

Multilingual Generative Language Models for Zero-Shot Cross-Lingual Event Argument Extraction

Kuan-Hao Huang^{*†} I-Hung Hsu^{*‡} Premkumar Natarajan[‡]
Kai-Wei Chang[†] Nanyun Peng^{‡†}

[†]Computer Science Department, University of California, Los Angeles

[‡]Information Science Institute, University of Southern California

{khhuang, kwchang, violetpeng}@cs.ucla.edu

{ihunghsu, pnataraj}@isi.edu

Abstract

We present a study on leveraging multilingual pre-trained *generative* language models for zero-shot cross-lingual event argument extraction (EAE). By formulating EAE as a *language generation* task, our method effectively encodes event structures and captures the dependencies between arguments. We design *language-agnostic templates* to represent the event argument structures, which are compatible with any language, hence facilitating the cross-lingual transfer. Our proposed model finetunes multilingual pre-trained generative language models to *generate* sentences that fill in the language-agnostic template with arguments extracted from the input passage. The model is trained on source languages and is then directly applied to target languages for event argument extraction. Experiments demonstrate that the proposed model outperforms the current state-of-the-art models on zero-shot cross-lingual EAE. Comprehensive studies and error analyses are presented to better understand the advantages and the current limitations of using generative language models for zero-shot cross-lingual transfer EAE.

1 Introduction

Event argument extraction (EAE) aims to recognize the entities serving as event arguments and identify their corresponding roles. As illustrated by the English example in Figure 1, given a trigger word “*destroyed*” for a *Conflict:Attack* event, an event argument extractor is expected to identify “*commando*”, “*Iraq*”, and “*post*” as the event arguments and predict their roles as “*Attacker*”, “*Place*”, and “*Target*”, respectively.

Zero-shot cross-lingual EAE has attracted considerable attention since it eliminates the requirement of labeled data for constructing EAE models in low-resource languages (Subburathinam et al., 2019; Ahmad et al., 2021; Nguyen and Nguyen,



Figure 1: An illustration of cross-lingual event argument extraction. Given sentences in arbitrary languages and their event triggers (*destroyed* and 起义), the model needs to identify arguments (*commando*, *Iraq* and *post* v.s. 军队, and 反对派) and their corresponding roles (*Attacker*, *Target*, and *Place*).

2021). In this setting, the model is trained on the examples in the *source* languages and directly tested on the instances in the *target* languages.

Recently, generation-based models¹ have shown strong performances on monolingual structured prediction tasks (Yan et al., 2021; Huang et al., 2021b; Paolini et al., 2021), including EAE (Li et al., 2021; Hsu et al., 2021). These works fine-tune pre-trained generative language models to generate outputs following designed templates such that the final predictions can be easily decoded from the outputs. Compared to the traditional classification-based models (Wang et al., 2019; Wadden et al., 2019; Lin et al., 2020), they better capture the structures and dependencies between entities, as the templates provide additional declarative information.

Despite the successes, the designs of templates in prior works are language-dependent, which makes it hard to be extended to the zero-shot cross-lingual transfer setting (Subburathinam et al., 2019; Ahmad et al., 2021). Naively applying such models trained on the source languages to the target languages usually generates *code-switching* outputs, yielding poor performance for zero-shot

¹We use pre-trained *generative* language models to refer to pre-trained models with encoder-decoder structure, such as BART (Lewis et al., 2020), T5 (Raffel et al., 2020), and mBART (Liu et al., 2020). For models adapting these pre-trained generative models to generate texts for downstream applications, we denote them as *generation-based* models.

*The authors contribute equally.

cross-lingual transfer,² as we will empirically show in Section 5.4. How to design *language-agnostic* generation-based models for zero-shot cross-lingual structured prediction problems is still an open question.

In this work, we present a study that leverage multilingual pre-trained generative models for zero-shot cross-lingual event argument extraction and propose X-GEAR (**Cross-lingual Generative Event Argument extractoR**). Given an input passage and a carefully designed prompt that contains an event trigger and the corresponding language-agnostic template, X-GEAR is trained to generate a sentence that fills in a language-agnostic template with arguments. X-GEAR inherits the strength of generation-based models that captures event structures and the dependencies between entities better than classification-based models. Moreover, the pre-trained decoder inherently identifies named entities as candidates for event arguments and does not need an additional named entity recognition module. The *language-agnostic templates* prevents the model from overfitting to the source language’s vocabulary and facilitates cross-lingual transfer.

We conduct experiments on two multilingual EAE datasets: ACE-2005 (Doddington et al., 2004) and ERE (Song et al., 2015). The results demonstrate that X-GEAR outperforms the state-of-the-art zero-shot cross-lingual EAE models. We further perform ablation studies to justify our design and present comprehensive error analyses to understand the limitations of using multilingual generation-based models for zero-shot cross-lingual transfer. Our code is available at <https://github.com/PlusLabNLP/X-Gear>

2 Related Work

Zero-shot cross-lingual structured prediction.

Zero-shot cross-lingual learning is an emerging research topic as it eliminates the requirement of labeled data for training models in low-resource languages (Ruder et al., 2021; Huang et al., 2021a). Various structured prediction tasks have been studied, including named entity recognition (Pan et al., 2017; Huang et al., 2019; Hu et al., 2020), dependency parsing (Ahmad et al., 2019b,a; Meng

et al., 2019), relation extraction (Zou et al., 2018; Ni and Florian, 2019), and event argument extraction (Subburathinam et al., 2019; Nguyen and Nguyen, 2021; Fincke et al., 2021). Most of them are *classification-based models* that build classifiers on top of a multilingual pre-trained *masked* language models. To further deal with the discrepancy between languages, some of them require additional information, such as bilingual dictionaries (Liu et al., 2019; Ni and Florian, 2019), translation pairs (Zou et al., 2018), and dependency parse trees (Subburathinam et al., 2019; Ahmad et al., 2021; Nguyen and Nguyen, 2021). However, as pointed out by previous literature (Li et al., 2021; Hsu et al., 2021), classification-based models are less powerful to model dependencies between entities compared to *generation-based models*.

Generation-based structured prediction. Several works have demonstrated the great success of generation-based models on monolingual structured prediction tasks, including named entity recognition (Yan et al., 2021), relation extraction (Huang et al., 2021b; Paolini et al., 2021), and event extraction (Du et al., 2021; Li et al., 2021; Hsu et al., 2021; Lu et al., 2021). Yet, as mentioned in Section 1, their designed generating targets are language-dependent. Accordingly, directly applying their methods to the zero-shot cross-lingual setting would result in less-preferred performance.

