

IELM: An Open Information Extraction Benchmark for Pre-Trained Language Models

Chenguang Wang¹, Xiao Liu², Dawn Song³

¹Washington University in St. Louis, ²Tsinghua University, ³UC Berkeley
chenguangwang@wustl.edu, liuxiao21@mails.tsinghua.edu.cn,
dawnsong@berkeley.edu

Abstract

We introduce a new open information extraction (OIE) benchmark for pre-trained language models (LM). Recent studies have demonstrated that pre-trained LMs, such as BERT and GPT, may store linguistic and relational knowledge. In particular, LMs are able to answer “fill-in-the-blank” questions when given a pre-defined relation category. Instead of focusing on pre-defined relations, we create an OIE benchmark aiming to fully examine the open relational information present in the pre-trained LMs. We accomplish this by turning pre-trained LMs into zero-shot OIE systems. Surprisingly, pre-trained LMs are able to obtain competitive performance on both standard OIE datasets (CaRB and Re-OIE2016) and two new large-scale factual OIE datasets (TAC KBP-OIE and Wikidata-OIE) that we establish via distant supervision. For instance, the zero-shot pre-trained LMs outperform the F1 score of the state-of-the-art supervised OIE methods on our factual OIE datasets without needing to use any training sets.¹

1 Introduction

Pre-trained language models (LM), such as BERT (Devlin et al., 2018) and GPT-3 (Brown et al., 2020), have revolutionized NLP over the last several years and advanced the state-of-the-art results in a wide set of downstream NLP tasks. Recent studies show that a considerable amount of linguistic (Hewitt and Manning, 2019; Clark et al., 2019) and relational knowledge (Petroni et al., 2019; Talmor et al., 2019; Jiang et al., 2020; Petroni et al., 2020) has been captured by the pre-trained LMs via pre-training on large-scale textual corpora. These approaches often design “fill-in-the-blank” questions based on pre-defined relations. For example, a question “Bob Dylan was born in

__” is manually created for the LMs to answer the “birthplace” relation of “Bob Dylan”.

Most existing approaches that evaluate what pre-trained LMs have learned are based on benchmarks with pre-defined relation categories. Yet, the benchmarks present two limitations. First, most benchmarks only cover a limited number of pre-defined relations. Therefore, it is unclear whether the pre-trained LMs have stored general open relation information. For example, the Google-RE in LAMA benchmark (Petroni et al., 2019) includes only three relations (i.e., “birthplace”, “birthdate”, and “death-place”), while there are hundreds of relations available in the real world scenario. Second, a majority of benchmarks evaluate LMs in a close manner. This means that the gold relation is given to the models. For example, “was born in” is given as the model’s input. Besides, the existing benchmarks often provide a single gold relation for each input sentence. However, an input sentence may indicate multiple relations, e.g., containing both “birthplace” and “birthdate” information about an argument or entity. We are curious: instead of the limited relational information, can we systematically benchmark the general information stored in the pre-trained LMs?

In this work, we set up a new open information extraction (OIE) benchmark, called IELM, towards testing the general and open relational information stored in pre-trained LMs. We refer to OIE as it is a task that is designed to extract open relations from massive corpora without requiring a pre-defined relation category. As shown in Figure 1, we successfully convert pre-trained LMs to zero-shot OIE systems. We apply them to two standard OIE datasets, including CaRB (Bhardwaj et al., 2019) and Re-OIE2016 (Stanovsky and Dagan, 2016; Zhan and Zhao, 2020), as well as two new large-scale factual OIE datasets in our IELM benchmark. We show that the zero-shot pre-trained LMs outperform the fully supervised state-of-the-arts on fac-

¹Our code and datasets are available at <https://github.com/cgraywang/IELM>.

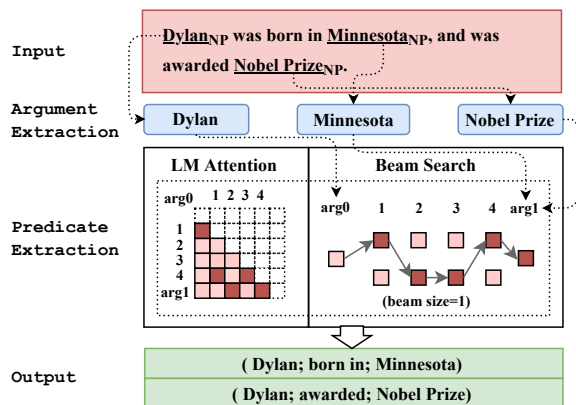


Figure 1: Summary of our approach. The zero-shot open information extraction system takes a noun phrase (NP) chunked sentence as input, and outputs a set of triples. The approach first conducts argument extraction by encoding NPs as argument pairs, then performs predicate extraction via decoding using the parameters (i.e., attention scores) of the pre-trained language models. The output extractions are then evaluated on our IELM benchmark.

tual OIE datasets. Standard OIE datasets rely on human annotations and often consist of thousands of gold triples and sentences. Unlike those datasets, we create two large-scale OIE datasets, namely TAC KBP-OIE and Wikidata-OIE, via distant supervision from knowledge graphs. For example, Wikidata-OIE is constructed via aligning English Wikipedia to Wikidata triples, resulting in millions of triples and documents. The design of zero-shot LMs for OIE is important: by encoding the noun chunks as arguments in the input, we only make use of the parameters of pre-trained LMs to decode the predicates (or relations) between the arguments. To the best of our knowledge, this is the first attempt to systematically evaluate pre-trained LMs in a zero-shot OIE setting. To summarize, our key contributions are the following.

1. We benchmark the general relational information in pre-trained LMs on our IELM benchmark. Besides two standard OIE datasets (CaRB and Re-OIE2016), we also create two large-scale factual OIE datasets for our benchmark. The two new OIE datasets are called TAC KBP-OIE and Wikidata-OIE, which are constructed via distant supervision from two knowledge graphs (TAC KBP and Wikidata). Our benchmark is a general OIE benchmark, helping the development of future OIE systems.
2. We enable the zero-shot capabilities of pre-trained LMs for OIE by encoding the arguments

in the input and decoding predicates using the parameters of pre-trained LMs. The pre-trained LMs are particularly good at recovering factual arguments and predicates.

3. We test the OIE performance of 6 pre-trained LMs (BERT and GPT-2 (Radford et al., 2019) families) and 14 OIE systems on IELM benchmark. The zero-shot LMs achieve state-of-the-art OIE performance on TAC KBP-OIE and Wikidata-OIE, even outperforming fully supervised OIE systems.

2 Language Models as Zero-Shot Information Extractors

For open information extraction (OIE), we take an input as a NP-chunked sentence and output a set of triples. Below is an example.

Input Dylan_{NP} was born in Minnesota_{NP}, and was awarded Nobel Prize_{NP}.

Output (Dylan; born in; Minnesota), (Dylan; awarded; Nobel Prize).

NP denotes the noun phrase.

2.1 Argument Extraction

Follow traditional linguistic OIE systems such as Stanford OpenIE (Angeli et al., 2015) and OpenIE5 (Saha et al., 2017, 2018), we treat each NP pair as an argument pair (e.g., “Dylan” and “Minnesota”). We then utilize the parameters of LMs to extract the predicates (e.g., “born in”) between the pair in the input as below.

