This formulation then allows us to capture cross-entity dependencies easily with our proposed generative model, a pre-trained sequence-to-sequence model integrated with a copy mechanism.

### 3.1 Template Generation Formulation

We frame the REE and RE tasks as template generation problem, as shown in Figure 2. A template is composed of slot names and slot values. For both tasks, slot names are entity types, and slot values are all entity mentions corresponding to such entity types. Similar to previous works on REE (Huang and Riloff, 2011; Du and Cardie, 2020a; Du et al., 2021), we only generate one template per document without differentiating which event template each entity mention associates with. In contrast, for RE, we generate multiple templates, each corresponding to a relation. A binary relation can be represented by a template of 2 slots, whereas a 4-ary relation forms a 4-slot template. A relation template consists of typed mentions of corresponding salient entities. After transforming REE and RE annotation to templates, each template can then

<sup>2</sup>Salient entities are entities needed to describe the results of corresponding scientific article.

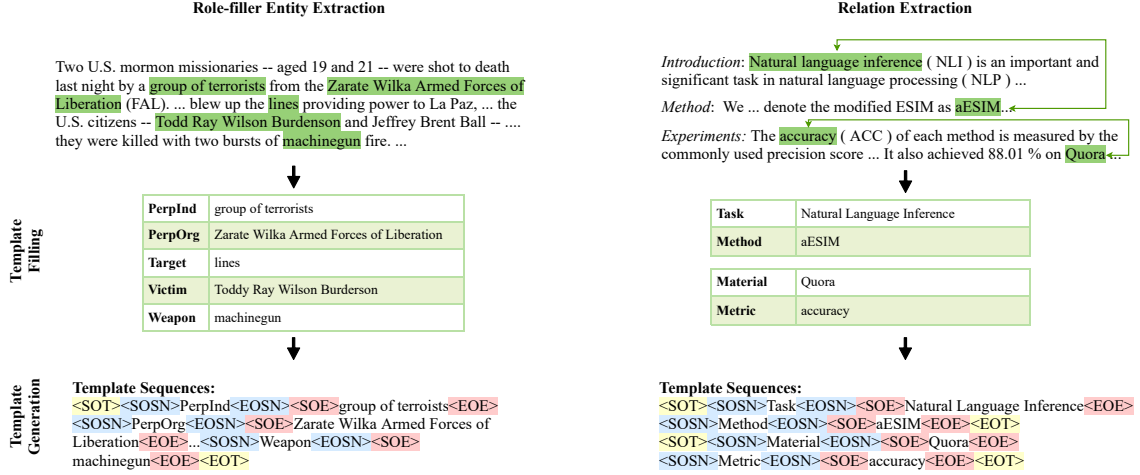


Figure 2: An overview of how document-level EE tasks can be transformed into template generation tasks. Special tags are defined as follows: <SOT>: start of template, <EOT>: end of template, <SOSN>: start of slot name, <EOSN>: end of slot name, <SOE>: start of entity (slot value), <EOE>: end of entity.

be transformed into *template sequences* with special tags delimiting templates, slot names, and slot values.

Formally, a document of tokens  $\mathcal{D} = \{\mathcal{D}_i\}_{i=1}^n$  may correspond to a decoding target of zero to many *template sequences*  $\{T_i\}_{i=1}^l$ . A *template sequence*  $T_i$  is characterized by multiple *slot sequences*  $\{S_{i,j}\}_{j=1}^m$ ,

$$T_i = \langle \text{SOT} \rangle S_{i,1}, \dots, S_{i,m} \langle \text{EOT} \rangle.$$

A *slot sequence*  $S_{i,j}$  is represented by slot names and entities,

$$S_{i,j} = \langle \text{SOSN} \rangle L \langle \text{EOSN} \rangle \langle \text{SOE} \rangle \mathcal{D}_1^{(e_k)} \dots \mathcal{D}_n^{(e_k)} \langle \text{EOE} \rangle.$$

where  $L$  is the slot name<sup>3</sup>, and  $\mathcal{D}_1^{(e_k)}, \dots, \mathcal{D}_n^{(e_k)}$  is the token sequence that correspond to one mention randomly sampled from entity  $e_k$ . Special tokens, such as <SOSN> and <EOSN>, are to indicate whether a tag-enclosed string is a slot name or an entity mention. In the first row of the REE example from Figure 2,  $L$  would be PERPIND and  $\mathcal{D}_1^{(e_k)}, \dots, \mathcal{D}_n^{(e_k)}$  are “group of terrorists”. Using this formulation, scalability challenges of modeling cross-entity dependencies is alleviated due to the significantly reduced distances between entities in *template sequences*.

### 3.2 Cross-attention Guided Template Generation

The template generation problem can be broken down into two sub-goals: (1) generating valid tem-

plate structures while capturing the dependencies between the input document and decoder targets, and (2) ensuring that salient mentions in the input document are correctly identified and outputted by the decoder. To achieve the first sub-goal, we leverage BART (Lewis et al., 2020), a pre-trained sequence-to-sequence model. The second sub-goal is achieved using a novel copy mechanism incorporated into BART.

