

# STAR: Boosting Low-Resource Information Extraction by Structure-to-Text Data Generation with Large Language Models

Mingyu Derek Ma<sup>1</sup>, Xiaoxuan Wang<sup>1</sup>, Po-Nien Kung<sup>1</sup>  
P. Jeffrey Brantingham<sup>2</sup>, Nanyun Peng<sup>1</sup>, Wei Wang<sup>1</sup>

<sup>1</sup>Department of Computer Science, University of California, Los Angeles

<sup>2</sup>Department of Anthropology, University of California, Los Angeles

{ma, xw27, ponienkung, violetpeng, weiwang}@cs.ucla.edu, branting@ucla.edu

## Abstract

Information extraction tasks such as event extraction require an in-depth understanding of the output structure and sub-task dependencies. They heavily rely on task-specific training data in the form of (passage, target structure) pairs to obtain reasonable performance. However, obtaining such data through human annotation is costly, leading to a pressing need for low-resource information extraction approaches that require minimal human labeling for real-world applications. Fine-tuning supervised models with synthesized training data would be a generalizable method, but the existing data generation methods either still rely on large-scale ground-truth data or cannot be applied to complicated IE tasks due to their poor performance. To address these challenges, we propose STAR, a data generation method that leverages Large Language Models (LLMs) to synthesize data instances given limited seed demonstrations, thereby boosting low-resource information extraction performance. Our approach involves generating target structures ( $Y$ ) followed by generating passages ( $X$ ), all accomplished with the aid of LLMs. We design fine-grained step-by-step instructions to obtain the initial data instances. We further reduce errors and improve data quality through self-reflection error identification and self-refinement with iterative revision. Our experiments show that the data generated by STAR significantly improve the performance of low-resource event extraction and relation extraction tasks, even surpassing the effectiveness of human-curated data. Human assessment of the data quality shows STAR-generated data exhibit higher passage quality and better align with the task definitions compared with the human-curated data.

## 1 Introduction

Information extraction (IE) aims to extract knowledge of certain perspectives from natural language and consolidate it into an output structure (Ma et al. 2021b). To induce the target structure, the IE models need to understand fine-grained task requirements and constraints. Taking event extraction (EE), which is a component for IE systems to identify event triggers, event types, and their related details as arguments, as an example, task-specific rules include the predicted spans should be subsequences of the input passage, and arguments should be participants or attributes of the event. EE

models are also expected to be aware of the dynamic skeleton of the event structure because the different predicted event types result in their respective sets of argument roles being filled in. Supervised models learn the implicit requirement and ontology knowledge from training data in the form of (passage, target structure) pairs (Ma et al. 2021a). Prompt-based inference-only approaches with Large Language Models (LLMs) are shown to be unable to solve these complicated IE tasks (Li et al. 2023; Gao et al. 2023; Han et al. 2023). In real-world applications, text from various sources and domains contains a broad range of output spaces and label definitions. It is costly and rigid to annotate sufficient training data, thus, performing IE given minimal seed data instances is of particular interest for realistic IE applications.

The limited resources make understanding the task requirement and structure skeleton even harder. Existing low-resource IE methods mostly leverage other tasks via transfer learning (Huang et al. 2018; Zhang, Wang, and Roth 2021) or reformulate the task into alternative data-rich tasks for indirect supervision (Sainz et al. 2022; Xu, Ma, and Chen 2023; Lu et al. 2022). These works heavily depend on the availability of the source tasks' data and the compatibility between tasks, limiting its application to richer and broader output label spaces and under-represented domains.

Synthesizing additional training data to fine-tune supervised models would be a generalizable method and has demonstrated its success for tasks like sentiment analysis (Ye et al. 2022) and relation extraction (Josifoski et al. 2023). Some works annotate unlabeled examples with existing models as weak annotators ( $X_{gold} \rightarrow Y$ ) (He et al. 2021; Chia et al. 2022), produce analogous input ( $X_{gold} \rightarrow X'$ ) (Kumar, Choudhary, and Cho 2020; Lee et al. 2021), or generate input assuming the labels are available ( $Y_{gold} \rightarrow X$ ) (Meng et al. 2022; Gao, Fisch, and Chen 2021; Josifoski et al. 2023). However, these works require ground-truth  $X_{gold}$  or  $Y_{gold}$ , limiting their generalizability and scalability. What's more, they are designed for classification or straightforward IE tasks, and the performance drops significantly as the task complexity increases. Applying them to complicated IE tasks with sub-task dependencies and dynamic output structures such as EE produces noisy data that may impair task performance.

In this paper, we present STAR, a Structure-to-Text Data GeneRation pipeline to produce dependable data instances

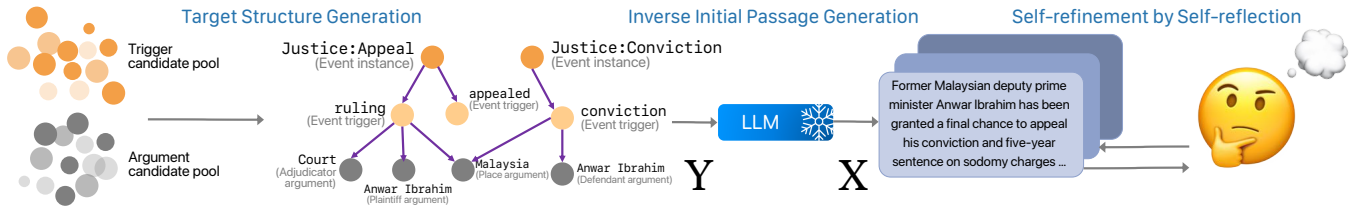


Figure 1: The STAR inverse data generation strategy using event extraction task as an example. We first generate target structures from valid trigger and argument candidates. Then we prompt the LLM with task instructions from different task granularities to generate the initial passage  $X_0$  containing the event information in the given target structure  $Y$ . Finally, we create self-reflection questions to prompt LLM to identify quality issues automatically and refine the passage with template-based hindsight feedback.