2.2 Predicate Extraction

The predicate extraction problem is formulated as extracting a set of sequences in the input that are associated with an argument pair. We particularly use the attention scores in a pre-trained LM to measure the relevance between a sequence and the argument pair. We frame the process as a search problem. Given an argument pair, we aim to search for the sequences with the largest attention scores connecting the pair. To compute a score for every possible sequence is computationally expensive especially when the sequence length is large, the exhaustive search is therefore intractable. We adapt beam search as an approximate strategy to efficiently explore the search space. Beam search maintains the k -best candidates. This means the time cost of beam search does not depend on the sequence length, but on the size of the beam and

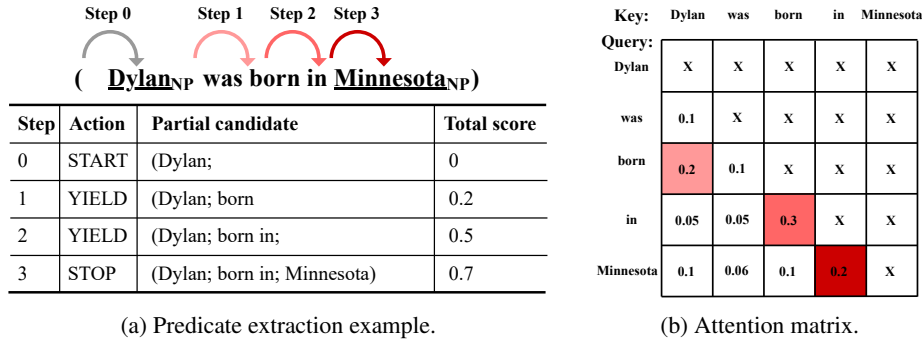


Figure 2: Illustration of predicate extraction with a pre-trained language model (LM). The upper part of (a) represents the general search steps of producing the triple $(Dylan; born\ in; Minnesota)$ from the input “ $Dylan_{NP}$ was born in $Minnesota_{NP}$ ” encoded with argument noun phrases (NP). The lower portion shows the corresponding step-by-step process. (b) shows the attention scores generated through the forward pass of the LM over the corresponding input.

the average length of the candidates. In general, the beam search starts with the first argument (e.g., “Dylan”). At each step, beam search simply selects top- k next tokens with the largest attention scores, and just keeps k partial candidates with the highest scores, where k is the beam size. When a candidate produces the second argument (e.g., “Minnesota”), the candidate is complete.

We show a running example as follows. Let’s first consider the search from left to right with beam size equal to 1. An example search process is shown in Figure 2. Given an argument pair “Dylan” and “Minnesota”, at each step, the search performs one of the following actions:

- **START** the search from first argument. The first argument is added as an initial candidate into the beam. In Figure 2(a), at step 0, “Dylan” is added into the beam. The total attention score is initialized to 0.
- **YIELD** a new partial candidate in the beam if the current candidate has not reached the second argument. This action conducts the following: The next largest attended token is appended to the end of the current candidate to yield the new candidate. The total score is increased by the associated attention score. At step 1 of Figure 2(a), “born” is appended to the current candidate to yield the partial candidate, since “born” has the largest attention score (0.2 as highlighted in Figure 2(b)) with “Dylan” in the attention matrix. The total score becomes 0.2. Note that we only consider the single head attention in this example for simplicity. “x” in Figure 2(b) marks the tokens (prior to the current token) that are not considered in the search to prevent searching

backward. Step 2 takes the same action, and the score becomes 0.5.

- **STOP** the search step if the candidate has reached the second argument, then add the candidate as a valid triple into the beam. As beam size equals to 1, $(Dylan; born\ in; Minnesota)$ is returned for the given pair. The final score of the triple is 0.7.

We also notice triples are often in reverse order in the sentence, thus enabling bidirectionality by running the algorithm in both directions (left to right and right to left). We merge the subwords as words, and only consider word-level attention. The beam search is implemented by the breadth-first search, which is efficient as the time complexity is $O(k \cdot d)$. d is the maximum depth of the search tree.

3 The IELM Benchmark

3.1 Datasets

3.1.1 Standard OIE

We adopt two standard OIE datasets below.

CaRB CaRB (Bhardwaj et al., 2019) is a crowd-sourced OIE dataset, where the input sentences are from the OIE2016 (Stanovsky and Dagan, 2016).

Re-OIE2016 Re-OIE2016 (Zhan and Zhao, 2020) is also generated based on the input sentences in the OIE2016, and is further enhanced by human annotations.

3.1.2 Factual OIE

In addition, we introduce two large-scale factual OIE datasets based on knowledge graphs (KG).

Method	AIDA	
	dev	test
Spitkovsky and Chang 2012	26.0	28.2
Kolitsas et al. 2018*	-	82.4
Ours	63.8	64.5

Table 1: Evaluation of unsupervised entity linking of Wikidata-OIE on AIDA benchmark. An asterisk (*) indicates a supervised method.

TAC KBP-OIE TAC Knowledge Base Population (KBP) Slot Filling is a task to search a document collection to fill in a target entity for predefined relations (slots) with a given entity in a reference KG. We adapt the dataset as an OIE dataset. In particular, we use a document sub-collection of the TAC KBP 2013 task (Surdeanu, 2013) as the input, and use the official human annotations regarding the documents as gold extractions.

Wikidata-OIE Besides TAC KBP-OIE, we create a larger factual OIE dataset based on the English Wikipedia. Different from TAC KBP, there are no gold triple annotations for Wikipedia. Since a large amount of Wikidata triples are from English Wikipedia, we create the dataset using distant supervision (Zhang et al., 2017) by aligning Wikidata triples to Wikipedia text. We employ an unsupervised entity linker based on a pre-built mention-to-entity dictionary (Spitkovsky and Chang, 2012) to extract potential gold arguments for scalability considerations. The entity linker links an arbitrary entity mention in a sentence to a Wikipedia anchor, which is further linked to a Wikidata entity. For each sentence in Wikipedia articles containing two linked arguments, if there is a Wikidata triple describing a relation holding the two arguments, we denote the Wikidata triple as a gold extraction.

Unlike TAC KBP-OIE which is built based on human annotations, Wikidata-OIE is derived from automatic annotations. Therefore, we evaluate our unsupervised entity linker on the standard AIDA benchmark (Hoffart et al., 2011) consisting of Wikipedia entities. Table 1 shows that it significantly improves the unsupervised performance (Spitkovsky and Chang, 2012) and reaches competitiveness with a supervised method (Kolitsas et al., 2018). Given the scale of Wikidata-OIE, we sacrifice acceptable effectiveness for efficiency.

The statistics of the datasets are shown in Table 2. For CaRB and Re-OIE2016, we report the statistics of the corresponding test sets. We include a dataset

Dataset	# of triples	# of arguments	# of predicates	# of documents
Re-OIE2016	1,508	3,328	1,506	595
CaRB	2,715	6,226	2,715	641
TAC KBP-OIE	27,655	39,661	41	3,877,207
Wikidata-OIE	27,368,562	6,047,494	1,156	6,047,494

Table 2: Dataset statistics of the IELM benchmark.

comparison in Appendix A.3.

3.2 Pre-Trained Language Models for OIE

Unidirectional Language Models Given an input sequence $\mathbf{x} = \{x_1, x_2, \dots, x_N\}$, unidirectional LMs assign a joint probability to the sequence by factorizing it as $p(\mathbf{x}) = \prod_t p(x_t | x_{t-1}, \dots, x_1)$, where $p(x_t | x_{t-1}, \dots, x_1) = \sigma(\mathbf{W}\mathbf{h}_t + \mathbf{b})$. \mathbf{h}_t is the output vector of a neural network at position t .