#### Seq2Seq Model for Template Generation

BART (Lewis et al., 2020) is a pre-trained language model that combines bidirectional and auto-regressive transformers. Pre-training with multiple denoising objectives, BART has demonstrated significant advantages in various text generation tasks, especially on summarization (Lewis et al., 2020)<sup>4</sup>. The template generation problem much resembles summarization, except that generated *template sequences* contain implicit structures. With the various denoising pre-training objectives, we believe that BART can capture the implicit structure within *template sequences*, effectively model the dependencies among predicted entities, and produce rich semantics to reason over between slot names and entities.

**Cross-attention guided copy mechanism** To enhance BART’s capabilities to identify salient mentions in the input documents, we incorporate a copy mechanism based on cross-attention. As cross-attentions often imply saliency of input to-

<sup>3</sup>Slot name corresponds to role in REE and entity type in RE.

<sup>4</sup>We have considered the SOTA abstractive summarization LM, PEGASUS (Zhang et al., 2019). Yet, the GPU memory consumption is too high for us to test it.

kens, a naive approach of computing copy distributions  $P_{\text{copy}}$  at time step  $t$  over the input tokens is taking the mean of the last decoder layer’s cross-attention across all heads, as mentioned in Xu et al. (2020),

$$\alpha_{t,h} = \text{softmax}\left(\frac{(W_s s_t)^T W_e e}{\sqrt{d_k}}\right) \quad (1)$$

$$P_{\text{copy}} = \frac{\sum_h \alpha_{t,h}}{|H|}, \quad (2)$$

where  $\alpha_{t,h}$  is the attention scores over input tokens at decoding step  $t$  for head  $h$ .  $W_s$  and  $W_e$  are the projection matrices for the encoder and the decoder.  $s_t$  is the decoder hidden states at step  $t$ , and  $e$  denotes the encoder hidden states.

However, recent studies have shown that attention heads are not equally important, and that some heads can be pruned out with a marginal decrease in overall performance (Voita et al., 2019; Michel et al., 2019). We hypothesize that the attention probabilities produced by insignificant attention heads may be noisy. Thus, computing copy distributions without these heads could improve the model’s ability to infer the importance of each token in the input document. Motivated by this hypothesis, we propose TOPK COPY, a copy mechanism where only the Top- $k$  important attention heads are used for computing copy distributions. Consider the formulation of multi-head attention, following the notation from Vaswani et al. (2017):

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \dots, \text{head}_h) W^O \quad (3)$$

$$\text{head}_i = \text{Attention}(Q W_i^Q, K W_i^K, V W_i^V). \quad (4)$$

$W_i^Q, W_i^K, W_i^V \in \mathbb{R}^{d_{\text{model}} \times d}$  are the projection matrices for computing attention.  $W^O \in \mathbb{R}^{hd_v \times d_{\text{model}}}$  is the matrix that allows interaction between different attention heads, where  $h$  is the number of heads. To determine the importance of each attention head, we first transform  $W^O$  to dimension  $h \times d_v \times d_{\text{model}}$  (Equation (5)), and then sum over the last two dimensions of  $W^O$  (Equation (6)),

$$W^O \in \mathbb{R}^{hd_v \times d_{\text{model}}} \rightarrow W^O \in \mathbb{R}^{h \times d_v \times d_{\text{model}}} \quad (5)$$

$$\text{score}_i = \sum_{j,k} |W_{i,j,k}^O|. \quad (6)$$

where  $\text{score}_i$  denotes the significance score for head  $i$ . We take the attention heads with Top- $k$  highest significance scores in the last cross-attention layer, and use the mean of the attention

probabilities outputted by these heads as the copy distribution as shown in equations 7 and 8,

$$K = \text{Top-}k(\text{score}) \quad (7)$$

$$P_{\text{copy}} = \frac{\sum_{h \in K} \alpha_{t,h}}{k}. \quad (8)$$

**Objective function.** The final probability  $P_{\text{final}}$  of a word  $w_t$  is a weighted sum of vocabulary distribution computed by BART  $P_{\text{vocab}}$  and copy distribution  $P_{\text{copy}}$ ,

$$P_{\text{final}}(w_t) = p_{\text{gen}} P_{\text{vocab}}(w_t) + (1 - p_{\text{gen}}) P_{\text{copy}}(w_t). \quad (9)$$

where  $p_{\text{gen}} \in [0, 1]$  is the generation probability computed by passing the dot product of the mean encoder hidden state  $\bar{e} = \frac{\sum_{i=0}^n e_i}{n}$  and decoder hidden state  $s_t$  at time step  $t$  through the sigmoid function  $\sigma$ ,

$$p_{\text{gen}} = \sigma(\bar{e} \cdot s_t) \quad (10)$$

Using the final probability distribution  $P_{\text{final}}$ , we can then compute the loss function as the average negative log likelihood of the target word  $y_t$  over all timesteps, following See et al. (2017),

$$\mathcal{L} = \frac{1}{T} \sum_{t=0}^T -\log P_{\text{final}}(y_t). \quad (11)$$