for low-resource IE. Instead of using existing models to produce silver target structures  $Y$  derived from the input passages  $X$  (i.e.  $X_{gold} \rightarrow Y$ ) to enrich the training data, we propose to generate data instances inversely by producing target structures from scratch first and then prompting the LLM to generate a passage (i.e.  $Y \rightarrow X$ ) containing the target structure information. This inverse design reformulates the synthetic data generation task from structure induction, where models often struggle, to conditional text generation, where LLMs excel. STAR contains three components. First, STAR generates diverse target structures  $Y$  from scratch, requiring minimal human efforts to initiate the data generation. In addition, with the ability to customize target structures and control their distribution, we can mitigate data imbalance and improve data diversity by producing target structures encompassing a broader range of triggers, event types and arguments, as well as their various combinations. Second, STAR performs instruction-guided data generation to prompt the LLM about the fine-grained task definition and constraints to produce passage  $X$ . Third, STAR detects errors in the generated data instances via self-reflection and provides hindsight natural language intervention to self-refine the generated data without additional human efforts.

Experimental results on event extraction (EE) and relation extraction (RE) tasks show that STAR is capable of generating human-level IE data instances given a couple of exemplar instances as demonstrations without the need for additional ground-truth passage or target structures. For EE, training the supervised model on STAR-generated data improves the argument classification sub-task by up to 12.91 points in F1 score on the ACE05 dataset, 2.9 points higher than using the same amount of human-curated data. For RE, we observe a 5.41-point F1 score improvement on the TACRED dataset, which is comparable to using the human-curated data. The improvements brought by STAR-generated data to multiple supervised models across multiple IE tasks demonstrate the generalizability and compatibility of STAR. Our manual data examination indicates STAR-produced data exhibits higher passage quality and better aligns with the task definitions compared with human-curated data.

We further conduct a detailed analysis of different methods for using LLMs to improve EE performance, and we show that training the supervised models on STAR-generated data yields at least 27.51 points higher F1 score for argument classification over the best inference-only LLM formulation.

Our contributions are three-fold: 1) We introduce STAR, an inverse structure-to-text IE data generation method with self-refinement; 2) We demonstrate the effectiveness, generalizability and compatibility of STAR by showcasing significant EE and RE performance enhancements and manual data quality examination; 3) We conduct a thorough analysis of different methods for using LLMs in EE, and the most effective approach identified is employing LLMs to generate training data for fine-tuning supervised models.

## 2 STAR: Structure-to-Text Data Generation

We introduce the method design of STAR as illustrated in Figure 1. We use event extraction (EE) as the exemplar IE task in this section as it covers more method details. The goal is to create  $N$  new data instances  $(X, Y)$  based on  $k$  demonstration instances to be used as additional training resources for supervised IE models. Each data instance is composed of a natural language **passage**  $X$  containing event information and a **target structure**  $Y$  containing 0 to any number of *events*, each contains an event trigger, its event type and 0 to any number of (argument mention, argument role) pairs. There are three steps: 1) target structure generation to prepare  $Y$  (§2.1), 2) initial passage generation to generate  $X_0$  (§2.2), and 3) self-refinement with self-reflection to revise  $X_0$  to  $X_t$  where  $t$  is the times of revision (§2.3). We introduce the adaptation of STAR to other IE tasks in §2.4.

### 2.1 Target Structure Generation

The output distribution of the structure prediction dataset largely determines the generalizability and robustness of the model fine-tuned on it. We first generate a pool of valid seed words for triggers of each event type, and for arguments of each (event type, argument role) combination. We then create a target structure  $Y$ . During the process of generating target structures, we particularly employ target structure distribution control to alleviate data imbalance and improve data diversity and comprehensiveness.

**Trigger candidate generation.** We prompt the LLM with 1) a definition of the selected event type; and 2) a few passages that contain event triggers of the selected event type as demonstrations. We use special tags to wrap the trigger word in the demonstration passages and prompt the LLM to continuously generate more passages with trigger words

wrapped within tag pairs. Then we parse the response and extract the trigger candidate words.

**Argument candidate generation.** We find reasonable arguments for a certain pair of event type and argument role by prompting LLM with: 1) a definition of the argument role under this specific event type; 2) the entity type we are looking for (e.g. a vehicle). The prompt would be like “Given the definition of Instrument argument as ‘The device used to inflict the harm’, what are some possible vehicle names that can be used as Instrument?” for INSTRUMENT argument of LIFE:INJURE event type. The allowed entity types of arguments are provided in an event ontology, and we merge generated word pools returned from separate queries if there are multiple allowed entity types for an argument role. For example, the ORIGIN argument of the MOVEMENT:TRANSPORT event type could be of entity types GPE (geopolitical entity), LOC (location), or FAC (facility name). We further parse the numbered/bullet lists generated by LLM to get argument candidate words.

**Creating target structure and distribution control.** We randomly sample trigger and argument candidates to create target structure  $Y$ . However, unbalanced label distribution and the prevalence of dominant labels pose challenges in many existing human-curated datasets (Ma et al. 2023b; Zhou, Kaneko, and Bollegala 2022; Cao et al. 2022; Zhao et al. 2018). In EE, a single dominant trigger word eclipses other relevant terms, leading to an unbalanced representation (Tong et al. 2022). We address the issues of imbalanced event type and trigger distribution by evenly generating data instances and significantly expanding the pool size for trigger candidates to 100, which is  $1.4\times$  to  $50\times$  larger for various event types compared to the human-curated ACE05 dataset. Additionally, we balance the **argument hallucination ratio** in the generated data. This involves ensuring that the generated dataset contains events with both many arguments and few arguments by uniformly replacing argument value with `None` across different argument hallucination ratios. Furthermore, we balance the **event density** in passage  $X$  by providing target structure  $Y$  with 0 to 5 events. Such target structure distribution controls from multiple perspectives ensure the EE models trained on the generated data points effectively learn features and patterns associated with different event and argument densities, enhancing robustness and generalizability.