We consider GPT-2 (Radford et al., 2019), where \mathbf{h}_t is produced by Transformer decoders (Vaswani et al., 2017). GPT-2 is pre-trained on WebText containing 40GB of text. We explore all four pre-trained GPT-2s with different model sizes: GPT-2 (117M), GPT-2_{MEDIUM} (345M), GPT-2_{LARGE} (774M), and GPT-2_{XL} (1558M).

Bidirectional Language Models Different from unidirectional LMs that predict the next word given the previous words, bidirectional LMs take both left and right context of the target word into consideration, formally, $p(x_t) = p(x_t | x_1, \dots, x_{t-1}, x_{t+1}, \dots, x_N)$.

We use BERT (Devlin et al., 2018) that enables bidirectional context modeling via a masked LM objective and utilizing the Transformer architecture. BERT is pre-trained on BooksCorpus and English Wikipedia. We use both its pre-trained settings: BERT_{BASE} (109M) and BERT_{LARGE} (335M).

3.3 Comparison Methods

We compare our method with a wide set of OIE systems including both neural and traditional linguistic OIE systems. Most OIE systems are based on supervised learning, which are indicated with asterisks (*) in Table 3. We provide details of the comparison systems in Appendix A.5.

3.4 Evaluation Method

3.4.1 Standard OIE

On CaRB and Re-OIE2016, we follow the original evaluation proposed in (Bhardwaj et al., 2019) and (Stanovsky and Dagan, 2016; Zhan and Zhao, 2020), and report precision, recall, F1, area under the curve (AUC) for compared OIE systems. AUC

is calculated from a plot of the precision and recall values for all potential confidence thresholds. The F1 is the maximum value among the precision-recall pairs. We follow the matching function proposed for each dataset, i.e., lexical match for Re-OIE2016, and tuple match for CaRB. The CaRB evaluation function is stricter as it penalizes long extractions.

3.4.2 Factual OIE

We report precision, recall, and F1 of the OIE systems on two large-scale factual OIE datasets: TAC KBP-OIE and Wikidata-OIE. We introduce *exact match* as the matching function for them as below.

Matching Function The matching functions for standard OIE datasets are generally flexible. For example, the lexical match of Re-OIE2016 judges an argument or predicate as correct if and only if it includes the syntactic head of the gold argument or predicate. Unlike these matching functions, our exact matching function requires both arguments and predicates are linked to the gold extractions.

For TAC KBP-OIE, we judge an argument to be correct if and only if it matches the name of the gold argument and the span position of the gold argument in the sentence. The main challenge is how to properly link a predicate, since there are often many ways to express it. We follow Stanford OpenIE (Angeli et al., 2015) to produce the predicate mapping between the OIE relations and TAC KBP predicates. A predicate is correct if the pair of OIE relation and gold predicate exists in the predicate mapping. The predicate mapping is constructed in two steps. First, a collection of predicate mappings was constructed by a single annotator in approximately a day. Second, predicate mappings were finalized through the following learning procedure. This process matches OIE relations to the TAC KBP predicates by searching for co-occurrent relations in a large distantly-labeled corpus, and decides pairs of OIE relations and TAC KBP predicates that have a high PMI². The basic idea is that the more often the argument pairs of the triples and TAC KBP triples are linked, the more likely the corresponding relations or predicates are linked to each other. Example predicate mappings are shown in Appendix A.4.

For Wikidata-OIE, we link an argument based on the entity linker used in Wikidata-OIE construction (Sec. 3.1). An argument is correct if the linked argument matches the name of the gold argument and

the span position of the gold argument in the sentence. The predicate mapping is bootstrapped from TAC KBP-OIE’s mapping. In addition, we normalize each predicate phrase of the triples by lemmatization, and removing inflection, auxiliary verbs, adjectives, adverbs. One author manually filters out the bad predicate pairs. This process takes approximately a day. A predicate is correct if the OIE to gold predicate pair exists in the bootstrapped predicate mapping. An annotator randomly subsampled and checked 100 aligned triple-sentence pairs and concluded with a 93% accuracy of extracted triples.

Metrics We use the official scorer of TAC KBP Slot Filling 2013 to calculate precision, recall, and F1 for TAC KBP-OIE. Besides, like previous OIE systems, for LMs, we adopt two constraints from ReVerb (Fader et al., 2011), the antecessor of OpenIE5: (i) the frequency of predicates must be above a threshold aiming to avoid triples to be over-specified, and (ii) a predicate must be a contiguous sequence in the sentence avoiding predicates that have no meaningful interpretation. We set these parameters according to Sec. 4.4.

We only report precision, recall, and F1 based on the parameter study in Sec. 4.4 for pre-trained LMs on the IELM benchmark. We do not compute AUC as pre-trained LMs are treated as zero-shot OIE systems. We therefore do not tune the results with respect to different confidence. Our main focus is to benchmark the OIE performance of the LMs under a unified setting. Another reason is that it is computationally expensive to get the AUC on the two large-scale datasets: TAC KBP-OIE and Wikidata-OIE. We also do not report AUC for the compared OIE systems on TAC KBP-OIE and Wikidata-OIE. Instead, we use the confidence threshold that obtains the best F1 on Re-OIE2016 to compute the scores.

4 Results

In this section, we show that pre-trained LMs are effective zero-shot OIE systems, and exceed the previous state-of-the-art OIE systems on our large-scale factual OIE datasets in IELM benchmark. To keep our evaluation as simple as possible, the hyper-parameters and settings are shared across datasets. More experimental details are described in the appendix.