## 4 Experimental Setup

### 4.1 Dataset and Evaluation Metric

Experiments are conducted on two English datasets: MUC-4 (1992) for the role-filler entity extraction task and SciREX (Jain et al., 2020) for the binary and 4-ary end-to-end relation extraction tasks. MUC-4 contains 1700 documents, with on average about 400 tokens per document. Documents are annotated with zero to multiple event templates. As per Du et al. (2021)’s pre-processing, we have a 13:2:2 split on the documents for train, development, and test, respectively. We evaluate the REE task on this dataset using the entity-level metric, CEAF-REE (Du et al., 2021). The metric aligns predicted entities with gold entities using Kuhn–Munkres algorithm (Kuhn, 1955; Munkres, 1957), where a predicted entity is considered correct if and only if its corresponding mentions are a subset of the aligned gold entity’s mentions.



Model	REE			Binary RE			4-ary RE		
	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1
NST (Du and Cardie, 2020a)	56.82	48.92	52.58	-	-	-	-	-	-
TANL (Paolini et al., 2021)	64.89	47.75	55.02	0.74	0.67	0.62	0.00	0.00	0.00
GRIT (Du et al., 2021)	64.19	47.36	54.50	-	-	-	-	-	-
DYGIE++ (Wadden et al., 2019)	57.04	46.77	51.40	2.9	12.8	3.8	-	-	-
SciREX-P (Jain et al., 2020)	-	-	-	6.5	<b>41.1</b>	9.6	0.7	<b>17.3</b>	0.8
TEMPGEN	<b>68.55</b>	<b>49.90</b>	<b>57.76</b>	<b>17.11</b>	13.56	<b>14.47*</b>	<b>3.19</b>	4.26	<b>3.55*</b>

Table 1: Performance comparison on role-filler entity extraction, binary and 4-ary relation extraction tasks. TANL results are re-implemented and evaluated by ourselves. TEMPGEN outperforms all previous systems on REE, binary RE, and 4-ary RE. Statistical significance over previous best systems computed using the paired bootstrap procedure (Berg-Kirkpatrick et al., 2012) is indicated with \* ( $p < .01$ ).

The SciREX dataset<sup>5</sup> consists of scientific articles, with entity, coreference, and relation annotations. With an average token count of about 5700, the articles are significantly longer than the documents in MUC-4. We use the pre-processed data from Jain et al. (2020), which contains 306 documents for training, 66 for validation, and 66 for testing. In contrast to conventional relation extraction datasets, such as ACE05, relations are not typed in SciREX. Hence, the official SciREX evaluator (Jain et al., 2020) only considers the correctness of predicted entities and entity types<sup>6</sup> in each relation. Predicted entities are aligned with gold entities based on mention overlap. When the entities are aligned, predicted relations are aligned with gold relations accordingly. A predicted relation is correct if and only if both the associated entities and the entity types match the aligned gold relation.

## 4.2 Baselines

We compare our method with the following competitive baseline systems.

**NST** (Du and Cardie, 2020a) builds multi-granularity representations on documents, and utilizes gate mechanism to fuse representations of different granularity.

**TANL** (Paolini et al., 2021) augments sequential labels with input sentences, allowing it to be applied to various structured prediction tasks<sup>7</sup>.

**GRIT** (Du et al., 2021) shares transformer parameters between the encoder and the pointer network decoder, and is the SOTA system for the REE task

on the MUC dataset.

**DYGIE++** (Wadden et al., 2019) is a span-based multi-task IE framework jointly trained on relation extraction, named entity recognition, and coreference resolution.

**SciREX-P** (Jain et al., 2020) is the SOTA framework for end-to-end binary and 4-ary relation extraction on SciREX. The pipeline is composed of 4 components: mention identification, mention clustering, salient entity cluster identification, and relation classification.

In terms of the pre-trained language models used, BERT-BASE (Devlin et al., 2019) is used for **NST**, **DYGIE++**, and **GRIT**. **SciREX-P** is fine-tuned on SciBERT (Beltagy et al., 2019). We replace T5 (Raffel et al., 2020) with BART-BASE for **TANL** for a fair comparison with our method.

## 4.3 Implementation details

The proposed models are optimized using AdamW (Loshchilov and Hutter, 2019) with learning rate  $5e-5$  and weight decay  $1e-5$ . We used grid search to find the best  $k$  for TOPK COPY and found that  $k = 10$  yields the best overall performance across REE and RE. The maximum input sequence length for RE and REE are 1024 and 512, respectively. During inference time, all generative models used beam search with a beam width of 4.