Prompting methods. There are growing interests recently to incorporate prompts on pre-trained language models in order to guide the models’ behavior or elicit knowledge (Peng et al., 2019; Sheng et al., 2020; Shin et al., 2020; Schick and Schütze, 2021; Qin and Eisner, 2021; Scao and Rush, 2021). Following the taxonomy in (Liu et al., 2021), these methods can be classified depending on whether the language models’ parameters are tuned and on whether trainable prompts are introduced. Our method belongs to the category that fixes the prompts and tunes the language models’ parameters. Despite the flourish of the research in prompting methods, there is only limited attention being put on multilingual tasks (Winata et al., 2021).

3 Zero-Shot Cross-Lingual Event Argument Extraction

We focus on zero-shot cross-lingual EAE. Given an input passage and an event trigger, an EAE

²For example, TANL (Paolini et al., 2021) is trained to generate “[Two soldiers|target] were attacked” to represent *Two soldiers* being a *target* argument. When directly applying it to Chinese, the ground truth for TANL becomes “[两位士兵|target]被攻击”, which is a sentence alternating between Chinese and English.

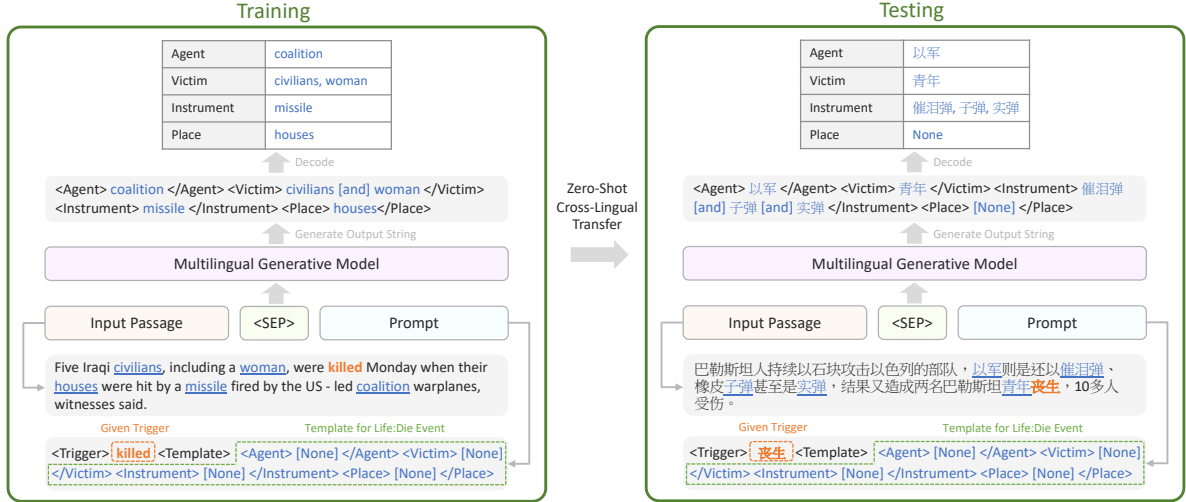


Figure 2: The overview of X-GEAR. Given an input passage and a carefully designed prompt containing an event trigger and a language-agnostic template, X-GEAR fills in the language-agnostic template with event arguments.

model identifies arguments and their corresponding roles. More specifically, as illustrated by the training examples in Figure 2, given an input passage x and an event trigger t (*killed*) belonging to an event type e (*Life:Die*), an EAE model predicts a list of arguments $\mathbf{a} = [a_1, a_2, \dots, a_l]$ (*coalition, civilians, woman, missile, houses*) and their corresponding roles $\mathbf{r} = [r_1, r_2, \dots, r_l]$ (*Agent, Victim, Victim, Instrument, Place*). In the zero-shot cross-lingual setting, the training set $X_{train} = \{(\mathbf{x}_i, \mathbf{t}_i, \mathbf{e}_i, \mathbf{a}_i, \mathbf{r}_i)\}_{i=1}^N$ belongs to the source languages while the testing set $X_{test} = \{(\mathbf{x}_i, \mathbf{t}_i, \mathbf{e}_i, \mathbf{a}_i, \mathbf{r}_i)\}_{i=1}^M$ are in the target languages.

Similar to monolingual EAE, zero-shot cross-lingual EAE models are expected to capture the dependencies between arguments and make structured predictions. However, unlike monolingual EAE, zero-shot cross-lingual EAE models need to handle the differences (e.g., grammar, word order) between languages and learn to transfer the knowledge from the source languages to the target languages.

4 Proposed Method: X-GEAR

We formulate zero-shot cross-lingual EAE as a language generation task and propose X-GEAR, a **C**ross-lingual **G**enerative **E**vent **A**rgument **E**xtracto**R** that is illustrated in Figure 2. There are two challenges raised by this formulation: (1) The input language may vary during training and testing; (2) The generated output strings need to be easily parsed into final predictions. Therefore, the output strings have to reflect the change of the input language accordingly while remaining well-

structured.

We address these challenges by designing *language-agnostic templates*. Specifically, given an input passage x and a designed prompt that contains the given trigger t , its event type e , and a *language-agnostic template*, X-GEAR learns to generate an output string that fills in the language-agnostic template with information extracted from input passage. The language-agnostic template is designed in a structured way such that parsing the final argument predictions \mathbf{a} and role predictions \mathbf{r} from the generated output is trivial. Moreover, since the template is language-agnostic, it facilitates cross-lingual transfer.

X-GEAR fine-tunes multilingual pre-trained generative models, such as mBART-50 (Tang et al., 2020) or mT5 (Xue et al., 2021), and augments them with a copy mechanism to better adapt to input language changes. We present its details as follows, including the language-agnostic templates, the target output string, the input format, and the training details.