Method	#Params	CaRB				Re-OIE2016				TAC KBP-OIE			Wikidata-OIE		
		P	R	F1	AUC	P	R	F1	AUC	P	R	F1	P	R	F1
MinIE (Gashteovski et al., 2017)	-	-	-	41.9	-	48.2	71.1	57.5	47.2	-	-	-	-	-	-
ClausIE (Del Corro and Gemulla, 2013)	-	-	-	44.9	22.4	-	-	64.2	46.4	-	-	-	-	-	-
OLLIE (Schmitz et al., 2012)	-	-	-	41.1	22.4	-	-	49.5	31.3	-	-	-	-	-	-
PropS (Stanovsky et al., 2016)	-	-	-	31.9	12.6	-	-	64.2	43.3	-	-	-	-	-	-
OpenIE4 (Christensen et al., 2011)	-	-	-	48.8	27.2	-	-	68.3	50.9	52.9	14.2	22.4	30.6	14.0	19.2
OpenIE5 (Saha et al., 2017, 2018)	-	-	-	48.0	25.0	61.1	76.5	67.9	45.8	57.0	14.6	23.2	26.3	14.4	18.6
Stanford OpenIE* (Angeli et al., 2015)	-	-	-	23.0	13.4	-	-	16.7	11.5	61.6	17.4	27.1	23.3	13.1	16.8
SenseOIE* (Roy et al., 2019)	640K	-	-	28.2	-	-	-	-	-	-	-	-	-	-	-
SpanOIE* (Zhan and Zhao, 2020)	963K	60.9	41.6	49.4	30.0	79.7	74.5	77.0	65.8	-	-	-	-	-	-
NeuralOIE* (Cui et al., 2018)	5M	-	-	51.6	32.8	79.2	77.5	78.4	73.0	-	-	-	-	-	-
Multi ² OIE* (Ro et al., 2020)	110M	60.9	45.8	52.3	32.6	86.9	81.0	83.9	74.6	-	-	-	-	-	-
RnnOIE* (Stanovsky et al., 2018)	965K	55.6	40.2	46.7	26.8	84.2	73.9	78.7	68.3	50.0	14.6	22.6	29.9	15.9	20.7
IMOJIE* (Kolluru et al., 2020b)	110M	64.7	45.6	53.5	33.3	88.1	67.1	76.2	63.1	58.2	14.9	23.8	31.2	16.2	21.3
OpenIE6* (Kolluru et al., 2020a)	220M	-	-	54.0	35.7	75.3	78.2	76.7	73.8	60.0	14.9	23.9	29.8	15.3	20.3
BERT _{BASE} (zero-shot) (ours)	109M	21.2	18.3	19.7	-	25.9	34.0	29.4	-	61.6	18.8	28.8	32.0	18.9	23.7
BERT _{LARGE} (zero-shot) (ours)	335M	22.4	20.2	21.2	-	30.7	38.5	34.1	-	61.7	19.0	29.1	32.3	19.1	24.0
GPT-2 (zero-shot) (ours)	117M	23.1	19.8	21.3	-	25.1	38.3	30.3	-	61.6	18.2	28.1	32.4	18.1	23.2
GPT-2 _{MEDIUM} (zero-shot) (ours)	345M	23.7	20.0	21.7	-	26.8	39.9	32.0	-	62.1	18.7	28.7	33.1	18.2	23.5
GPT-2 _{LARGE} (zero-shot) (ours)	774M	24.2	20.5	22.2	-	27.4	41.6	33.0	-	62.4	19.0	29.1	34.2	18.3	23.8
GPT-2 _{XL} (zero-shot) (ours)	1558M	24.5	21.0	22.7	-	29.3	43.4	35.0	-	62.7	19.5	29.7	35.7	18.5	24.4

Table 3: Compare the quality of different OIE systems. An asterisk (*) indicates a supervised method.

4.1 OIE Results

Table 3 shows the results. While zero-shot OIE systems synthesized by pre-trained LMs obtain notably lower scores compared to previous OIE systems on standard OIE datasets, they outperform the previous OIE systems on factual OIE datasets. We also find that larger LMs obtain improved results on all datasets. For example, BERT_{LARGE} outperforms BERT_{BASE} due to its larger model size. GPT-2s share similar trends. This is because larger LMs store richer relational information. This finding is consistent with previous studies (Petroni et al., 2019, 2020).

4.1.1 Standard OIE

The main reasons for the degraded performance of pre-trained LMs on standard OIE datasets are three-fold. First, the comparison methods mainly involve supervised systems that are trained on OIE datasets, which are denoted with asterisks (*) in Table 3. Besides the supervised systems, the remaining comparison systems all require human involvement, such as providing linguistic patterns for the extraction. In contrast, the pre-trained LMs are used as zero-shot OIE systems without using any training sets. Second, the zero-shot OIE modules still have room to improve. For example, approximately 30.0% of the argument extraction errors are due to the spaCy noun chunker. 16.9% of the gold extractions contain predicates outside the argument pairs. The current predicate extraction only allows searching between the arguments, and thus cannot handle such cases. Third, standard OIE benchmarks

such as CaRB and Re-OIE2016 mainly examine the general information extraction capability. The zero-shot approach is not able to recall the information of interest in LMs. We might need a specific module (e.g., ranking) to locate such information. Interestingly, pre-trained LMs achieve comparable performance with supervised Stanford OpenIE. The result indicates pre-trained LMs contain informative patterns that are useful for OIE.

4.1.2 Factual OIE

As shown in Table 3, the best zero-shot OIE system based on GPT-2_{XL} obtains a +2.6% and a +3.1% absolute F1 improvement on TAC KBP-OIE and Wikidata-OIE respectively over the previous supervised state-of-the-art. Due to the computation cost of OIE systems (Sec. 4.3), we only select several best performed OIE systems on the standard OIE datasets for the large-scale OIE experiments including: linguistic OIE systems (OpenIE4, OpenIE5, Stanford OpenIE) and neural OIE systems (RnnOIE, IMOJIE, OpenIE6).

Compared to the results on standard OIE datasets, pre-trained LMs consistently achieve state-of-the-art performance on both datasets. Both datasets emphasize measuring factual arguments and predicates in the reference KGs. Previous studies (Petroni et al., 2019, 2020) show that LMs have stored a considerable amount of factual information via pre-training on large-scale text. We draw the same conclusion. To the best of our knowledge, our IELM benchmark is the first benchmark that includes factual OIE datasets. More importantly, both

Method	MinIE	ClausIE	OLLIE	PropS	OpenIE 4	OpenIE 5	Stanford OpenIE	SpanOIE	RnnOIE	
Sentences/sec.	8.9	4.0	14.5	4.6	20.1	3.1	2.5	19.4	149.2	
Method	NeuralOIE	IMOJIE	Multi ² OIE	OpenIE 6	BERT _{BASE}	BERT _{LARGE}	GPT-2	GPT-2 _{MEDIUM}	GPT-2 _{LARGE}	GPT-2 _{XL}
Sentences/sec.	11.5	2.6	21.2	142.0	16.2	11.9	13.9	12.7	11.5	11.0

Table 4: Runtime on Re-OIE2016.

linguistic and neural OIE systems are derived from manually designed linguistic patterns or learned patterns. The result shows that the pre-trained attention weights capture a more flexible set of factual patterns. The result also suggests that our approach is capable of using such patterns. In order to scale our approach to large-scale datasets, the argument and predicate extraction are both efficient by design. In particular, the beam search for predicate extraction is efficient in exploring the relational sequences in the input sentence. Besides, the attention scores used in the beam search are produced via a single forward pass of the pre-trained LM over the input sentence without fine-tuning.

Moreover, we find that BERT LMs outperform GPT-2 LMs under similar model sizes. On both datasets, BERT_{BASE} performs better than GPT-2 in F1, and BERT_{LARGE} outperforms GPT-2_{MEDIUM} in F1. This is mainly because the recall of BERT LMs is higher than that of corresponding GPT-2 LMs. The result indicates that the Cloze-style loss function (i.e., masked LM) of BERT is more effective and flexible in recovering information than the autoregressive LM objective. We also notice that the precision of GPT-2 LMs is higher than that of BERT LMs. The reason is that the autoregressive LM objective captures more accurate information than Cloze-style loss does by preventing extra noise (e.g., masked tokens) in pre-training.

Pre-trained LMs achieve competitive precision, e.g., the precision is greater than 60% on TAC KBP-OIE. However, only moderate recalls are obtained. Therefore, improving recall is clearly the future direction. We find that both argument and predicate extraction can be further improved. For example, the main cause of the moderate recall is the incorrect arguments caused by spaCy noun chunks as summarized in Sec. 4.2. Besides, we can incorporate predicates that are not between the argument pairs into the extractions, as we observe a number of gold triples are in inverted sentences. We also notice that the F1 gain over previous state-of-the-arts on TAC KBP-OIE is smaller compared to that on Wikidata-OIE. A larger text corpus, e.g.,

Wikipedia, provides more information. We could improve the recall by running on larger corpora such as WebText2 and Common Crawl (Raffel et al., 2019; Brown et al., 2020) to collect more triples.