# 5 Results and Analysis

## 5.1 Main Results

Table 1 summarizes the main results on role-filler entity extraction, binary, and 4-ary relation extraction. TEMPGEN establishes new state-of-the-art scores on all three tasks, outperforming the previous best models by an absolute F1 of 3.26%, 4.8%, and 2.7%. The improvements demonstrate the effectiveness of our approach in formulating

<sup>5</sup><https://github.com/allenai/SciREX>

<sup>6</sup>There are 4 entity types: MATERIAL, METRIC, TASK, and METHOD.

<sup>7</sup>Since the source code of TANL has not been released by the time we conducted experiments, we re-implemented it by closely following the method described in Paolini et al. (2021).

Model	REE	Binary RE	4-ary RE
TEMPGEN	<b>57.76</b>	<b>14.47</b>	<b>3.55</b>
→ NAIVE COPY	56.45	11.22	1.42
→ SAGCopy	54.47	11.17	1.41
w/o TOPK COPY	55.76	12.63	3.00
numeric slot name	56.31	8.22	0.85

Table 2: Ablation study on removing and replacing different components of TEMPGEN.

document-level EE tasks into template generation tasks. Although TEMPGEN scores the highest F1 across all three tasks, SCIREX-P does achieve the highest recall on both RE tasks. This can be explained by the fact that our model can only encode the first 1024 sub-tokens of each SCIREX document, which is merely 17% of the average sub-token count per document. This makes it challenging for TEMPGEN to identify relations that lie in the latter 83% of each document. In the future, we can extend BART’s positional embedding matrix to enable TEMPGEN to encode longer documents. Additionally, we set the maximum input sequence length to 512 for TEMPGEN for fairer comparisons with SCIREX-P. We obtain F1 scores of 11.94% and 2.18% on binary and 4-ary relation extraction, respectively. This confirms the advantage of our model on the relation extraction tasks.

While TANL performs worse than our model on REE, it is still able to achieve a higher score than GRIT. This suggests that augmenting decoding targets with label names provides useful semantics, whereas adding input documents to decoding targets may not yield better results in the REE task. We also observe that TANL scores extremely low on both RE tasks, where 58% of the binary relations and 26 % of the 4-ary relations in the decoding targets are filtered out due to exceeding maximum sequence length of BART. Out of the remaining relations, 57% of the binary relations and 78% of the 4-ary relations have at least one entity removed in the decoding targets due to its long distance from the first-appearing entity<sup>8</sup>, suggesting that TANL’s poor performance on RE tasks is due to scarcity of gold labels. This reflects that TANL is ill-suited for document-level EE tasks.

We observe extremely low performances across all systems on both tasks of SCIREX, even though TEMPGEN outperforms the baseline systems significantly. This is mainly caused by the characteristics of the SCIREX dataset. First, syntactic

<sup>8</sup>Please refer to Appendix B for more details.

.....The police also stopped 8,000 cars in the search for assassins, who are presumably **members** of the maoist terrorist organization shining path.....  
The dircote (counterterrorism division) has identified one of the **terrorists** as Gerardo Olivos Silva through a composite made from witness' reports.....

Generated PerpInd entity: terrorists

(a) Copy distribution produced by NAIVE COPY.

.....The police also stopped 8,000 cars in the search for assassins, who are presumably **members** of the maoist terrorist organization shining path.....  
The dircote (counterterrorism division) has identified one of the **terrorists** as Gerardo Olivos Silva through a composite made from witness' reports.....

Generated PerpInd entity: Gerardo Olivos Silva

(b) Copy distribution produced by TOPK COPY.



(c) The darker the color, the higher the probability.

Figure 3: TOPK COPY produces a more reliable copy distribution  $P_{copy}$  than that computed by NAIVE COPY in a MUC-4 example. Given an input document and decoded tokens “<s><SOT><SOSN>PERPIND<EOSN><SOE>”, the gold PERPIND entity is “Gerardo Olivos Silva”. However, “terrorist” is assigned the highest copy probability computed by NAIVE COPY, leading to incorrect entity extracted. Conversely, TOPK COPY assigns the highest  $P_{copy}$  to the head token of the gold entity, “Ger”, resulting in successful extraction of the correct entity eventually.

characteristics specific to scientific journals, such as algorithm blocks, result in the unusually long sequences in the SCIREX dataset despite best parsing efforts. Additionally, another feature frequently seen in scientific journals is the use of table and figure captions. Since captions are not included as part of the input text, the number of accepted relations decreases drastically.