## 2.2 Instruction-Guided Passage Generation

We use task instruction from multiple task granularities to provide recipes for the LLM to generate passages containing structured event information. The instruction is appended to  $k$  in-context learning examples verbalized by our instance verbalizer. Finally, we provide the verbalized target structure information based on the target structure  $Y$  to prompt the LLM to generate the initial passage  $X_0$ .

**Task-level instruction.** We provide task-related instruction following the annotation guideline curated by experts to guide the human annotation process of the ACE05 dataset (Doddington et al. 2004). Specifically, we provide:

1) a definition of “event”, “trigger”, “participant arguments” and “attribute arguments”; 2) an overall task requirement that the goal is to generate a sentence containing the event trigger words and arguments; 3) hallucination clarification that instructs the model *not* to generate arguments of certain roles if we explicitly provide that “the argument is None”; 4) multiple event clarification that information from multiple events should be contained in a *single* passage.

**Event type-level instruction.** In this segment, we introduce meta-information provided by pre-defined event ontology for a specific event type, including the name and definition of the event type and *each* possible argument roles. We provide all possible argument roles instead of the ones with existing values to ensure the generated passage  $X$  does not contain hallucinated arguments that should not appear according to the output structure  $Y$ .

**Instance-level verbalizer.** We verbalize exemplar data instances and target structure  $Y$  into natural language sequences with three segments: 1) the number of events in the passage; 2) the content of the event target structure; 3) the passage  $X$  with tags wrapping triggers and arguments to explicit hint the LLM about the roles and positions of the keywords, e.g. “<Plaintiff>He</Plaintiff> threatened to <Trigger>sue</Trigger> the company.” could provide an explicit indicator to the LLM that “he” is served as a PLAIN-TIFF argument for the event triggered by “sue”.

## 2.3 Self-refinement by Self-reflection

After the initial passage  $X_0$  is generated, the self-refinement mechanism evaluates the quality and identifies potential errors and further improves the initial generation results through iterative updates (Madaan et al. 2023; Wang et al. 2023; Liu, Sferrazza, and Abbeel 2023). In the  $t$ -th refinement iteration, we first identify the potential quality issues of  $X_{t-1}$  from a diverse set of quality dimensions (e.g. the passage contains CRIME argument information, but it should be “None” according to  $Y$ ), then the issues are **feedback** to the LLM by providing a template-based natural language intervention (e.g. “The passage contains a hallucinated argument CRIME incorrectly, remove CRIME information for event triggered by ‘jailed’.”) along with the generated passage of the previous iteration  $X_{t-1}$ , so that the LLM could **refine** the passage and produce  $X_t$ .

We define a set of quality dimensions and their intervention template manually. For EE, they include 1) whether the trigger/argument mention is a subsequence of the passage; 2) whether a trigger is used to initiate an occurrence; 3) whether an argument is used as an event participant or attribute of the specific event; 4) whether the argument is serving the required argument role; 5) whether the passage contains information that could serve as an argument that should not appear; 6) whether POS tags of the argument mentions in the passage context match the provided ones.

For each quality dimension, we query LLM with questions like “Is ‘Syria’ a DESTINATION argument describing the event triggered by ‘flee’?”. We then standardize LLM’s response to a binary error identification flag by checking whether the response entails a confirmative phrase “Yes,

it is.” with a Natural Language Inference model fine-tuned on MultiNLI (Williams, Nangia, and Bowman 2018) based on BART-large (Lewis et al. 2020). If a quality issue is flagged, we use the intervention template corresponding to the selected quality dimension as part of the feedback to the LLM for iterative revision. Such a self-reflection design makes the self-refinement process generalizable and robust since the entire error identification process through self-reflection and the revision process are done by the LLM itself without external add-on components.

## 2.4 Adaptation to Relation Extraction

RE’s relation type would be the equivalent concept of “event type” in EE. For target structure generation, we generate entity candidates using seed data instances’ entities as in-context examples. We then randomly pair entity candidates and assign a relation between the two entities. For initial passage generation, we use relation type definition instead. For self-refinement, we use the quality dimensions: 1) whether the given entities are contained in the generated passage, 2) whether there is a relation between them, and 3) whether they hold the certain relation provided in  $Y$ .

## 3 Experiments on Event Extraction

To evaluate the efficiency of STAR-generated data for event extraction, we compare the performance of supervised EE models with and without using the STAR-generated data.

### 3.1 Baselines

We use two types of models as our baselines: the *inference-only methods*, and the *supervised models* fine-tuned on data created by various *data creation strategies*.

**Inference-only EE methods.** We use LLM GPT-3.5 (OpenAI. 2022) and GPT-4 (OpenAI. 2023) to perform inference.<sup>1</sup> We adopt different EE input-target formulations to prompt LLMs, including formulations inspired by generative supervised models (1-3) and LLM prompting methods specifically designed for EE proposed by recent works (4-6). The formulations include: 1) **Examples & IO (Text2Event)** (Lu et al. 2021) uses a concise but *unnatural* template to represent event structure. 2) **Examples & IO (DEGREE)** (Hsu et al. 2022) generates a filled-in *natural* language template. 3) **Examples & IO (DICE)** (Ma et al. 2023c) is similar to DEGREE but uses separate queries for different argument roles. 4) **Task Instruction** (Li et al. 2023) provides task description and pre-defined event type *names*. 5) **Instruction+Examples** (Gao et al. 2023) provides event type *definitions* and positive and negative examples, in addition to the task description. 6) **Code4Struct** (Wang, Li, and Ji 2023) formulates task definition, event type definition and examples in Python code. For baselines 1-3, we follow the original input and target formulations and additionally provide  $k$  demonstration input-target pairs contained in the input prompt for in-context learning. Baselines 4 and 5 only support Tri-I and Tri-C, and baseline 6 only supports Arg-I and Arg-C.