4.2 Error Analysis

There is still significant room to improve the results. We argue that we are measuring a lower bound for what LMs know. To further understand the shortcomings of the current method, we conduct an error analysis of the errors in precision on all datasets. We choose BERT_{LARGE} for the study. We sample 100 documents from the Wikidata-OIE dataset, and manually check the reasons for the errors. We find 33.1% of the errors are caused by *incorrect arguments*, while the predicate phrases are correct. The errors are due to the incorrect noun chunks detected by the spaCy. 18.3% of the errors are due to the *missing pairs in predicate mapping*. We also note that approximately 23.8% of the errors are actually *correct triples that are not covered by Wikidata*. For example, (*Bob_Dylan, residence, Nashville*) does not exist in Wikidata, but it is a correct triple. The rest of the errors made by BERT_{LARGE} are *incorrect predicate phrases*, such as uninformative phrases. We find similar errors are made by BERT_{LARGE} on other datasets. Based on the above analysis, enhancing argument detection and predicate mapping is helpful to further improve the results.

4.3 Runtime Analysis

The runtime of OIE systems is crucial in practice. We test the runtime of different OIE systems on Re-OIE2016. The results are in Table 4. We find ours is competitive in terms of efficiency given the size of the models.

4.4 Parameter Study

We study the effects of the key parameters using BERT_{BASE} on TAC KBP-OIE as shown in Figure 3. We randomly sample 20% of the oracle query entities (provided by TAC KBP) as a hold-out

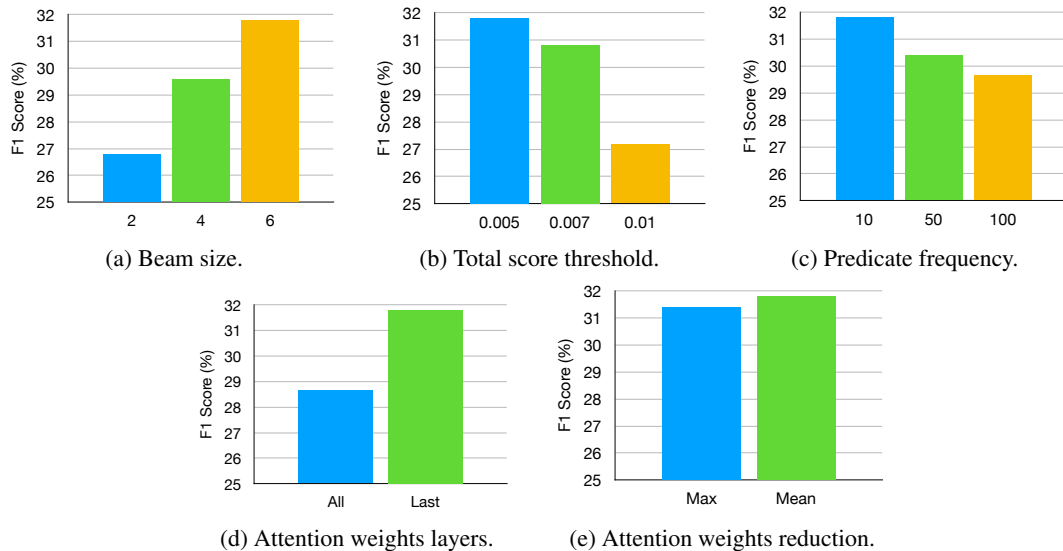


Figure 3: Parameter study with BERT_{BASE} on TAC KBP-OIE hold-out subset.

dataset to tune the parameters, and use the best parameter setting achieved for all experiments. When studying the effect of a certain parameter, we keep the remaining parameters as default. We use F1 to measure the effects. Additional details are described in Appendix A.2.3.

5 Related Work

Pre-trained language models (LM), e.g., BERT (Devlin et al., 2018), GPT (Radford et al., 2018, 2019), and large LMs over 100B parameters (Brown et al., 2020; Chowdhery et al., 2022; Zeng et al., 2022) contain growing amount of linguistic and factual knowledge obtained via pre-training on large-scale corpora. To evaluate their abilities, researchers have created many knowledge benchmarks. LAMA leverages manually created prompts (Petroni et al., 2019, 2020). Recent studies have also developed soft prompts (Liu et al., 2021; Zhong et al., 2021) for fact retrieval. KILT (Petroni et al., 2021) proposes a knowledge-intensive benchmark concerning several downstream tasks to evaluate LMs’ ability in capturing knowledge. Wang et al. (2022) have utilized a set of knowledge-intensive structure prediction tasks to evaluate the knowledge in pre-trained LMs. Shen et al. (2022) have adapted KG completion as a benchmark to evaluate LMs. Besides relational knowledge, closed-book OpenQA (Roberts et al., 2020) benchmarks (in which LMs answer the open-domain questions without retrieving contexts) have also been adopted as a way to evaluate LMs’ knowledge. While the existing benchmarks evaluate LMs in an implicit

way, the main difference is that our benchmark explicitly and interpretably evaluates triples from the textual corpora extracted using model parameters (e.g. attentions). In the field of neural network interpretation (Linzen et al., 2016; Adi et al., 2016; Tenney et al., 2019), in particular the pre-trained deep LM analysis, substantial recent work focuses on both visualizing and analyzing the attention (Vig, 2019; Jain and Wallace, 2019; Clark et al., 2019; Michel et al., 2019; Ramsauer et al., 2020). Instead of analyzing or visualizing, our benchmark quantitatively evaluates the relational information with respect to open information extraction.

Open information extraction systems, e.g., OLLIE (Schmitz et al., 2012), Reverb (Fader et al., 2011), Stanford OpenIE (Angeli et al., 2015), OpenIE 5 (Saha et al., 2017, 2018), RnnOIE (Stanovsky et al., 2018), and OpenIE 6 (Kolluru et al., 2020a) aim to extract triples from web corpora for open schema KGs. Besides, NELL (Carlson et al., 2010), DeepDive (Niu et al., 2012), Knowledge Vault (Dong et al., 2014) extract information based on a fixed schema or ontology, where humans help improve the accuracy of the extractions. Probase (Wu et al., 2012) produces taxonomies instead of rich typed relations in general KGs. Our benchmark first evaluates LMs’ unsupervised information extraction ability on common open information extraction datasets such as CaRB (Bhardwaj et al., 2019) and Re-OIE2016 (Zhan and Zhao, 2020), and then aligns the extracted triples to KG triples for large-scale knowledge extraction benchmark construction. Our algorithm is similar to

the generation algorithm in DeepEx (Wang et al., 2021). The focus of this work is to benchmark the zero-shot OIE performance of pre-trained LMs on both standard and factual OIE datasets. To further improve the OIE performance, the ranking module in DeepEx can be useful. The structure pre-training proposed in (Wang et al., 2022) can also be helpful.

6 Conclusion

We benchmark the general relational information in pre-trained language models (LM) in an open information extraction (OIE) setup. We find that the pre-trained LMs contain a considerable amount of open relational information through large-scale evaluation on both standard OIE datasets and newly created large-scale factual OIE datasets in our IELM benchmark. We are able to turn pre-training LMs into zero-shot OIE systems to efficiently deliver the benchmark results. The reach of this result is broad and has potential downstream utility for deep neural network interpretation, information extraction, and knowledge graph construction. Although the results are promising, we argue that our results just indicate a lower bound about what the LMs have. We hope our results will foster further research in the LM OIE benchmark direction.