4.1 Language-Agnostic Template

We create one language-agnostic template T_e for each event type e , in which we list all possible associated roles³ and form a unique HTML-tag-style template for that event type e . For example, in Figure 2, the *Life:Die* event is associated with four roles: *Agent, Victim, Instrument, and Place*. Thus, the template for *Life:Die* events is designed as:

³The associated roles can be obtained by skimming training data or directly from the annotation guideline if provided.

```
<Agent> [None] </Agent><Victim> [None] </Victim>
<Instrument> [None] </Instrument><Place> [None] </Place>.
```

For ease of understanding, we use English words to present the template. However, these tokens ([None], <Agent>, </Agent>, <Victim>, etc.) are encoded as special tokens⁴ that the pre-trained models have never seen and thus their representations need to be learned from scratch. Since these special tokens are not associated with any language and are not pre-trained, they are considered as *language-agnostic*.

4.2 Target Output String

X-GEAR learns to generate target output strings that follow the form of language-agnostic templates. To compose the target output string for training, given an instance ($\mathbf{x}, \mathbf{t}, \mathbf{e}, \mathbf{a}, \mathbf{r}$), we first pick out the language-agnostic template T_e for the event type \mathbf{e} and then replace all “[None]” in T_e with the corresponding arguments in \mathbf{a} according to their roles \mathbf{r} . If there are multiple arguments for one role, we concatenate them with a special token “[and]”. For instance, the training example in Figure 2 has two arguments (*civilians* and *woman*) for the *Victim* role, and the corresponding part of the output string would be

```
<Victim> civilians [and] woman </Victim>.
```

If there are no corresponding arguments for one role, we keep “[None]” in T_e . By applying this rule, the full output string for the training example in Figure 2 becomes

```
<Agent> coalition </Agent><Victim> civilians[and]
woman </Victim><Instrument> missile </Instrument>
<Place> houses </Place>.
```

Since the output string is in the HTML-tag style, we can easily decode the argument and role predictions from the generated output string via a simple rule-based algorithm.

4.3 Input Format

As we mentioned previously, the key for the generative formulation for zero-shot cross-lingual EAE is to guide the model to generate output strings in the desired format. To facilitate this behavior, we feed the input passage \mathbf{x} as well as a *prompt* to X-GEAR, as shown by Figure 2. The *prompt* contains all

⁴In fact, the special tokens can be replaced by any other format, such as <-token1-> or </-token1->. Here, we use <Agent> and </Agent> to highlight that arguments between these two special tokens are corresponding to the *Agent* role.

valuable information for the model to make predictions, including a trigger \mathbf{t} and a language-agnostic template T_e . Notice that we do not *explicitly* include the event type \mathbf{e} in the prompt because the template T_e *implicitly* contains this information. In Section 6.1, we will show the experiments on explicitly adding event type \mathbf{e} to the prompt and discuss its influence on the cross-lingual transfer.

4.4 Training

To enable X-GEAR to generate sentences in different languages, we resort multilingual pre-trained generative model to be our base model, which models the conditional probability of generating a new token given the previous generated tokens and the input context to the encoder c , i.e.,

$$P(x|c) = \prod_i P_{gen}(x_i|x_{<i}, c),$$

where x_i is the output of the decoder at step i .

Copy mechanism. Although the multilingual pre-trained generative models can generate sequences in many languages, solely relying on them may result in generating hallucinating arguments (Li et al., 2021). Since most of the tokens in the target output string appear in the input sequence,⁵ we augment the multilingual pre-trained generative models with a copy mechanism to help X-GEAR better adapt to the cross-lingual scenario. Specifically, we follow See et al. (2017) to decide the conditional probability of generating a token t as a weighted sum of the vocabulary distribution computed by multilingual pre-trained generative model P_{gen} and copy distribution P_{copy}

$$P_{X-GEAR}(x_i = t|x_{<i}, c) = w_{copy} \cdot P_{copy}(t) + (1 - w_{copy}) \cdot P_{gen}(x_i = t|x_{<i}, c)$$

where $w_{copy} \in [0, 1]$ is the copy probability computed by passing the decoder hidden state at time step i to a linear layer. As for P_{copy} , it refers to the probability over input tokens weighted by the cross-attention that the last decoder layer computed (at time step i). Our model is then trained end-to-end with the following loss:

$$\mathcal{L} = -\log \sum_i P_{X-GEAR}(x_i|x_{<i}, c).$$

5 Experiments

5.1 Datasets

We consider two commonly used event extraction datasets: ACE-2005 and ERE. We consider En-

⁵Except for the special tokens [and] and [None].

glish, Arabic, and Chinese annotations for ACE-2005 (Doddington et al., 2004) and follow the preprocessing in Wadden et al. (2019) to keep 33 event types and 22 argument roles. ERE (Song et al., 2015) is created by the Deep Exploration and Filtering of Test program. We consider its English and Spanish annotations and follow the preprocessing in Lin et al. (2020) to keep 38 event types and 21 argument roles. Detailed statistics and preprocessing steps about the two datasets are in Appendix A.

Notice that prior works working on the zero-shot cross-lingual transfer of event arguments mostly focus on event argument role labeling (Subburathinam et al., 2019; Ahmad et al., 2021), where they assume ground truth entities are provided during both training and testing. In their experimental data splits, events in a sentence can be scattered in all training, development, and test split since they treat each event-entity pair as a different instance. In this work, we consider event argument extraction (Wang et al., 2019; Wadden et al., 2019; Lin et al., 2020), which is a more realistic setting.

5.2 Evaluation Metric

We follow previous work (Lin et al., 2020; Ahmad et al., 2021) and consider the *argument classification F1 score* to measure the performance of models. An argument-role pair is counted as correct if both the argument offsets and the role type match the ground truth. Given the ground truth arguments \mathbf{a} , ground truth roles \mathbf{r} , predicted arguments $\tilde{\mathbf{a}}$, and predicted roles $\tilde{\mathbf{r}}$, the argument classification F1 score is defined as the F1 score between the set $\{(\mathbf{a}_i, \mathbf{r}_i)\}$ and the set $\{(\tilde{\mathbf{a}}_j, \tilde{\mathbf{r}}_j)\}$. For every model, we experiment with three different random seeds and report the average results.

5.3 Compared Models

We compare the following models and their implementation details are listed in Appendix B.

- **OneIE** (Lin et al., 2020), the state-of-the-art for monolingual event extraction, is a classification-based model trained with multitasking, including entity extraction, relation extraction, event extraction, and *event argument extraction*. We simply replace its pre-trained embedding with XLM-RoBERTa-large (Conneau et al., 2020) to fit the zero-shot cross-lingual setting. Note that the multi-task learning makes OneIE require *additional annotations*, such as named entity annotations and relation annotations.