<sup>1</sup>We use gpt-3.5-turbo-0301 and gpt-4-0314.

**Supervised EE models.** We use two representative EE models as the testbed to evaluate the quality of the generated data. **OneIE** (Lin et al. 2020) is a multi-task sequence-tagging model trained with global features based on RoBERTa-large (Liu et al. 2019). Note that we remove the human-labeled data for tasks other than EE (*e.g.* entity identification, relation extraction) to enable a fair comparison. **DEGREE** (Hsu et al. 2022) is a prompt-based model that fills in event type-specific human written templates based on a BART-large pre-trained model (Lewis et al. 2020).

**Data creation strategies.** Besides STAR, we introduce two other approaches to obtain training data. **Weakly Supervision:** we use the best inference-only model for EE (*i.e.* Inst.+Examples for Tri-I and Tri-C, and Code4Struct for Arg-I and Arg-C) to predict event structure  $Y'$  from passage  $X$ , and the  $(X, Y')$  pairs are used as training data. **Human:** we use human-curated data instances randomly sampled from the ACE05 dataset, which requires much more human annotators’ efforts, as an ideal but unrealistic setting.

### 3.2 Experimental Setup

We denote  $k$  as the number of demonstration examples of each event type used as in-context demonstrations for the inference-only methods and data creation strategies, and  $N$  as the number of data instances per event type created by data creation strategies. Supervised models are trained on  $k + N$  data instances per event type. We use the full test set for evaluation. We use the event ontology and data instances from the widely used sentence-level English event extraction dataset ACE05 (Doddington et al. 2004).

We follow previous EE works (Lin et al. 2020) and report F1 scores for four tasks.<sup>2</sup> 1) Trigger Identification: identified trigger span is correct. 2) Trigger Classification: its predicted *event type* is also correct. 3) Argument Identification: identified argument span is correct. 4) Argument Classification: its predicted *argument role* is also correct. Note that each task is dependent on the output of the previous task. We report the medium result for three runs of different random seeds.

### 3.3 Effectiveness of Data Generation

Table 1 shows the EE performance when using various numbers ( $k$ ) of demonstrations and Figure 2 further shows the effects of various amounts ( $N$ ) of augmented data instances.

**STAR boosts low-resource EE performance.** Table 1 shows that data generated by STAR significantly improve the supervised models’ performance (line 9 vs 11-12, and line 14 vs 16-17) across all tasks. The F1 scores of Arg-C are improved by 9.69 and 12.91 respectively for OneIE and DEGREE when  $k = 10$ .

**Data produced by STAR is more effective than human-curated ones given sufficient examples.** Compared to human-curated data instances (line 13 and 18 in Table 1), training supervised models with STAR-generated data leads

<sup>2</sup>The extracted trigger/argument has to match the ground truth *exactly*, instead of head word matching or coreference matching used in works (Li, Ji, and Han 2021; Wang, Li, and Ji 2023).

#		$k = 0$ 510			$k = 0$ 510			$k = 0$ 510			$k = 0$ 510			
		Trigger Iden.			Trigger Clas.			Argument Iden.			Argument Clas.			
<i>LLM</i>		<i>Formulation</i>		<i>Inference-only Methods</i>										
1	GPT-3.5	E&IO (Text2Event)	0.00	9.23	11.30	0.00	2.12	3.47	0.00	0.87	1.03	0.00	0.31	0.44
2		E&IO (DEGREE)	0.00	14.39	17.52	0.00	3.17	6.21	0.00	1.02	2.47	0.00	0.92	1.98
3		E&IO (DICE)	0.00	15.13	16.94	0.00	4.11	7.09	0.00	0.71	1.65	0.00	0.33	0.97
4		Task Inst. <sup>§</sup>	18.31	18.31	18.31	8.37	8.37	8.37	—	—	—	—	—	—
5		Inst.+Examples	29.44	47.24	59.71	21.56	40.57	53.29	—	—	—	—	—	—
6		Code4Struct	—	—	—	—	—	—	12.33	18.34	23.74	9.72	14.85	19.10
7	GPT-4	Inst.+Examples	<b>34.31</b>	<b>52.55</b>	<b>62.12</b>	<b>27.35</b>	<b>46.57</b>	<b>56.46</b>	—	—	—	—	—	—
8		Code4Struct	—	—	—	—	—	—	<b>17.51</b>	<b>24.50</b>	<b>27.62</b>	<b>11.89</b>	<b>24.28</b>	<b>25.48</b>
<i>Supervised Models</i> ( $N = 50$ except line 9 & 14)														
<i>EE Model</i>		<i>Data Creation</i>												
9	OneIE	None ( $N = 0$ )	0.00	57.24	60.55	0.00	52.38	54.84	0.00	29.06	36.45	0.00	25.85	33.56
10		Weak Sup.	29.48	49.23	51.66	23.61	45.02	45.23	16.19	24.35	26.84	10.47	19.14	22.94
11		STAR (GPT-3.5)	42.61	63.08	64.12	36.65	56.61	57.29	30.32	39.76	43.40	24.36	36.17	40.93
12		STAR (GPT-4)	<b>45.42</b>	<b>64.63</b>	<b>66.77</b>	<b>39.15</b>	<b>58.84</b>	<b>60.76</b>	<b>32.23</b>	<b>42.76</b>	<b>46.22</b>	<b>27.47</b>	<b>39.53</b>	<b>43.25</b>
13		Human <sup>†§</sup>	65.62	65.62	65.62	60.10	60.10	60.10	44.76	44.76	44.76	41.60	41.60	41.60
14	DEGREE	None ( $N = 0$ )	0.00	55.62	57.65	0.00	50.69	52.49	0.00	31.77	42.29	0.00	30.19	40.08
15		Weak Sup.	27.51	46.48	49.70	22.23	41.65	43.55	18.14	32.53	33.33	13.45	27.38	30.01
16		STAR (GPT-3.5)	43.74	61.39	<u>63.57</u>	38.90	56.41	<u>59.10</u>	32.32	48.73	<u>53.06</u>	28.21	46.55	<u>50.97</u>
17		STAR (GPT-4)	<b>46.69</b>	<b>64.47</b>	<b>65.17</b>	<b>41.75</b>	<b>59.92</b>	<b>61.42</b>	<b>35.85</b>	<b>51.92</b>	<b>54.56</b>	<b>32.09</b>	<b>50.74</b>	<b>52.99</b>
18		Human <sup>†§</sup>	63.49	63.49	63.49	58.86	58.86	58.86	52.47	52.47	52.47	50.09	50.09	50.09