## 5.2 Performance Analysis

**Ablation Study** We conducted ablation studies by replacing the TOPK COPY module with other copy mechanisms. **NAIVE COPY** refers to computing copy distributions with the attentions from all cross-attention heads. **SAGCopy** (Xu et al., 2020) utilizes encoder self-attention to compute centrality scores for measuring the saliency of each input token. As shown in Table 2, we found that NAIVE COPY leads to performance drop on all three tasks, especially on binary and 4-ary relation extractions. **NAIVE COPY** achieving scores even lower than fine-tuning BART alone (i.e. w/o TOPK COPY) reflects that copy mechanisms may mislead models to copy incorrect input tokens. A qualitative example of the difference between TOPK COPY and NAIVE COPY demonstrated in Figure 3 validates our hypothesis. Quantitatively, examining MUC-4




Figure 4: Effect of  $K$ . We compare the test set F1 score on MUC-4 REE, SciREX binary and 4-ary RE tasks with regard to different  $K$ .  $K = 0$  is equivalent to removing TOPK COPY.

test set predictions, there are 79 cases where TOPK COPY corrects the misguidance of Naive COPY, while only 32 cases where new errors are introduced by TOPK COPY. For both REE and RE, adding SAGCopy leads to performance drop, suggesting that the centrality scores of input tokens may not be an ideal feature for these tasks.

We also experimented with replacing the original slot names with numeric slot names (i.e. converting PERPIND to <ROLE\_1>, PERPORG to <ROLE\_2>, and etc). This conversion removes the semantics of slot names in the decoding targets. While little performance drop was observed on the REE task, using numeric slot names resulted in the worst performance on binary and 4-ary relation extraction tasks, which could be a result of strong slot dependencies in RE in comparison with REE. In RE, slots are directly semantically related to other slots in each template whereas slots in REE are relatively independent. This shows that slot name semantics are useful for template generation tasks with strong slot dependencies in each template. Finally, we conducted ablation studies on different variations of templates as decoding targets. Specifically, three variations are tested on the REE task: (1) We merge entities of the same role names into the same “slot”. (e.g. transforming the decoding targets from “<SOSN>PerpInd<EOSN> <SOE>Alice<EOE> <SOSN>PerpInd<EOSN> <SOE>Bob<EOE>” to “<SOSN>PerpInd<EOSN> <SOE>Alice; Bob<EOE>”). (2) Based on (1), all slot names, such as “PerpInd” and “PerpOrg”, are replaced with the same special token “<ROLE>”. (3) We use the same decoding targets as GRIT’s. These three settings achieve test set F1 scores of 56.65, 54.16, and 52.55, respectively. The results suggest that differentiating entities with different entity types helps improve the performance. Furthermore, comparing with the results in Table 1, we found that GRIT performs better than our system, reflecting that a pointer network-based model, which has with smaller search space than ours, is



Figure 5: REE test set performance on MUC-4 with regard to different amount of training data.

more advantageous when using the same decoding targets.

**Impact of the Amount of Training Data** To test the data efficiency of our approach, we compared TEMPGEN and TEMPGEN - TOPK COPY with GRIT on the REE task using different amount of MUC training data. As seen in Figure 5, both TEMPGEN and TEMPGEN - TOPK COPY outperform GRIT across all settings with a slightly larger performance margin in low resource settings. This indicates that our approach is more data-efficient compared to the previous SOTA system on REE.

**Impact of  $K$  Cross-attention Heads** Figure 4 shows our model’s change in performance conditioned on various values of  $K$  in the TOPK COPY mechanism. Consistent with our results in Section 5.1, we see that removing some of the cross-attention gain heads (12  $\rightarrow$  10) can lead to performance gain due to the filtered noise brought by unimportant attention heads. However, we noticed a drop in performance across all three tasks for lower values of  $K$ , suggesting that beneficial cross-attention heads are removed. Interestingly, performance drops immediately as  $K$  decreases below 10, suggesting that only a small portion of the cross-attention heads are unimportant. The trend is consistent with Michel et al. (2019)’s results where pruning cross-attention heads to a certain extent can easily result in performance drop. Additionally, the model with no copy mechanism ( $K = 0$ ) outperforms the model with few attention heads

Miguel Soler Rodriguez, El Espectador's circulation manager and Martha Luz Lopez, the correspondent's office administrator, were killed today .... the main offices in Bogota of El Espectador, ..., were partially destroyed.

	Victim	Target
<b>Gold</b>	<ul style="list-style-type: none"> <li>Miguel Soler Rodriguez</li> <li>Martha Luz Lopez</li> </ul>	<ul style="list-style-type: none"> <li>the main offices in Bogota of El Espectador</li> </ul>
<b>Ours</b>	<ul style="list-style-type: none"> <li>Miguel Soler Rodriguez</li> <li>Martha Luz Lopez</li> </ul>	<ul style="list-style-type: none"> <li>El Espectador's office</li> </ul>
<b>GRIT</b>	<ul style="list-style-type: none"> <li>El Espectador</li> </ul>	<ul style="list-style-type: none"> <li>Miguel Soler Rodriguez</li> <li>Martha Luz Lopez</li> </ul>

Figure 6: An example showing how GRIT misidentifies the VICTIM entities and TARGET entities, likely due to the lack of role type semantics. Here, VICTIM entities are the people attacked, and TARGET entities are the objects compromised.