- **CL-GCN** (Subburathinam et al., 2019) is a classification-based model for cross-lingual event argument role labeling (EARL). It considers *dependency parsing annotations* to bridge different languages and use GCN layers (Kipf and Welling, 2017) to encode the parsing information. We follow the implementation of previous work (Ahmad et al., 2021) and add two GCN layers on top of XLM-RoBERTa-large. Since CL-GCN focuses on EARL tasks, which assume the ground truth entities are available during testing, we add one name entity recognition module jointly trained with CL-GCN.
- **GATE** (Ahmad et al., 2021), the state-of-the-art model for zero-shot cross-lingual EARL, is a classification-based model which considers *dependency parsing annotations* as well. Unlike CL-GCN, it uses a Transformer layer (Vaswani et al., 2017) with modified attention to encode the parsing information. We follow the original implementation and add two GATE layers on top of pre-trained multilingual language models.⁶ Similar to CL-GCN, we add one name entity recognition module jointly trained with GATE.
- **TANL** (Paolini et al., 2021) is a generation-based model for monolingual EAE. Their predicted target is a sentence that embeds labels into the input passage, such as [Two soldiers|target] were attacked, which indicates that “Two soldiers” is a “target” argument. To adapt TANL to zero-shot cross-lingual EAE, we change its pre-trained generative model from T5 (Raffel et al., 2020) to mT5-base (Xue et al., 2021).
- **X-GEAR** is our proposed model. We consider three different pre-trained generative language models: mBART-50-large (Tang et al., 2020), mT5-base, and mT5-large (Xue et al., 2021).

5.4 Results

Table 1 and Table 2 list the results on ACE-2005 and ERE, respectively, with all combinations of source languages and target languages. Note that all the models have similar numbers of parameters

⁶To better compare our method with this strong baseline, we consider three different pre-trained multilingual language models for GATE – (1) XLM-RoBERTa-large (2) mBART-50-large (3) mT5-base. For mBART-50-large and mT5-base, we follow BART’s recipe (Lewis et al., 2020) to extract features for EAE predictions. Specifically, the input passage is fed into both encoder and decoder, and the final token representations are elicited from the decoder output.

Model	# of parameters	en	en	en	ar	ar	ar	zh	zh	zh	avg
		↓ en	↓ zh	↓ ar	↓ ar	↓ en	↓ zh	↓ zh	↓ en	↓ ar	
OneIE (XLM-R-large) (Lin et al., 2020)	~570M	63.6	42.5	37.5	57.8	27.5	<u>31.2</u>	<u>69.6</u>	51.5	31.1	45.8
CL-GCN (XLM-R-large) (Subburathinam et al., 2019)	~570M	59.8	29.4	25.0	47.5	25.4	19.4	62.2	40.8	23.3	37.0
GATE (XLM-R-large) (Ahmad et al., 2021)	~590M	67.0	49.2	<u>44.5</u>	59.6	27.6	26.3	70.6	46.7	37.3	47.6
GATE (mBART-50-large)	~630M	65.5	43.0	38.9	58.5	27.5	26.1	65.9	45.3	30.2	44.5
GATE (mT5-base)	~590M	59.8	47.7	32.6	45.4	20.7	21.0	64.0	35.3	22.8	38.8
TANL (mT5-base) (Paolini et al., 2021)	~580M	59.1	38.6	29.7	50.1	18.3	16.9	65.2	33.3	18.3	36.6
X-GEAR (mBART-50-large)	~610M	<u>68.3</u>	48.9	37.8	59.8	<u>30.5</u>	29.2	63.6	45.9	32.3	46.2
X-GEAR (mT5-base)	~580M	67.9	<u>53.1</u>	42.0	<u>66.2</u>	27.6	30.5	69.4	<u>52.8</u>	32.0	<u>49.1</u>
X-GEAR (mT5-large)	~1230M	71.2	54.0	44.8	68.9	32.1	33.3	68.9	55.8	<u>33.1</u>	51.3

Table 1: Average results in argument classification F1(%) of ACE-2005 with three different seeds. The best is in bold and the second best is underlined. “en \Rightarrow zh” denotes models transferring from en to zh. Compared with models using similar numbers of parameters, X-GEAR (mT5-base) outperforms baselines. To test the influence of using larger pre-trained generative models, we add X-GEAR (mT5-large), which achieves even better results.

Model	en	en	es	es	avg
	↓ en	↓ es	↓ es	↓ en	
OneIE (XLM-R-large)	64.4	56.8	64.8	56.9	60.7
CL-GCN (XLM-R-large)	61.9	51.9	62.9	48.5	55.9
GATE (XLM-R-large)	66.4	61.5	63.0	56.5	61.9
TANL (mT5-base)	65.9	40.3	58.6	47.4	53.1
X-GEAR (mBART-50-large)	69.5	57.3	63.9	58.9	62.4
X-GEAR (mT5-base)	<u>69.8</u>	57.9	<u>66.1</u>	<u>59.0</u>	<u>63.2</u>
X-GEAR (mT5-large)	72.9	<u>59.7</u>	67.4	64.1	66.0

Table 2: Average results in argument classification F1(%) of ERE with three different seeds. The best is in bold and the second best is underlined. “en \Rightarrow es” denotes that models transfer from en to es.

except for X-GEAR with mT5-large.

Comparison to prior generative models. We first observe that TANL has poor performance when transferring to different languages. The reason is that its language-dependent template makes TANL easily generate code-switching outputs,⁷ which is a case that pre-trained generative model rarely seen, leading to poor performance. In contrast, X-GEAR considers the language-agnostic templates and achieves better performance for zero-shot cross-lingual transfer.

Comparison to classification models. X-GEAR with mT5-base outperforms OneIE, CL-GCN, and GATE on almost all the combinations of the source language and the target language. This suggests that our proposed method is indeed a promising approach for zero-shot cross-lingual EAE.

It is worth noting that OneIE, CL-GCN, and GATE require an additional pipeline named entity recognition module to make predictions. Moreover, CL-GCN and GATE need additional dependency

⁷Such as the example shown in footnote 2.

parsing annotations to align the representations of different languages. On the contrary, X-GEAR is able to leverage the learned knowledge from the pre-trained generative models, and therefore no additional modules or annotations are needed.