7 Limitations

For the limitations of our method, the argument extraction module of our algorithm relies on a third-party noun chunker. As reported, the noun chunker introduces the majority of the errors in our extraction results. A limitation in our benchmark is that we have not conducted a large-scale manual evaluation of our factual OIE datasets (TAC KBP-OIE and Wikidata-OIE). The main focus of our study is to provide a large-scale OIE benchmark. As a result, this makes our benchmark more challenging to be used than standard OIE datasets in terms of computation costs and infrastructure. Finally, we have only benchmarked BERT and GPT-2 on our datasets. Future work could include testing a wide range of language models on our benchmark.

8 Ethical Considerations

We hereby acknowledge that all of the co-authors of this work are aware of the provided *ACM Code of Ethics* and honor the code of conduct. This work is about benchmarking the zero-shot OIE capability of pre-trained language models including BERT and GPT. Our ethical considerations and the work's

underlying future impacts are discussed in the following perspectives. Language models are known to present potential risks and limitations (Brown et al., 2020), and the corpus used in pre-training (such as Wikipedia) may introduce unwanted biases and toxicity. We do not anticipate the production of harmful outputs after using our method or datasets, especially for vulnerable populations.

9 Environmental Impact

We adopt the pre-trained language models BERT (Devlin et al., 2018) and GPT-2 series (Radford et al., 2019) in our IELM benchmark. The models' carbon footprints are estimated to be 22–28 kilograms (Gibney, 2022). Additionally, The focus of this study is to test the zero-shot OIE ability of pre-trained language models. We do not train language models on massive datasets. Instead, we only do inference on a few evaluation datasets. This cost is less than 0.1% energy than that of their pre-training. This demonstrates that developing proper zero-shot learning strategies for large language models can not only deepen our understanding of their latent mechanisms, but also further reduce the energy consumption and environmental impacts that language models with ever-growing size may cause.

Acknowledgements

We would like to thank the anonymous reviewers for their suggestions and comments. This material is in part based upon work supported by Berkeley DeepDrive and Berkeley Artificial Intelligence Research.

References

- Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2016. Fine-grained analysis of sentence embeddings using auxiliary prediction tasks. *arXiv preprint arXiv:1608.04207*.
- Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D Manning. 2015. Leveraging linguistic structure for open domain information extraction. In *ACL*, pages 344–354.
- Sangnie Bhardwaj, Samarth Aggarwal, and Mausam Mausam. 2019. Carb: A crowdsourced benchmark for open ie. In *EMNLP*, pages 6262–6267.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda

- Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R Hruschka Jr, and Tom M Mitchell. 2010. Toward an architecture for never-ending language learning. In *AAAI*, 3.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- Janara Christensen, Stephen Soderland, and Oren Etzioni. 2011. An analysis of open information extraction based on semantic role labeling. In *KCAP*, pages 113–120.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D Manning. 2019. What does bert look at? an analysis of bert’s attention. *arXiv preprint arXiv:1906.04341*.
- Lei Cui, Furu Wei, and Ming Zhou. 2018. Neural open information extraction. In *ACL*, pages 407–413.
- Luciano Del Corro and Rainer Gemulla. 2013. Clausie: clause-based open information extraction. In *WWW*, pages 355–366.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Xin Dong, Evgeniy Gabilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmman, Shaohua Sun, and Wei Zhang. 2014. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *KDD*, pages 601–610.
- Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011. Identifying relations for open information extraction. In *EMNLP*, pages 1535–1545.
- Kiril Gashteovski, Rainer Gemulla, and Luciano del Corro. 2017. Minie: minimizing facts in open information extraction. *ACL*.
- Elizabeth Gibney. 2022. How to shrink ai’s ballooning carbon footprint. *Nature*, 607(7920):648–648.
- John Hewitt and Christopher D Manning. 2019. A structural probe for finding syntax in word representations. In *NAACL*, pages 4129–4138.
- Johannes Hoffart, Mohamed Amir Yosef, Ilaria Bordino, Hagen Fürstenau, Manfred Pinkal, Marc Spaniol, Bilyana Taneva, Stefan Thater, and Gerhard Weikum. 2011. Robust disambiguation of named entities in text. In *EMNLP*, pages 782–792.
- Sarthak Jain and Byron C Wallace. 2019. Attention is not explanation. *arXiv preprint arXiv:1902.10186*.
- Zhengbao Jiang, Frank F Xu, Jun Araki, and Graham Neubig. 2020. How can we know what language models know? *TACL*, 8:423–438.
- Nikolaos Kolitsas, Octavian-Eugen Ganea, and Thomas Hofmann. 2018. End-to-end neural entity linking. *arXiv preprint arXiv:1808.07699*.
- Keshav Kolluru, Vaibhav Adlakha, Samarth Aggarwal, Soumen Chakrabarti, et al. 2020a. Openie6: Iterative grid labeling and coordination analysis for open information extraction. In *EMNLP*, pages 3748–3761.
- Keshav Kolluru, Samarth Aggarwal, Vipul Rathore, Soumen Chakrabarti, et al. 2020b. Imojie: Iterative memory-based joint open information extraction. In *ACL*, pages 5871–5886.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. Assessing the ability of lstms to learn syntax-sensitive dependencies. *TACL*, pages 521–535.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021. Gpt understands, too. *arXiv:2103.10385*.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? In *NIPS*, pages 14014–14024.
- Feng Niu, Ce Zhang, Christopher Ré, and Jude Shavlik. 2012. Elementary: Large-scale knowledge-base construction via machine learning and statistical inference. *International Journal on Semantic Web and Information Systems (IJSWIS)*, 8(3):42–73.
- Fabio Petroni, Patrick Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2020. How context affects language models’ factual predictions. *arXiv preprint arXiv:2005.04611*.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Maillard, et al. 2021. Kilt: a benchmark for knowledge intensive language tasks. In *NAACL*, pages 2523–2544.
- Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2019. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, (8):9.

- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *CoRR*, abs/1910.10683.
- Hubert Ramsauer, Bernhard Schödl, Johannes Lehner, Philipp Seidl, Michael Widrich, Lukas Gruber, Markus Holzleitner, Milena Pavlović, Geir Kjetil Sandve, Victor Greiff, et al. 2020. Hopfield networks is all you need. *arXiv preprint arXiv:2008.02217*.
- Youngbin Ro, Yookyung Lee, and Pilsung Kang. 2020. Multi²oie: Multilingual open information extraction based on multi-head attention with bert. In *Findings of EMNLP*, pages 1107–1117.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? *arXiv preprint arXiv:2002.08910*.
- Arpita Roy, Youngja Park, Taesung Lee, and Shimej Pan. 2019. Supervising unsupervised open information extraction models. In *EMNLP*, pages 728–737.
- Swarnadeep Saha, Harinder Pal, et al. 2017. Bootstrapping for numerical open ie. In *ACL*, pages 317–323.
- Swarnadeep Saha et al. 2018. Open information extraction from conjunctive sentences. In *COLING*, pages 2288–2299.
- Michael Schmitz, Stephen Soderland, Robert Bart, Oren Etzioni, et al. 2012. Open language learning for information extraction. In *EMNLP*, pages 523–534.
- Jianhao Shen, Chenguang Wang, Linyuan Gong, and Dawn Song. 2022. Joint language semantic and structure embedding for knowledge graph completion. In *COLING*.
- Valentin I Spitzkovsky and Angel X Chang. 2012. A cross-lingual dictionary for english wikipedia concepts.
- Gabriel Stanovsky and Ido Dagan. 2016. Creating a large benchmark for open information extraction. In *EMNLP*, pages 2300–2305.
- Gabriel Stanovsky, Jessica Fidler, Ido Dagan, and Yoav Goldberg. 2016. Getting more out of syntax with props. *arXiv preprint arXiv:1603.01648*.
- Gabriel Stanovsky, Julian Michael, Luke Zettlemoyer, and Ido Dagan. 2018. Supervised open information extraction. In *NAACL*, pages 885–895.
- Mihai Surdeanu. 2013. Overview of the tac2013 knowledge base population evaluation: English slot filling and temporal slot filling. *TAC*, page 2.
- Alon Talmor, Yanai Elazar, Yoav Goldberg, and Jonathan Berant. 2019. olympics—on what language model pre-training captures. *arXiv preprint arXiv:1912.13283*.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R Bowman, Dipanjan Das, et al. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. *arXiv preprint arXiv:1905.06316*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 5998–6008.
- Jesse Vig. 2019. Visualizing attention in transformer-based language models. *arXiv preprint arXiv:1904.02679*.
- Chenguang Wang, Xiao Liu, Zui Chen, Haoyun Hong, Jie Tang, and Dawn Song. 2021. Zero-shot information extraction as a unified text-to-triple translation. In *EMNLP*.
- Chenguang Wang, Xiao Liu, Zui Chen, Haoyun Hong, Jie Tang, and Dawn Song. 2022. DeepStruct: Pre-training of language models for structure prediction. In *ACL*.
- Wentao Wu, Hongsong Li, Haixun Wang, and Kenny Q Zhu. 2012. Probase: A probabilistic taxonomy for text understanding. In *SIGMOD*, pages 481–492.
- Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. Glm-130b: An open bilingual pre-trained model. *arXiv preprint arXiv:2210.02414*.
- Junlang Zhan and Hai Zhao. 2020. Span model for open information extraction on accurate corpus. In *AAAI*, pages 9523–9530.
- Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D Manning. 2017. Position-aware attention and supervised data improve slot filling. In *EMNLP*.
- Zexuan Zhong, Dan Friedman, and Danqi Chen. 2021. Factual probing is [mask]: Learning vs. learning to recall. In *NAACL*, pages 5017–5033.

A The IELM Benchmark Details

Additional details of our open information extraction (OIE) benchmark IELM are described in this section.

A.1 Wikidata-OIE

In this section, we describe some technical details regarding the construction and evaluation of Wikidata-OIE.

A.1.1 Entity Linking

We use an unsupervised entity linker for both Wikidata-OIE dataset construction and OIE evaluation. The entity linker is originally developed in (Spitkovsky and Chang, 2012), which is based on a mention-to-entity dictionary. We build an enhanced dictionary as follows: we add new Wikipedia anchors to the dictionary which results in 26 million entries compared to the original 21 million entries. Then a Wikipedia anchor to the Wikidata item dictionary is used to further link the entities (or arguments) to Wikidata. If an argument is a pronoun, we further use neuralcoref² for coreference resolution.

A.1.2 Predicate Mapping

The predicate mapping of Wikidata-OIE is constructed offline using the method in Sec. 3.4. In more detail, we randomly sampled a hold-out dataset including 2,000 documents from English Wikipedia for the bootstrapped predicate mapping construction based on the TAC KBP mapping (Angeli et al., 2015). To filter out the wrong predicate pairs, we manually check whether the top predicate phrases are true.

A.1.3 Gold Triples

For gold triples in Wikidata-OIE, we only preserve those triples describing predicates between arguments that can be linked to corresponding Wikipedia anchors. We rule out triples of attributes about arguments and triples of auxiliary predicates (such as *topic’s main category.P901*) and finally result in 27,368,562 gold triple extractions.

A.1.4 Evaluation

Given the large number of source sentences and gold triples in Wikidata-OIE, a MongoDB database is maintained to store the gold triples to enable an efficient evaluation.

A.2 Zero-Shot Language Model Based Open Information Extraction

In this section, we introduce additional details about how we adapt pre-trained language models (LM) as zero-shot OIE systems.

A.2.1 Argument Extraction

We use spaCy noun chunker³ to annotate the noun phrases in the sentences.

²<https://github.com/huggingface/neuralcoref>

³<https://spacy.io/usage/linguistic-features/#noun-chunks>

Algorithm 1 Beam search with attention scores.

Input: Argument pair (arg_0, arg_1) , sentence s , attention matrix \mathbf{A}_s , action manager $\mathcal{O} = \{\text{START}, \text{YIELD}, \text{STOP}\}$, beam size k

Output: Triples $T_{(arg_0, arg_1)}$

- 1: $T_{(arg_0, arg_1)} \leftarrow \{\text{START}(arg_0)\}$ \triangleright Start by adding the first argument as a candidate in the beam
- 2: **while** $\exists c \in T_{(arg_0, arg_1)}[\mathcal{O}(c) = \text{YIELD}]$ **do**
- 3: $\tilde{T}_{(arg_0, arg_1)} \leftarrow \emptyset$ \triangleright Initialize a new beam
- 4: **for each** $c \in T_{(arg_0, arg_1)}$ **do**
- 5: **if** $\mathcal{O}(c) = \text{YIELD}$ **then**
- 6: $\tilde{T}_{(arg_0, arg_1)} \leftarrow \tilde{T}_{(arg_0, arg_1)} \cup \{\text{YIELD}(c, s, \mathbf{A}_s)\}$ \triangleright Yield a new candidate if not reached the second argument
- 7: **else**
- 8: $\tilde{T}_{(arg_0, arg_1)} \leftarrow \tilde{T}_{(arg_0, arg_1)} \cup \{\text{STOP}(c, t)\}$ \triangleright Stop then produce a valid triple if reached the second argument
- 9: **end if**
- 10: **end for**
- 11: $T_{(arg_0, arg_1)} \leftarrow \text{TOP}(k, \tilde{T}_{(arg_0, arg_1)})$ \triangleright Maintain k -best candidates in the beam
- 12: **end while**
- 13: **return** $T_{(arg_0, arg_1)}$

A.2.2 Predicate Extraction

We first describe predicate extraction introduced in Sec. 2.2 in detail.

- **Beam Search.** The inputs of the search algorithm are an argument pair (arg_0, arg_1) , a sentence s , an attention matrix \mathbf{A}_s of s . Both arg_0 and arg_1 are identified by the noun chunker in s . \mathbf{A}_s is the attention matrix associated with s from the forward pass of an LM without fine-tuning. The search gets started by adding the first argument arg_0 as the initial candidate in the beam. While there are still new candidates waiting to be yielded, the search continues, and the top k candidates sorted by the attention scores are maintained in the beam. The details of the proposed beam search are described in Algorithm 1. In practice, we implement an action manager \mathcal{O} to decide which action to take at each step. Given a candidate c in the beam, $\mathcal{O}(c) = \text{START}$ always happens at the beginning of the search. If c has not reached the second argument arg_1 yet, $\mathcal{O}(c) = \text{YIELD}$. Otherwise, $\mathcal{O}(c) = \text{STOP}$.
- **Implementation Details.** For Wikidata-OIE, we randomly split the English Wikipedia data into 20 partitions, and map the data partitions to 20 distributed servers to run. Each server is configured with four Tesla K80 12Gs. We set the max sequence length to 256, and batch size as 32 for BERT_{LARGE} and 4 for GPT-2_{XL}. We use implementations of pre-trained LMs in the

Transformers package ⁴. We use spaCy sentencizer ⁵ to segment the documents into sentences. BERT_{LARGE} takes approximately 48 hours, and GPT-2_{XL} costs around 96 hours. The resulting triples from the 20 servers are then reduced to a data server. The batch sizes of BERT_{BASE}, GPT-2, GPT-2_{MEDIUM}, GPT-2_{LARGE} are 64, 32, 16, 8 respectively.