( $K \in \{2, 4, 6\}$ ), suggesting that the copy distributions obtained from not sufficiently informative cross-attentions can mislead the model.

### 5.3 Qualitative Analysis

The following qualitative analysis provides intuition for our model’s ability to capture dependencies across entities and utilize slot name semantics.

**Cross-entity Dependencies** To validate our approach’s capability to capture cross-entity dependencies, we considered binary relations on SCIREX where at least one of the associated entities is involved in multiple relations. The dependencies among entities are better captured by the model that predicts fewer unlikely relations. Comparing the test set outputs of TEMPGEN and SCIREX-P, we see that 13131 errors made by SCIREX-P are corrected by our model, which only introduces 604 errors. This result demonstrates the strength of TEMPGEN in modeling cross-entity dependencies.

**Importance of Label Semantics** Comparing the test set predictions between TEMPGEN and GRIT on the MUC-4 REE task, we see that our approach better distinguishes confusing entities such as VICTIM and TARGET entities. As shown in the example in Figure 6, GRIT incorrectly predicts the two victims, “Miguel Soler Rodriguez” and “Martha Luz Lopez”, as TARGET entities. It also misidentifies “El Espectador”, a newspaper company, as a victim of the attack. In contrast, TEMPGEN is able to recognize the roles of the two victims. Even though it’s not an exact match, the predicted TARGET entity had a correctly identified role type with similar semantic meaning compared to the gold label.




Figure 7: Inference time comparison on the SCIREX 4-ary RE task.

## 6 Inference time comparison

As discussed earlier, TEMPGEN can significantly reduce the exponential computational complexity of document-level N-ary relation identification. To illustrate this, we compared the inference time between TEMPGEN and two other systems, TANL and SCIREX-P, on the SCIREX 4-ary RE task. As shown in Figure 7, TEMPGEN drastically shortens the inference time by around 39 times compared to SCIREX-P. TANL also runs much faster than SCIREX-P, but is still around 4 times slower than TEMPGEN. This is resulted from the fact that TANL generates the entire input document in addition to entity and relation labels, which is much longer than TEMPGEN’s generated sequences.

## 7 Number of Parameters

Figure 8 shows the number of parameters of different models. GRIT, with the same size of BERT-BASE, has the least number of parameters among all models. DYGIE++ and SCIREX-P have slightly more parameters than GRIT due to the additional linear layers for constructing classifiers. The two generative models, TANL and TEMPGEN, have the most parameters, thanks to the larger vocab size ( $30522 \rightarrow 50265$ ), larger positional embedding matrix ( $512 \rightarrow 1024$ ), and cross-attention modules in BART-BASE.




Figure 8: Number of parameters of different systems.

## 8 Related Works

In the following sections, we will first discuss a few important works on the REE task and document-



level RE task. Then, we will dive into a few works that uses a similar sequence generative approach to various document-level IE tasks.

### 8.1 Role-filler Entity Extraction

Document-level REE has been explored in recent works using a variety of model architectures. [Du and Cardie \(2020b\)](#) formulates the task as a sequence tagging problem, and trains layered classifiers as sequence readers on multiple granularities. In contrast, GRIT ([Du et al., 2021](#)) formulates the problem as sequence generation, and employs a single transformer layer whose parameters are shared between encoder and decoder to enrich semantics in the shared parameters. A pointer selection network is used for the final layer of decoding.

### 8.2 Document-level Relation Extraction

Due to long-term dependencies that often span over hundreds of tokens, capturing entity relations have proven to be a challenging task. One approach was constructing a document-level graph from sentence encoding, then extracting entity relations from edge representations in the graph ([Christopoulou et al., 2019](#)). Other works such as [Jia et al. \(2019\)](#) layer classifiers in a pipeline architecture to obtain hierarchical representation of  $N$ -ary relations.

### 8.3 IE as Sequence Generation

Recently, there has been an increasing number of works framing information extraction tasks as sequence generation problem. [Zeng et al. \(2018\)](#) formed triple extraction as a sequence generation task and adopted a RNN-based model with copy mechanisms. To encourage the faithfulness of the extracted triplets, [Ye et al. \(2021\)](#) designed a triplet contrastive training objective. These works focus on sentence-level triplet extraction, while our work extracts role-filler entities and entity relations at the document level. [Li et al. \(2021\)](#); [Hsu et al. \(2021\)](#) formulates the document-level event argument extraction task as a conditional generation problem by providing event ontology. However, their work cannot be applied to REE or RE due to the lack of ontology for role-filler entities and relations. [Du et al. \(2021\)](#) relied on a pointer-network-based decoder ([Vinyals et al., 2015](#)) to extract event role-filler entities, and the parameters of BERT ([Devlin et al., 2019](#)) is shared between the encoder and the decoder. Nevertheless, their method cannot incorporate role labels, whereas our approach can take advantage of the label semantics.