Comparison to different pre-trained generative language models. Interestingly, using mT5-base is more effective than using mBART-50-large for X-GEAR, although they have a similar amount of parameters. We conjecture that the use of special tokens leads to this difference. mBART-50 has different begin-of-sequence (BOS) tokens for different languages. During generation, we have to specify which BOS token we would like to use as the start token. We guess that this language-specific BOS token makes mBART-50 harder to transfer the knowledge from the source language to the target language. Unlike mBART-50, mT5 does not have such language-specific BOS tokens. During generation, mT5 uses the padding token as the start token to generate a sequence. This design is more general and benefit zero-shot cross-lingual transfer.

Larger pre-trained models are better. Finally, we demonstrate that the performance of X-GEAR can be further boosted with a larger pre-trained generative language model. As shown by Table 1 and Table 2, X-GEAR with mT5-large achieves the best scores on most of the cases.

6 Analysis

6.1 Ablation Studies

Copy mechanism. We first study the effect of the copy mechanism. Table 3 lists the performance of X-GEAR with and without copy mechanism. It shows improvements in adding a copy mechanism

Model	en	ar	zh	xx	xx	xx	avg
	↓ xx	↓ xx	↓ xx	↓ en	↓ ar	↓ zh	
mBART-50-large - w/o copy	51.6	39.8	47.2	48.2	43.2	47.2	46.2
	50.9	42.2	49.6	50.6	43.5	48.7	47.6
mT5-base - w/o copy	54.3	41.4	51.4	49.4	46.7	51.0	49.1
	52.1	39.5	47.6	48.1	42.7	48.5	46.4
mT5-large - w/o copy	56.7	44.8	52.6	53.0	48.9	52.1	51.3
	55.1	45.0	51.5	52.0	46.3	53.2	50.5

Table 3: Ablation study on copy mechanism for ACE-2005. “en ⇒ xx” indicates the average of “en ⇒ en”, “en ⇒ zh”, and “en ⇒ ar”.

when using mT5-large and mT-base. However, interestingly, adding a copy mechanism is not effective for mBART-50. We conjecture that this is because the pre-trained objective of mBART-50 is denoising autoencoding (Liu et al., 2020), and it has already learned to copy tokens from the input. Therefore, adding a copy mechanism is less useful. In contrast, the pre-trained objective of mT5 is to only generate tokens been masked out, resulting in lacking the ability to copy input. Thus, the copy mechanism becomes beneficial for mT5.

Including event type in prompts. In Section 4, we mentioned that the designed prompt for X-GEAR consists of only the input sentence and the language-agnostic template. In this section, we discuss whether *explicitly* including the event type information in the prompt is helpful. We consider three ways to include the event type information:

- **English tokens.** We put the English version of the event type in the prompt even if we are training or testing on non-English languages, for example, using *Attack* for the event type *Attack*.
- **Translated tokens.** For each event type, we prepare the translated version of that event type token. For example, both *Attack* and 攻击 represents the *Attack* event type. During training or testing, we decide the used token(s) according to the language of the input passage. Since all the event types are written in English in ACE-2005 and ERE, we use an off-the-self machine translation tool to perform the translation.
- **Special tokens.** We create a special token for every event type and let the model learn the representations of the special tokens from scratch. For instance, we use <-attack-> to represent the *Attack* event type.

Table 4 shows the results. In most cases, including event type information in the prompt decreases

Model	en	ar	zh	xx	xx	xx	avg
	↓ xx	↓ xx	↓ xx	↓ en	↓ ar	↓ zh	
X-GEAR (mT5-base)	54.3	41.4	51.4	49.4	46.7	51.0	49.1
w/ English Tokens	53.3	39.3	52.3	49.2	46.5	49.2	48.3
w/ Translated Tokens	51.7	40.4	52.2	49.8	45.6	48.8	48.1
w/ Special Tokens	52.3	39.7	51.8	49.0	45.4	49.3	47.9

Table 4: Ablation study on including event type information in prompts for ACE-2005. “en ⇒ xx” indicates the average of “en ⇒ en”, “en ⇒ zh”, and “en ⇒ ar”.

Model	en	ar	zh	xx	xx	xx	avg
	↓ xx	↓ xx	↓ xx	↓ en	↓ ar	↓ zh	
X-GEAR (mT5-base)	54.3	41.4	51.4	49.4	46.7	51.0	49.1
w/ random order 1	54.4	38.9	50.8	48.7	45.1	50.1	48.0
w/ random order 2	52.1	40.4	51.4	48.3	45.9	49.7	48.0
w/ random order 3	53.7	40.8	50.7	50.8	45.8	48.6	48.4

Table 5: Ablation study on different orders of roles in templates for ACE-2005. “en ⇒ xx” indicates the average of “en ⇒ en”, “en ⇒ zh”, and “en ⇒ ar”.

the performance. One reason is that one word in a language can be mapped to several words in another language. For example, the *Life* event type is related to *Marry*, *Divorce*, *Born*, and *Die* four sub-event types. In English, we can use just one word *Life* to cover all four sub-event types. However, In Chinese, when talking about *Marry* and *Divorce*, *Life* should be translated to “生活”; when talking about *Born* and *Die*, *Life* should be translated to “生命”. This mismatch may cause the performance drop when considering event types in prompts. We leave how to efficiently use event type information in the cross-lingual setting as future work.

Influence of role order in templates. The order of roles in the designed language-agnostic templates can potentially influence performance. When designing the templates, we intentionally make the order of roles close to the order in natural sentences.⁸ To study the effect of different orders, we train X-GEAR with templates with different random orders and report the results in Table 5. X-GEAR with random orders still achieve good performance but slightly worse than the original order. It suggests that X-GEAR is not very sensitive to different templates while providing appropriate order of roles can lead to a small improvement.

Using English tokens instead of special tokens for roles in templates. In Section 4, we mentioned that we use language-agnostic templates

⁸For example, types related to subject and object are listed first and types related to methods and places are listed last.




Figure 3: Distribution of errors that made by X-GEAR (mT5-base). **Left:** The distribution for our model that transfers from Arabic to English; **Right:** The distribution for our model trained on Chinese and tested on English.