Table 1: Event extraction performance (F1, %) when using  $k$  seed data instances per event type. *Inference-only methods* and *data creation methods* use the same set of  $k$  examples for each event type to prompt LLM to perform EE and generate data instances respectively. *Supervised EE models* are trained on  $k + N$  data instances per event type. Boldface indicates the best performance among each group (line 1-8, 9-13 and 14-18) without additional human efforts. † and gray background indicate using ACE05 human-curated data sampled, thus it is not comparable with other lines. Underlining indicate STAR-generated data improve EE performance more than human-curated data. There is no difference when using various  $k$  for lines indicated by § as no seed data instances are used. We use the trigger and event type produced by the best upstream model (*i.e.* Inst.+Examples) as inputs to Code4Struct to compensate for its lack of event detection capabilities.

to better performance when 5 or more demonstrations are used for STAR and the supervised model is DEGREE except for Arg-I (indicated by underlined results in line 17). We also observe a similar trend for OneIE when 10 demonstrations are used for STAR (underlined results in line 12). This indicates we could boost low-resource EE performance as if we had additional ground-truth data without paying the human annotation efforts. Figure 2 further shows the superiority of STAR-generated data over human-curated ones regardless of the number of augmented data instances ( $N$ ) using 10 demonstrations ( $k = 10$ ) for all supervised models.

**Noisy data generated by weak supervision impairs performance.** The performance of the best existing few-shot method used for weakly supervised data generation is still not satisfactory (line 7-8), leading to poor data quality and further impairing the performance of the supervised models by at least 10 points decline in F1 scores for Arg-C (line 9/14 vs 10/15). Shown by the downward curves in Figure 2, the performance decay is larger when more noisy data is used. This demonstrates that data generation methods that induce semi-supervised label  $Y'$  from  $X$  tend to perform poorly for complicated IE tasks like EE.

### 3.4 Best Recipe to Use LLM for EE

**Tunable supervised models are much better than inference-only methods.** Supervised models (without data augmentation) yield a 14.6-point higher F1 score than the best inference-only method (line 8 vs 14) for Arg-C when  $k = 10$ . The inference-only methods do not fully utilize the limited resources because they cannot learn from given demonstrations across all types. They suffer from the limitations of understanding the target structure from task description or inducing such patterns from in-context examples. We observe that supervised methods are better in handling low-resource EE, and their performances are further improved by training on additional data generated by STAR.

**When prompting LLM for EE, examples are crucial and code formulation helps.** Among various formulations for the inference-only methods, we observe that the natural language template (used in DEGREE) is more beneficial than the unnatural one (used in Text2Event) due to its better utilization of the pertaining knowledge. We also observe the importance of considering dependencies among arguments. This is evidenced by the superior performance of DEGREE




Figure 2: Event extraction performance (F1, %) when the EE models are trained on  $N$  augmented training data on top of 10 data points ( $k = 10$ ) for each event type. We observe that performance gain brought by STAR-generated data is magnified as the data augmentation scales up with a larger  $N$ , and data generated by STAR is even more effective than human-curated ones. We use GPT-3.5 version STAR for this set of experiments.

over DICE, which queries each argument role separately for argument extraction. The low scores of Text2Event, DEGREE, and DICE, which all use examples *without* task description, highlight the importance of describing requirements and rules in the instruction. It is worth noting that there is a substantial 45-point increase in F1 score from “Task Inst.” to “Inst.+Examples” when  $k = 10$  in the trigger classification task. This result highlights the effectiveness of the in-context examples in providing the LLM with the target structure knowledge. Additionally, we note that code formulation outperforms natural language, aligning with findings from previous studies (Wang, Li, and Ji 2023).

### 3.5 Effects of Demonstration Quantity $k$ and Data Quantity $N$

We observe that demonstration examples for the data generation pipeline are crucially important to bring in target structure information and showcase the task constraints in action. When no demonstration example is provided (*i.e.*  $k = 0$ ), the data generated by STAR could still significantly boost the supervised models’ performance, but it is much worse than training with human-curated data (line 12/17 vs 13/18 of Table 1). This indicates that even the state-of-the-art LLM GPT-4 is not capable of generating human-level EE data instances without demonstrations. The more demonstrations used in the data generation pipeline, the better the performance. Figure 2 shows that the performance gain brought by augmented data is magnified as the data augmentation scales up. Due to resource limitation, we use a maximum  $k$  of 10 and  $N$  of 50 in the experiments, and we anticipate the EE performance to further improve with larger  $k$  and/or  $N$ .

### 3.6 Ablation Studies

We investigate the influence of the design choices in Table 2.