[Paolini et al. \(2021\)](#) uses a very similar generative approach, which constructs decoder targets by inserting text markers and labels around entity mentions in the input sentence. The key idea is that augmenting the decoder targets with original input sentence and labels provides stronger semantics to the model. Unfortunately, modeling cross-entity dependencies remains a challenge as entities are further apart in their decoding targets. We instead transform annotations into *template sequences* as decoding targets, where distances between entities are significantly shortened. Thus, our approach alleviates the scalability challenge of capturing cross-entity dependencies at the scale of documents. Additionally, our approach differs in that the length of our decoder targets is significantly shorter, allowing the non-truncated decoder targets to fit in pre-trained language models. In contrast, for their method, the gold decoder targets are guaranteed to be longer than corresponding input document. Since the length of input tokens are often greater than the max sequence length of pre-trained language models for document-level EE, a great portion of the gold labels will be skipped using [Paolini et al. \(2021\)](#)'s method.

## 9 Conclusion

We have proposed TEMPGEN, a framework that frames document-level REE and RE tasks as a template generation task. A copy mechanism that takes the top- $k$  important cross-attentions as copy distributions is incorporated into BART for capturing key information in the input document. Experimental results on MUC-4 and SCIREX showed that TEMPGEN outperforms prior approaches on role-filler entity extraction and end-to-end document-level relation extraction tasks. Under different amount of training data, TEMPGEN demonstrates robustness across all settings, while being advantageous in lower-resource regime.

## 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

1992. *Fourth Message Understanding Conference (MUC-4): Proceedings of a Conference Held in McLean, Virginia, June 16-18, 1992*.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). 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 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Taylor Berg-Kirkpatrick, David Burkett, and Dan Klein. 2012. [An empirical investigation of statistical significance in NLP](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 995–1005, Jeju Island, Korea. Association for Computational Linguistics.
- Fenia Christopoulou, Makoto Miwa, and Sophia Ananiadou. 2019. [Connecting the dots: Document-level neural relation extraction with edge-oriented graphs](#). 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 4925–4936, Hong Kong, China. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. *Proceedings of the 2019 Conference of the North*.
- Xinya Du and Claire Cardie. 2020a. [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 and Claire Cardie. 2020b. Document-level event role filler extraction using multi-granularity contextualized encoding. *arXiv preprint arXiv:2005.06579*.
- Xinya Du, Alexander M. Rush, and Claire Cardie. 2021. Grit: Generative role-filler transformers for document-level event entity extraction. Online. EACL.
- I Hsu, Kuan-Hao Huang, Elizabeth Boschee, Scott Miller, Prem Natarajan, Kai-Wei Chang, Nanyun Peng, et al. 2021. Event extraction as natural language generation. *arXiv preprint arXiv:2108.12724*.
- Kung-Hsiang Huang and Nanyun Peng. 2021. Document-level event extraction with efficient end-to-end learning of cross-event dependencies. In *The 3rd Workshop on Narrative Understanding (NAACL 2021)*.
- Ruihong Huang and Ellen Riloff. 2011. [Peeling back the layers: Detecting event role fillers in secondary contexts](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1137–1147, Portland, Oregon, USA. 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.
- Harold W Kuhn. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Sha Li, Heng Ji, and Jiawei Han. 2021. [Document-level event argument extraction by conditional generation](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. [Are sixteen heads really better than one?](#) In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- James Munkres. 1957. Algorithms for the assignment and transportation problems. *Journal of the society for industrial and applied mathematics*, 5(1):32–38.
- Benjamin E. Nye, Jay DeYoung, Eric Lehman, Ani Nenkova, Iain J. Marshall, and Byron C. Wallace. 2020. [Understanding clinical trial reports: Extracting medical entities and their relations](#).
- Giovanni Paolini, Ben Athiwaratkun, Jason Krone, Jie Ma, Alessandro Achille, RISHITA ANUBHAI, Cicero Nogueira dos Santos, Bing Xiang, and Stefano Soatto. 2021. [Structured prediction as translation between augmented natural languages](#). In *International Conference on Learning Representations*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.

Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#).

Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2, NIPS’15*, page 2692–2700, Cambridge, MA, USA. MIT Press.

Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.

David Wadden, Ulme Wennberg, Yi Luan, and Hannaneh Hajishirzi. 2019. [Entity, relation, and event extraction with contextualized span representations](#). 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 5784–5789, Hong Kong, China. Association for Computational Linguistics.

Zequ Wu, Rik Koncel-Kedziorski, Mari Ostendorf, and Hannaneh Hajishirzi. 2020. [Extracting summary knowledge graphs from long documents](#).

Song Xu, Haoran Li, Peng Yuan, Youzheng Wu, Xiaodong He, and Bowen Zhou. 2020. [Self-attention guided copy mechanism for abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1355–1362, Online. Association for Computational Linguistics.

Hongbin Ye, Ningyu Zhang, Shumin Deng, Mosha Chen, Chuanqi Tan, Fei Huang, and Huajun Chen. 2021. [Contrastive triple extraction with generative transformer](#).