Model	en	ar	zh	xx	xx	xx	avg
	↓ xx	↓ xx	↓ xx	↓ en	↓ ar	↓ zh	
X-GEAR (mT5-base)	54.3	41.4	51.4	49.4	46.7	51.0	49.1
w/ English Tokens	51.4	39.3	49.7	46.6	44.7	49.0	46.8

Table 6: Comparison of using English tokens and special tokens for roles in templates. “en => xx” indicates the average of “en => en”, “en => zh”, and “en => ar”.

to facilitate the cross-lingual transfer. To further validate the effectiveness of the language-agnostic template. We conduct experiments using English tokens as the templates. Specifically, we set format

```
Agent: [None] <SEP> Victim: [None] <SEP> Instrument:
[None] <SEP> Place: [None]
```

to be the template for *Life:Die* events. Hence, for non-English instances, the targeted output string is a code-switching sequence. Table 6 lists the results. We can observe that applying language-agnostic templates bring X-GEAR 2.3 F1 scores improvements in average.

6.2 Error Analysis

We perform error analysis on X-GEAR (mT5-base) when transferring from Arabic to English and transferring from Chinese to English. For each case, we sample 30 failed examples and present the distribution of various error types in Figure 3.

Errors on both monolingual and cross-lingual models. We compare the predicted results from X-GEAR(ar => en) with X-GEAR(en => en), or from X-GEAR(zh => en) with X-GEAR(en => en). If their predictions are similar and both of them

are wrong when compared to the gold output, we classify the error into this category. To overcome the errors in this category, the potential solution is to improve monolingual models for EAE tasks.

Over-generating. Errors in this category happen more often in X-GEAR(ar => en). It is likely because the entities in Arabic are usually much longer than that in English when measuring by the number of sub-words. Based on our statistics, the average entity span length is 2.85 for Arabic and is 2.00 for English (length of sub-words). This leads to the natural for our X-GEAR(ar => en) to overly generate some tokens even though they have captured the correct concept. An example is that the model predicts “*The EU foreign ministers*”, while the ground truth is “*ministers*”.

Label disagreement on different language splits. The annotations for the ACE dataset in different language split contain some ambiguity. For example, given sentence “*He now also advocates letting in U.S. troops for a war against Iraq even though it is a fellow Muslim state.*” and the queried trigger “*war*”, the annotations in English tends to label *Iraq* as the *Place* where the event happen, while similar situations in other languages will mark *Iraq* as the *Target* for the war.

Grammar difference between languages. An example for this category is “... *Blackstone Group would buy Vivendi’s theme park division, including Universal Studios Hollywood ...*” and the queried trigger “*buy*”. We observe that X-GEAR(ar => en) predicts *Videndi* as the *Artifact* been sold and *division* is the *Seller*, while X-GEAR(en => en)

can correctly understand that *Vivendi* are the *Seller* and *division* is the *Artifact*. We hypothesize the reason being the differences between the grammar in Arabic and English. The word order of the sentence “*Vivendi’s theme park division*” in Arabic is reversed with its English counterpart, that is, “*theme park division*” will be written before “*Vivendi*” in Arabic. Such difference leads to errors in this category.

Generating words not appearing in the passage.

In X-GEAR(zh \Rightarrow en), we observe several cases that generate words not appearing in the passage. There are two typical situations. The first case is that X-GEAR(zh \Rightarrow en) mixes up singular and plural nouns. For example, the model generates “*studios*” as prediction while only “*studio*” appears in the passage. This may be because Chinese does not have morphological inflection for plural nouns. The second case is that X-GEAR(zh \Rightarrow en) will generate random predictions in Chinese.

Generating correct predictions but in Chinese.

This is a special case of “*Generating words not appearing in the passage*”. In this category, we observe that although the prediction is in Chinese (hence, a wrong prediction), it is correct if we translate the prediction into English.

6.3 Constrained Decoding

Among all the errors, we highlight two specific categories — “*Generating words not appearing in the passage*” and “*Generating correct predictions but in Chinese*”. These errors can be resolved by applying constrained decoding (Cao et al., 2021) to force all the generated tokens to appear input.

Table 7 presents the result of X-GEAR with constrained decoding. We observe that adapting such constraints indeed helps the cross-lingual transferability, yet it also hurts the performance in some monolingual cases. We conduct a qualitative inspection of the predictions. The observation is that constrained decoding algorithm although guarantees all generated tokens appearing in the input, the coercive method breaks the overall sequence distribution that learned. Hence, in many monolingual examples, once one of the tokens is corrected by constrained decoding, its following generated sequence changes a lot, while the original predicted suffixed sequence using beam decoding are actually correct. This leads to a performance decrease.⁹

⁹Indeed, a similar situation happens to cross-lingual cases;

Model	monolingual	cross-lingual	average all
X-GEAR (mBART-50-large)	63.9	37.4	46.2
w/ constrained decoding	62.4	37.6	45.9
X-GEAR (mT5-base)	67.8	39.7	49.1
w/ constrained decoding	67.0	39.9	48.9
X-GEAR (mT5-large)	69.7	42.2	51.3
w/ constrained decoding	68.8	43.1	51.6

Table 7: Results of applying constrained decoding. Breakdown numbers can be found in Appendix C. Based on whether the training languages are the same between training and testing, we classify the results into *monolingual* and *cross-lingual*, and we report the corresponding average for each category.

7 Conclusion

We present the first generation-based models for zero-shot cross-lingual event argument extraction. To overcome the discrepancy between languages, we design language-agnostic templates and propose X-GEAR, which well capture output dependencies and can be used without additional named entity extraction modules. Our experimental results show that X-GEAR outperforms the current state-of-the-art, which demonstrates the potential of using a language generation framework to solve zero-shot cross-lingual structured prediction tasks.

Acknowledgments

We thank anonymous reviewers for their helpful feedback. We thank the UCLA PLUSLab and UCLA-NLP group for the valuable discussions and comments. We also thank Steven Fincke, Shantanu Agarwal, and Elizabeth Boschee for their help on data preparation in Arabic. This work is supported in part by the Intelligence Advanced Research Projects Activity (IARPA), via Contract No. 2019-19051600007, and research awards sponsored by CISCO and Google.

Ethics Considerations

Our proposed models are based on the multilingual pre-trained language model that is trained on a large text corpus. It is known that the pre-trained language model could capture the bias reflecting the training data. Therefore, our models can potentially generate offensive or biased content learned by the pre-trained language model. We suggest carefully examining the potential bias before deploying our model in any real-world applications.

however, since the original performance for cross-lingual transfer is not high enough, the benefits of correcting tokens are more significant than this drawback.