**Structure generation method.** The diversity and balanced distribution of the generated target structures produced by our target structure generation component (§2.1) result in an almost 4-point higher Arg-C F1 score compared to using human-annotated target structures sampled from the ACE05 dataset. Our target structure generation component yields highly diverse structures  $Y$  without guaranteeing their factuality and commonsensical nature, suggesting that diversity outweighs factuality in terms of impact on performance.

#	Method Variant	Tri-I	Tri-C	Arg-I	Arg-C
Target Structure $Y$ Generation Methods					
1	Ground-truth $Y^\dagger$	59.21	54.03	44.13	41.77
2	LLM generation	60.52	54.80	48.00	45.73
Error Identification Strategies					
3	None	58.33	52.94	42.18	39.85
4	Rule-based checking	59.42	53.37	43.77	41.26
5	Self-reflection (NLI)	59.89	54.02	46.11	43.54
6	Self-reflection (LLM)	60.52	54.80	48.00	45.73

Table 2: Ablation study on DEGREE’s EE results while using 10 demonstrations with 10 additional generated data instances per event type ( $k = 10$ ,  $N = 10$ ).

**Error identification strategy.** We also investigate the error identification capabilities of our self-reflection module with LLM as the backbone (§2.3). We compare it with two alternative methods and we utilize the same template to provide feedback on the identified errors. **Rule-based checking** uses heuristics to check whether a trigger/argument is a subsequence of the generated passage and uses an external NER module (Yamada et al. 2020) to check whether a trigger/argument functions as the desired entity type. **Self-reflection (NLI)** uses the generated passage as the premier and a statement of a quality dimension as the hypothesis. We use entailment prediction of the NLI module used in §2.3 to identify whether a certain quality issue exists. The results are in lines 4-6 of Table 2. Our observations demonstrate both alternative methods help (line 3 vs 4-5), and self-reflection with LLM exhibits the highest effectiveness in error identification (line 6). Notably, the results underscore the effectiveness of the self-reflection design, resulting in a substantial 6-point increase in the F1 score for Arg-C without the need for additional annotation efforts.

## 4 Experiments on Relation Extraction

To assess the generalizability of our proposed method, we conduct experiments on the sentence-level relation extraction (RE) task. We use relation definitions and seed examples in the widely-used TACRED dataset (Zhang et al. 2017). In this task, we generate (subject, relation,

#	RE Model	Data Gen	$N = 0$	10	40
1	GPT-3.5	—	27.91	27.91	27.91
2	SURE	Weak Sup.	27.61	28.02	28.32
3		STAR (GPT-3.5)		<b>30.50</b>	<b>33.02</b>
4		Human <sup>†</sup>		30.11	35.62
5	GenPT	Weak Sup.	33.38	30.93	30.29
6		STAR (GPT-3.5)		<b>34.55</b>	<b>37.01</b>
7		Human <sup>†</sup>		36.74	37.61

Table 3: Relation extraction performance (%) when the RE models are trained on  $N$  augmented training data on top of 10 seed data instances ( $k = 10$ ) for each relation type.

object) tuples from scratch, providing additional training data. The RE task aims to identify the relation between the given subject and object entities within a context passage. We train two representative RE models on the generated data instances. **SURE** (Lu et al. 2022) converts the task into a summarization formulation to leverage the indirect supervision with PEGASUS-large as the pre-trained encoder (Zhang et al. 2020). **GenPT** (Han et al. 2022) transforms RE into an infilling problem with a RoBERTa-large model as backbone (Liu et al. 2019). We use the same set of data creation baselines as in the EE experiments. For the weakly supervised baseline, we use SURE’s formulation by querying the LLM once for each relation label, and the model outputs whether the verbalized relation between head and tail entities is entailed by the passage with  $k$  additional in-context examples, which is the most empirically superior RE formulation in our exploration. We report micro F1 score across all relations (except for the “no relation” class) following prior works (Lu et al. 2022, 2023).

Table 3 presents the RE performance. The STAR-generated data significantly enhances the performance compared to  $N = 0$ , with improvements of 5.4 and 3.6 F1 points when using SURE and GenPT respectively.

## 5 Quality Verification

Two annotators who are familiar with the EE task manually assess the quality of the EE data generated by STAR and curated by humans sampled from the ACE05 dataset in terms of 3 passage-related quality dimensions (grammatical, informative and commonsensical levels), and 5 task-specific quality dimensions (whether the trigger, event type, argument, and argument role provided in the generated target structure  $Y$  is being used in the generated passage  $X$  following their definitions) for data instances in 100 sentences. We report the average scores of the two annotators.

Table 4 shows both sets of data demonstrate high satisfactory levels. STAR-generated data exhibits higher passage quality and better follows the task definition for most metrics, suggesting STAR produces EE annotations with comparable or even better quality than human annotators.

Quality Dimension	STAR	Human
Grammaticality of the generated passage ( $X$ )	<b>96</b>	90
Informativeness of the generated passage ( $X$ )	<b>79</b>	78
Commonsense of the generated passage ( $X$ )	<b>95</b>	93
Trigger span describes event occurrence	<b>99</b>	<b>99</b>
Event follows event type definition	<b>99</b>	97
Argument span describes an event	<b>100</b>	99
Argument associated with correct trigger	<b>98</b>	95
Argument follows role definition	98	<b>99</b>

Table 4: Human quality assessment satisfactory rate (%).