Xiangrong Zeng, Daojian Zeng, Shizhu He, Kang Liu, and Jun Zhao. 2018. [Extracting relational facts by an end-to-end neural model with copy mechanism](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 506–514, Melbourne, Australia. Association for Computational Linguistics.

Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2019. [Pegasus: Pre-training with extracted gap-sentences for abstractive summarization](#).

## A REE Performance Breakdown

Table 3 demonstrates the per-role performance comparison between TEMPGEN and other baselines. We observe that:

- TEMPGEN achieves the best precision across all roles.
- Except for PERPIND, TEMPGEN obtain substantial improvement in F1 over other baselines.
- While TEMPGEN has higher precision over GRIT in extracting PERPIND entities, it scores slightly lower in recall, leading to worse F1 performance.

## B TANL Decoding Target Formulation

In this section, we illustrate how we formulate the TANL (Paolini et al., 2021) decoding targets for REE and RE. The formulation for REE is simple due to its similarity to the NER task. We produce REE decoding targets exactly the same way as how NER decoding targets are formed in Paolini et al. (2021). Given the REE example in Figure 2, the corresponding TANL decoding target is:

Two U.S. mormon missionaries – aged 19 and 21 – were shot to death last night by [a group of terrorists| *PerpInd*] from the [Zarate Wilka Armed Forces of Liberation| *PerpOrg*] (FAL) ... blew up the [lines| *Target*] providing power to La Paz, ... the U.S. citizens – [Todd Ray Wilson Burdenson| *Victim*] and Jeffrey Brent Ball – ... they were killed with two bursts of [machinegun| *Weapon*] fire ...

As for RE, we follow how Paolini et al. (2021) handles *nested entities and multiple relations*, but we made a small modification on decoding targets. Since SCIREX does not contain relation type annotation, we use the related entities' types as the relation type in the decoding targets. With their formulation, the decoding target is created by inserting each relation annotation around the first-appearing entity in the input document. Take the RE instance in Figure 2 as an example. The corresponding TANL decoding target would be:

Introduction: [Natural language inference | *Task* | *Method* = aESIM] (NLI) is an important and significant task in natural language processing (NLP)...

Method: We ... denote the modified ESIM as aESIM ...

Experiments: The [accuracy | *Metric* | *Material* = Quora] (ACC) of each method is measured by the commonly used precision score ... It also achieved 88.01 % on Quora ...

## C Hardware and Software configurations

All experiments are conducted on a CentOS Linux 7 (Core) machine with NVIDIA RTX 2080. We use PyTorch 1.6.0 with CUDA 10.1 as the Deep Learning framework and utilize Transformers 4.3.0 to load all pre-trained language models.

## D Implementation Details

We conducted grid search to find the best learning rate over  $\{1 \times 10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}, 7 \times 10^{-5}, 9 \times 10^{-5}\}$  using **TEMPGEN w/o TOPK COPY** on the MUC-4 REE task. The best learning rate,  $5 \times 10^{-5}$ , is fixed for all other experiments. Models are trained for 150 epochs for REE and binary RE experiments, and 50 epochs for 4-ary RE experiments. To reproduce our results, please follow the README.md file in <https://github.com/PlusLabNLP/TempGen>. The weights of the trained models are also included for reproduction purposes.

## E Validation Performance

For all reported test set results in Table 1, the corresponding development set performance are listed in Table 4.



Model	PERPIND	PERPORG	TARGET	VICTIM	WEAPON
DYGIE++ (Wadden et al., 2019)	48.39/ 32.61/ 43.32	56.00/ 34.15/ 42.42	53.49/ 50.74/ 52.08	60.00/ 66.32/ 63.00	57.14/ 53.33/ 55.17
GRIT (Du et al., 2021)	65.48/ 39.86/ <b>49.55</b>	66.04/ 42.68/ 51.85	55.05/ 44.12 / 48.98	76.32/ 61.05/ 67.84	61.82/ 56.67 / 59.13
TEMPGEN	67.12/ 35.51/ 46.45	67.12/ 59.76/ <b>63.23</b>	64.13/ 43.38/ <b>51.75</b>	77.22/ 64.21/ <b>70.11</b>	67.27/ 61.67/ <b>64.35</b>

Table 3: Performance breakdown with regard to each role in CEAF-REE (Precision/ Recall /F1) on the MUC-4 REE task.

	Role-filler Entity Extraction			Binary Relation Extraction			4-ary Relation Extraction		
Model	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
TANL	58.42	<b>46.74</b>	51.93	3.12	2.11	2.39	0.00	0.00	0.00
TEMPGEN	<b>61.34</b>	46.11	<b>52.64</b>	<b>22.04</b>	<b>19.24</b>	<b>19.60</b>	<b>1.38</b>	<b>2.77</b>	<b>1.85</b>

Table 4: Corresponding development set performance of the reported test set results in Table 1.