References

- Wasi Uddin Ahmad, Nanyun Peng, and Kai-Wei Chang. 2021. GATE: graph attention transformer encoder for cross-lingual relation and event extraction. In *Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI)*.
- Wasi Uddin Ahmad, Zhisong Zhang, Xuezhe Ma, Kai-Wei Chang, and Nanyun Peng. 2019a. Cross-lingual dependency parsing with unlabeled auxiliary languages. In *The 2019 SIGNLL Conference on Computational Natural Language Learning (CoNLL)*.
- Wasi Uddin Ahmad, Zhisong Zhang, Xuezhe Ma, Edouard H. Hovy, Kai-Wei Chang, and Nanyun Peng. 2019b. On difficulties of cross-lingual transfer with order differences: A case study on dependency parsing. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Nicola De Cao, Ledell Wu, Kashyap Popat, Mikel Artetxe, Naman Goyal, Mikhail Plekhanov, Luke Zettlemoyer, Nicola Cancedda, Sebastian Riedel, and Fabio Petroni. 2021. Multilingual autoregressive entity linking. *arXiv preprint arXiv:2103.12528*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- George R. Doddington, Alexis Mitchell, Mark A. Przybocki, Lance A. Ramshaw, Stephanie M. Strassel, and Ralph M. Weischedel. 2004. The automatic content extraction (ACE) program - tasks, data, and evaluation. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC)*.
- Xinya Du, Alexander M. Rush, and Claire Cardie. 2021. GRIT: generative role-filler transformers for document-level event entity extraction. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.
- Steven Fincke, Shantanu Agarwal, Scott Miller, and Elizabeth Boschee. 2021. Language model priming for cross-lingual event extraction. *arXiv preprint arXiv:2109.12383*.
- I-Hung Hsu, Kuan-Hao Huang, Elizabeth Boschee, Scott Miller, Prem Natarajan, Kai-Wei Chang, and Nanyun Peng. 2021. DEGREE: A data-efficient generative event extraction model. *arXiv preprint arXiv:2108.12724*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*.
- Kuan-Hao Huang, Wasi Uddin Ahmad, Nanyun Peng, and Kai-Wei Chang. 2021a. Improving zero-shot cross-lingual transfer learning via robust training. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Kung-Hsiang Huang, Sam Tang, and Nanyun Peng. 2021b. Document-level entity-based extraction as template generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Xiaolei Huang, Jonathan May, and Nanyun Peng. 2019. What matters for neural cross-lingual named entity recognition: An empirical analysis. In *2019 Conference on Empirical Methods in Natural Language Processing (EMNLP), short*.
- Thomas N. Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations (ICLR)*.
- 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, ACL 2020, Online, July 5-10, 2020*.
- Sha Li, Heng Ji, and Jiawei Han. 2021. Document-level event argument extraction by conditional generation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- 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 (ACL)*.
- Jian Liu, Yubo Chen, Kang Liu, and Jun Zhao. 2019. Neural cross-lingual event detection with minimal parallel resources. 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)*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*.

- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. Multilingual denoising pre-training for neural machine translation. *Trans. Assoc. Comput. Linguistics*, 8:726–742.
- Yaojie Lu, Hongyu Lin, Jin Xu, Xianpei Han, Jialong Tang, Annan Li, Le Sun, Meng Liao, and Shaoyi Chen. 2021. Text2event: Controllable sequence-to-structure generation for end-to-end event extraction. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL/IJCNLP)*.
- Tao Meng, Nanyun Peng, and Kai-Wei Chang. 2019. Target language-aware constrained inference for cross-lingual dependency parsing. In *2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Minh Van Nguyen and Thien Huu Nguyen. 2021. Improving cross-lingual transfer for event argument extraction with language-universal sentence structures. In *Proceedings of the Sixth Arabic Natural Language Processing Workshop*.
- Jian Ni and Radu Florian. 2019. Neural cross-lingual relation extraction based on bilingual word embedding mapping. 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)*.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. Cross-lingual name tagging and linking for 282 languages. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Giovanni Paolini, Ben Athiwaratkun, Jason Krone, Jie Ma, Alessandro Achille, Rishita Anubhai, Cícero Nogueira dos Santos, Bing Xiang, and Stefano Soatto. 2021. Structured prediction as translation between augmented natural languages. In *9th International Conference on Learning Representations (ICLR)*.
- Hao Peng, Ankur P. Parikh, Manaal Faruqui, Bhuwan Dhingra, and Dipanjan Das. 2019. Text generation with exemplar-based adaptive decoding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. Stanza: A python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations (ACL)*.
- Guanghui Qin and Jason Eisner. 2021. Learning how to ask: Querying lms with mixtures of soft prompts. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- 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. *J. Mach. Learn. Res.*
- Sebastian Ruder, Noah Constant, Jan Botha, Aditya Siddhant, Orhan Firat, Jinlan Fu, Pengfei Liu, Junjie Hu, Dan Garrette, Graham Neubig, and Melvin Johnson. 2021. XTREME-R: towards more challenging and nuanced multilingual evaluation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Teven Le Scao and Alexander M. Rush. 2021. How many data points is a prompt worth? In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Timo Schick and Hinrich Schütze. 2021. Exploiting cloze-questions for few-shot text classification and natural language inference. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.
- 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, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*.
- Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. 2020. Towards controllable biases in language generation. In *the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)-Findings, long*.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Zhiyi Song, Ann Bies, Stephanie M. Strassel, Tom Riese, Justin Mott, Joe Ellis, Jonathan Wright, Seth Kulick, Neville Ryant, and Xiaoyi Ma. 2015. From light to rich ERE: annotation of entities, relations, and events. In *Proceedings of the The 3rd Workshop on EVENTS: Definition, Detection, Coreference, and Representation, (EVENTS@HLP-NAACL)*.
- Ananya Subburathinam, Di Lu, Heng Ji, Jonathan May, Shih-Fu Chang, Avirup Sil, and Clare R. Voss. 2019. Cross-lingual structure transfer for relation and 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)*.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. Multilingual translation with extensible multilingual pretraining and finetuning. *arXiv preprint arXiv:2008.00401*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017 (NeurIPS)*.
- 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)*.
- Xiaozhi Wang, Ziqi Wang, Xu Han, Zhiyuan Liu, Juanzi Li, Peng Li, Maosong Sun, Jie Zhou, and Xiang Ren. 2019. HMEAE: hierarchical modular event argument 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)*.
- Genta Indra Winata, Andrea Madotto, Zhaojiang Lin, Rosanne Liu, Jason Yosinski, and Pascale Fung. 2021. Language models are few-shot multilingual learners. *arXiv preprint arXiv:2109.07684*.
- Haoran Xu, Seth Ebner, Mahsa Yarmohammadi, Aaron Steven White, Benjamin Van Durme, and Kenton W. Murray. 2021. Gradual fine-tuning for low-resource domain adaptation. *arXiv preprint arXiv:2103.02205*.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. mt5: A massively multilingual pre-trained text-to-text transformer. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*.
- Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. 2021. A unified generative framework for various NER subtasks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL/IJCNLP)*.
- Bowei Zou, Zengzhuang Xu, Yu Hong, and Guodong Zhou. 2018. Adversarial feature adaptation for cross-lingual relation classification. In *Proceedings of the 27th International Conference on Computational Linguistics (COLING)*.