## 6 Related Works

### 6.1 Low-Resource Information Extraction

Existing low-resource IE methods mostly borrow supervision signals from other tasks, such as Abstract Meaning Representation (Huang et al. 2018) or Semantic Role Labeling (Zhang, Wang, and Roth 2021), via transfer learning or reformulate the task as other data-rich tasks, such as Natural Language Inference (NLI) (Xu, Ma, and Chen 2023; Sainz et al. 2022), summarization (Lu et al. 2022) or QA (Lyu et al. 2021; Ma et al. 2023a), via indirect supervision. Recent works explore prompting LLM with task instruction and examples to predict event structure from input passage by prompting task requirement (Li et al. 2023), or providing examples along with task instruction (Gao et al. 2023; Xu et al. 2023). However, their performance is much worse than fine-tuning a model with training data for complicated IE tasks like EE, which motivates us to focus on generating training data for supervised IE models.

### 6.2 Data Generation

Existing works perform data augmentation by producing analogous input ( $X_{gold} \rightarrow X'$ ) (Kumar, Choudhary, and Cho 2020; Lee et al. 2021), and annotating unlabeled examples with existing models ( $X_{gold} \rightarrow Y$ ) (He et al. 2021) on tasks like textual similarity (Schick and Schütze 2021), relation extraction (Chia et al. 2022), text classification, QA, NLI (Ye et al. 2022) and more (Wang et al. 2023; Tang et al. 2023). Another line of prior work assumes that the labels are given and it prompts an LLM to generate input ( $Y_{gold} \rightarrow X$ ), such as given sentiment labels (Meng et al. 2022; Gao, Fisch, and Chen 2021) or relation triples (Josifoski et al. 2023). However, these works generate data for tasks with significantly simpler output structures than many IE tasks like EE. Furthermore, they require existing ground-truth label data to perform data generation while our work aims to generate both  $Y$  and  $X$  from scratch.

## 7 Conclusion

We present STAR, an inverse data generation pipeline designed for low-resource IE that generates output structure first and then curates input passage containing structure content with self-refinement capabilities to fix self-identified error cases. Experimental results on EE and RE show the generated data instances significantly improve the performance.



## Acknowledgments

Utkarsh Lal and Michael M. Song made significant contributions to the quality verification section. Many thanks to lab members at UCLA SCAI and UCLANLP for their suggestions, and to the anonymous reviewers for their feedback. This effort was sponsored by the Defense Advanced Research Project Agency (DARPA) grant HR00112290103 and AFOSR MURI via grant #FA9550-22-1-0380.

## References

- Cao, Y.; Pruksachatkun, Y.; Chang, K.-W.; Gupta, R.; Kumar, V.; Dhamala, J.; and Galstyan, A. 2022. On the Intrinsic and Extrinsic Fairness Evaluation Metrics for Contextualized Language Representations. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*.
- Chia, Y. K.; Bing, L.; Poria, S.; and Si, L. 2022. Relation-Prompt: Leveraging Prompts to Generate Synthetic Data for Zero-Shot Relation Triplet Extraction. In *Findings of the Association for Computational Linguistics: ACL 2022*.
- Doddington, G.; Mitchell, A.; Przybicki, M.; Ramshaw, L.; Strassel, S.; and Weischedel, R. 2004. The Automatic Content Extraction (ACE) Program – Tasks, Data, and Evaluation. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC’04)*.
- Gao, J.; Zhao, H.; Yu, C.; and Xu, R. 2023. Exploring the Feasibility of ChatGPT for Event Extraction. arxiv:2303.03836.
- Gao, T.; Fisch, A.; and Chen, D. 2021. Making Pre-trained Language Models Better Few-shot Learners. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*.
- Han, J.; Zhao, S.; Cheng, B.; Ma, S.; and Lu, W. 2022. Generative Prompt Tuning for Relation Classification. *arXiv preprint arXiv:2210.12435*.
- Han, R.; Peng, T.; Yang, C.; Wang, B.; Liu, L.; and Wan, X. 2023. Is Information Extraction Solved by ChatGPT? An Analysis of Performance, Evaluation Criteria, Robustness and Errors. arxiv:2305.14450.
- He, X.; Nassar, I.; Kiros, J. R.; Haffari, G.; and Norouzi, M. 2021. Generate, Annotate, and Learn: Generative Models Advance Self-Training and Knowledge Distillation.
- Hsu, I.-H.; Huang, K.-H.; Boschee, E.; Miller, S.; Nataraajan, P.; Chang, K.-W.; and Peng, N. 2022. DEGREE: A Data-Efficient Generation-Based Event Extraction Model. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Huang, L.; Ji, H.; Cho, K.; Dagan, I.; Riedel, S.; and Voss, C. 2018. Zero-Shot Transfer Learning for Event Extraction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Josifoski, M.; Sakota, M.; Peyrard, M.; and West, R. 2023. Exploiting Asymmetry for Synthetic Training Data Generation: SynthIE and the Case of Information Extraction. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.
- Kumar, V.; Choudhary, A.; and Cho, E. 2020. Data Augmentation Using Pre-trained Transformer Models. In *Proceedings of the 2nd Workshop on Life-long Learning for Spoken Language Systems*.
- Lee, K.; Guu, K.; He, L.; Dozat, T.; and Chung, H. W. 2021. Neural data augmentation via example extrapolation. *arXiv preprint 2102.01335*.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 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*.
- Li, B.; Fang, G.; Yang, Y.; Wang, Q.; Ye, W.; Zhao, W.; and Zhang, S. 2023. Evaluating ChatGPT’s Information Extraction Capabilities: An Assessment of Performance, Explainability, Calibration, and Faithfulness. arxiv:2304.11633.
- Li, S.; Ji, H.; and Han, J. 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*.
- Lin, Y.; Ji, H.; Huang, F.; and Wu, L. 2020. A Joint Neural Model for Information Extraction with Global Features. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Liu, H.; Sferrazza, C.; and Abbeel, P. 2023. Chain of Hindsight Aligns Language Models with Feedback. arxiv:2302.02676.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. arxiv:1907.11692.
- Lu, K.; Hsu, I.-H.; Zhou, W.; Ma, M. D.; and Chen, M. 2022. Summarization as Indirect Supervision for Relation Extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2022*.
- Lu, K.; Hsu, I.-H.; Zhou, W.; Ma, M. D.; and Chen, M. 2023. Multi-hop Evidence Retrieval for Cross-document Relation Extraction. In *Findings of the Association for Computational Linguistics: ACL 2023*, 10336–10351.
- Lu, Y.; Lin, H.; Xu, J.; Han, X.; Tang, J.; Li, A.; Sun, L.; Liao, M.; and Chen, S. 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 (Volume 1: Long Papers)*.
- Lyu, Q.; Zhang, H.; Sulem, E.; and Roth, D. 2021. Zero-Shot Event Extraction via Transfer Learning: Challenges and Insights. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*.