A Dataset Statistics and Data Preprocessing

Table 8 presents the detailed statistics for the ACE-2005 dataset and ERE dataset.

For the English and Chinese splits in ACE-2005, we use the setting provided by Wadden et al. (2019) and Lin et al. (2020), respectively. As for Arabic part, we adopt the setup proposed by Xu et al. (2021). Observing that part of the sentence breaks made from Xu et al. (2021) being extremely long for pretrained models to encode, we perform additional preprocessing and postprocessing procedures for Arabic data. Specifically, we split Arabic sentences into several portions that any of the portion is shorter than 80 tokens. Then, we map the models' predictions of the split sentences back to the original sentence during postprocessing.

B Implementation Details

We describe the implementation details for all the models as follows:

- **OneIE** (Lin et al., 2020). We use their provided code¹⁰ to train the model with the provided default settings. It is worth mention that for the Arabic split in the ACE-2005 dataset, OneIE is trained with only entity extraction, event extraction, and event argument extraction since there is no relation labels in Xu et al. (2021)'s preprocessing script. All other parameters are set to the default values.
- **CL-GCN** (Subburathinam et al., 2019). We refer the released code from Ahmad et al. (2021)¹¹ to re-implement the CL-GCN method. Specifically, we adapt the baseline framework that described and implemented in OneIE's code (Lin et al., 2020), but we remove its relation extraction module and add two layers of GCN on top of XLM-RoBERTa-large. The pos-tag and dependency parsing annotations are obtained by applying Stanza (Qi et al., 2020). All other parameters are set to be the same as the training of OneIE.
- **GATE** (Ahmad et al., 2021). We refer the official released code from Ahmad et al. (2021) to re-implement GATE. Similar to CL-GCN, we adapt the baseline framework that described and implemented in OneIE's code, but we remove

its relation extraction module and add two layers of GATE on top of XLM-RoBERTa-large, mT5, or mBART-50-large. The pos-tag and dependency parsing annotations are also obtained by applying Stanza (Qi et al., 2020). The hyperparameter of δ in GATE is set to be [2, 2, 4, 4, ∞ , ∞ , ∞ , ∞]. All other parameters are set to be the same as the training of OneIE.

- **TANL** (Paolini et al., 2021). To adapt TANL to zero-shot cross-lingual EAE, we adapt the public code¹² and replace its pre-trained based model T5 (Raffel et al., 2020) with mT5-base (Xue et al., 2021). All other parameters are set to their default values.
- **X-GEAR** is our proposed model. We consider three different pre-trained generative language models: mBART-50-large (Tang et al., 2020), mT5-base, and mT5-large (Xue et al., 2021). When fine-tune the pre-trained models, we set the learning rate to 10^{-4} for mT5, and 10^{-5} for mBART-50-large. The batch size is set to 8. The number of training epochs is 60.

C Constrained Decoding Detailed Results

Table 9 shows the detailed results for X-GEAR using constrained decoding algorithm during testing time. We directly apply constrained decoding algorithms on the trained models we have in Table 1.

¹⁰<http://blender.cs.illinois.edu/software/oneie/>

¹¹<https://github.com/wasiahmad/GATE>

¹²<https://github.com/amazon-research/tanl>

Dataset	Lang.	Train			Dev			Test		
		#Sent.	#Event	#Arg.	#Sent.	#Event	#Arg.	#Sent.	#Event	#Arg.
ACE-2005	en	17172	4202	4859	923	450	605	832	403	576
	ar	2722	1743	2506	289	117	174	272	198	287
	zh	6305	2926	5581	486	217	404	482	190	336
ERE	en	14734	6208	8924	1209	525	730	1161	551	882
	es	4582	3131	4415	311	204	279	323	255	354

Table 8: Dataset statistics of ACE-2005 and ERE.

Model	en	en	en	ar	ar	ar	zh	zh	zh	avg (mono.)	avg (cross.)	avg (all)
	↓ en	↓ zh	↓ ar	↓ ar	↓ en	↓ zh	↓ zh	↓ en	↓ ar			
X-GEAR (mBART-50-large)	68.3	48.9	37.7	59.8	30.5	29.2	63.6	45.9	32.3	63.9	37.4	46.2
w/ constrained decoding	68.0	49.1	37.8	59.5	30.6	29.2	59.7	47.7	31.3	62.4	37.6	45.9
X-GEAR (mT5-base)	67.9	53.1	42.0	66.2	27.6	30.5	69.4	52.8	32.0	67.8	39.7	49.1
w/ constrained decoding	67.9	53.1	42.0	66.2	27.8	30.4	66.7	53.1	33.1	67.0	39.9	48.9
X-GEAR (mT5-large)	71.2	54.0	44.8	68.9	32.1	33.3	68.9	55.8	33.1	69.7	42.2	51.3
w/ constrained decoding	71.2	54.8	45.6	68.9	32.0	33.3	66.2	57.7	35.0	68.8	43.1	51.6

Table 9: The detailed breakdown results for applying constrained decoding on X-GEAR. The avg(mono.) column represents the results that average over values in $en \Rightarrow en$, $zh \Rightarrow zh$, and $ar \Rightarrow ar$. The avg(cross.) column represents the results that average over values in $en \Rightarrow zh$, $en \Rightarrow ar$, $zh \Rightarrow en$, $zh \Rightarrow ar$, $ar \Rightarrow en$, and $ar \Rightarrow zh$.