- Ma, M. D.; Chen, M.; Wu, T.-L.; and Peng, N. 2021a. HyperExpan: Taxonomy Expansion with Hyperbolic Representation Learning. In *Findings of the Association for Computational Linguistics: EMNLP 2021*.
- Ma, M. D.; Kao, J.-Y.; Gao, S.; Gupta, A.; Jin, D.; Chung, T.; and Peng, N. 2023a. Parameter-Efficient Low-Resource Dialogue State Tracking by Prompt Tuning. In *Proc. InterSpeech 2023*.
- Ma, M. D.; Kao, J.-Y.; Gupta, A.; Lin, Y.-H.; Zhao, W.; Chung, T.; Wang, W.; Chang, K.-W.; and Peng, N. 2023b. Mitigating Bias for Question Answering Models by Tracking Bias Influence. arxiv:2310.08795.
- Ma, M. D.; Sun, J.; Yang, M.; Huang, K.-H.; Wen, N.; Singh, S.; Han, R.; and Peng, N. 2021b. EventPlus: A Temporal Event Understanding Pipeline. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations*.
- Ma, M. D.; Taylor, A.; Wang, W.; and Peng, N. 2023c. DICE: Data-Efficient Clinical Event Extraction with Generative Models. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Madaan, A.; Tandon, N.; Gupta, P.; Hallinan, S.; Gao, L.; Wiegrefe, S.; Alon, U.; Dziri, N.; Prabhunoye, S.; Yang, Y.; Welleck, S.; Majumder, B. P.; Gupta, S.; Yazdanbakhsh, A.; and Clark, P. 2023. Self-Refine: Iterative Refinement with Self-Feedback. arxiv:2303.17651.
- Meng, Y.; Huang, J.; Zhang, Y.; and Han, J. 2022. Generating Training Data with Language Models: Towards Zero-Shot Language Understanding. arxiv:2202.04538.
- OpenAI. 2022. Chatgpt: Optimizing language models for dialogue. <https://openai.com/blog/chatgpt/>.
- OpenAI. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Sainz, O.; Gonzalez-Dios, I.; Lopez de Lacalle, O.; Min, B.; and Agirre, E. 2022. Textual Entailment for Event Argument Extraction: Zero- and Few-Shot with Multi-Source Learning. In *Findings of the Association for Computational Linguistics: NAACL 2022*.
- Schick, T.; and Schütze, H. 2021. Generating Datasets with Pretrained Language Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Tang, R.; Han, X.; Jiang, X.; and Hu, X. 2023. Does Synthetic Data Generation of LLMs Help Clinical Text Mining? arxiv:2303.04360.
- Tong, M.; Xu, B.; Wang, S.; Han, M.; Cao, Y.; Zhu, J.; Chen, S.; Hou, L.; and Li, J. 2022. DocEE: A Large-Scale and Fine-grained Benchmark for Document-level Event Extraction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Wang, X.; Li, S.; and Ji, H. 2023. Code4Struct: Code Generation for Few-Shot Event Structure Prediction. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Wang, Y.; Kordi, Y.; Mishra, S.; Liu, A.; Smith, N. A.; Khashabi, D.; and Hajishirzi, H. 2023. Self-Instruct: Aligning Language Models with Self-Generated Instructions. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Williams, A.; Nangia, N.; and Bowman, S. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*.
- Xu, J.; Ma, M. D.; and Chen, M. 2023. Can NLI Provide Proper Indirect Supervision for Low-resource Biomedical Relation Extraction? In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Xu, J.; Ma, M. D.; Wang, F.; Xiao, C.; and Chen, M. 2023. Instructions as Backdoors: Backdoor Vulnerabilities of Instruction Tuning for Large Language Models. arxiv:2305.14710.
- Yamada, I.; Asai, A.; Shindo, H.; Takeda, H.; and Matsumoto, Y. 2020. LUKE: Deep Contextualized Entity Representations with Entity-aware Self-attention. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Ye, J.; Gao, J.; Li, Q.; Xu, H.; Feng, J.; Wu, Z.; Yu, T.; and Kong, L. 2022. ZeroGen: Efficient Zero-shot Learning via Dataset Generation. arxiv:2202.07922.
- Zhang, H.; Wang, H.; and Roth, D. 2021. Zero-Shot Label-Aware Event Trigger and Argument Classification. In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021*.
- Zhang, J.; Zhao, Y.; Saleh, M.; and Liu, P. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, 11328–11339. PMLR.
- Zhang, Y.; Zhong, V.; Chen, D.; Angeli, G.; and Manning, C. D. 2017. Position-Aware Attention and Supervised Data Improve Slot Filling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*.
- Zhao, J.; Wang, T.; Yatskar, M.; Ordonez, V.; and Chang, K.-W. 2018. Gender Bias in Coreference Resolution: Evaluation and Debiasing Methods. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*.
- Zhou, Y.; Kaneko, M.; and Bollegala, D. 2022. Sense Embeddings are also Biased – Evaluating Social Biases in Static and Contextualised Sense Embeddings. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.