A.2.3 Parameter Settings

We then discuss the parameter setup of our OIE systems as below.

The parameter settings are shared across all OIE datasets. All the choices are based on the parameter study in Sec. 4.4. The beam size of Algorithm 1 is set to 6. The attention score threshold is set to 0.005, and the number of relation/predicate frequencies is set to 10. To generate the attention weight matrix A_s of a sentence, we reduce the weights of every attention head in the last layer of pre-trained LMs using the mean operator. We analyze the effects of various parameters below.

Figure 3(a) illustrates the effects of various beam sizes in Algorithm 1. We find that in general, the larger the beam size is, the better F1 the setting achieves. This is because our method is able to reserve more potentially correct triples when more candidates are allowed. However, F1 improvement gradually becomes subtle, while the computation costs increase more significantly. For efficiency consideration, we do not explore larger beam sizes. We set the beam size as 6.

Figure 3(b) compares the effect of different thresholds of the total score. We set the threshold as 0.005 since it achieves the best result. Note that the summed attention score is normalized by the length of the triple to penalize the cumbersome triples. The threshold is effective. This is mainly because of the relational information contained in the self-attention matrix: the score in the attention matrix is representing the chance of the triples to be the true triples based on the stored information. Figure 3(c) shows the impact of the predicate frequency threshold in identifying common predicates. The best result is achieved when it equals 10. This shows that while our method mostly identifies frequent predicates, it is also able to capture some rare predicates.

Figure 3(d) shows the comparison between the attention weights of the last layer and the mean of

all layers. The attention weights of the last layer perform better. This is due to the attention weights in lower layers being low-level linguistic knowledge according to (Clark et al., 2019; Ramsauer et al., 2020), which are less relevant to the relational information. Figure 3(e) compares the impact of different attention reductions, i.e., mean, max, over the attention heads of the last layer. We find the “mean” performs better. The reason is that the token often intensively attends to several specific tokens in the sequence (Michel et al., 2019), and the “mean” operator is less sensitive to such biases.

A.3 The Number of Predicates of Standard and Factual OIE Datasets

Note that there are more predicates in standard OIE datasets than that in factual datasets. This is because, for standard OIE, predicates are open and not attached to a certain schema. These predicates were extracted from the input sentences and are usually natural language utterances. For factual OIE, the predicates are unified into a fixed KG schema. For example, for a person’s birthplace, there are multiple natural language expressions like “was born in” or “gave birth” in standard OIE datasets, while only a single “birth_place” predicate exists in the factual OIE sets.

A.4 Predicate Mapping Examples

We show example predicate mappings in a dictionary below.

- `per:city_of_birth`: born in, born at, born, birth city, hometown.
- `org:founded_by`: established by, founded by, founded, founder, co-founder of.

where the keys are KG predicates, e.g., `per:city_of_birth` and `org:founded_by`. The values are the corresponding OIE relations.

A.5 Comparison Systems

We compare our zero-shot OIE systems with the following OIE systems.

A.5.1 Neural OIE Systems

The following neural network based systems are selected.

- SenseOIE (Roy et al., 2019)⁶ learns to ensemble various previous unsupervised OIE systems’ ex-

⁴<https://github.com/huggingface/transformers>

⁵<https://spacy.io/api/sentencizer>

⁶The code is not available.

tractions using supervised learning to combine their strengths.

- SpanOIE (Zhan and Zhao, 2020)⁷ presents the Re-OIE2016 datasets for a more rigorous evaluation and a span-based (instead of sequence labeling) extraction model.
- RnnOIE (Stanovsky et al., 2018)⁸ is one of the state-of-the-art OIE systems. It uses LSTM to model the OIE problem as a sequence tagging problem, and is trained on a large-scale OIE training set.
- NeuralOIE (Cui et al., 2018)⁹ is an encoder-decoder based architecture that adopts the copy mechanism to conduct OIE.
- IMOJIE (Kolluru et al., 2020b)¹⁰ is a sequence generation based OIE model that uses BERT at encoding time.
- Multi²OIE (Ro et al., 2020)¹¹ models OIE as a sequence labeling problem that combines BERT with multi-head attention blocks.
- OpenIE6 (Kolluru et al., 2020a)¹² is one of the state-of-the-art OIE systems. It treats OIE as a 2-D grid labeling task, and trains a BERT family architecture for the task.

Note that while our methods are zero-shot without needing to use the specific training sets, all the neural OIE systems are supervised on corresponding training sets.

A.5.2 Linguistic OIE Systems

We also compare our systems with the following linguistic pattern based systems developed prior to the use of neural networks.

- MinIE (Gashteovski et al., 2017)¹³ proposes to minimize facts in OIE by representing information by annotations rather than extraction and removing redundant specific information.
- ClausIE (Del Corro and Gemulla, 2013)¹⁴ is a clause-based approach by first identifying lin-

⁷https://github.com/zhanjunlang/Span_OIE

⁸<https://github.com/gabrielStanovsky/supervised-oie>

⁹We use the BERT implementation available at <https://github.com/dair-iitd/imojie>.

¹⁰<https://github.com/dair-iitd/imojie>

¹¹<https://github.com/youngbin-ro/Multi2OIE>

¹²<https://github.com/dair-iitd/openie6>

¹³<https://github.com/uma-pi1/minie>

¹⁴<https://www.mpi-inf.mpg.de/departments/databases-and-information-systems/research/ambiverse-nlu/clausie>

guistic structure and then their information and attributes.

- OLLIE (Schmitz et al., 2012)¹⁵ uses contextual sentence decomposition to conduct OIE.
- PropS (Stanovsky et al., 2016)¹⁶ proposes proposition structure which is implied from syntax using dependency trees.
- OpenIE4 (Christensen et al., 2011)¹⁷ is the successor to OLLIE using similar argument and relation expansion heuristics to create OIE extractions from semantic role labeling frames.
- OpenIE5 (Saha et al., 2017, 2018)¹⁸ is one of the state-of-the-art OIE systems, which is the successor to OLLIE, and it improves extractions from noun relations, numerical sentences, and conjunctive sentences depending on the linguistic patterns.
- Stanford OpenIE (Angeli et al., 2015)¹⁹ leverages POS tag and dependency parser, and generates self-contained clauses from long sentences to extract the triples.

B The TAC KBP-OIE and Wikidata-OIE Datasets

We show samples of our zero-shot OIE extractions and the gold triples on both TAC KBP-OIE and Wikidata-OIE datasets.

B.1 TAC KBP-OIE

OIE Extractions and Gold Extractions We randomly sample 100 documents from the TAC KBP-OIE corpus, then sample sentences from those documents. The uncurated triples and the corresponding gold triples of the sampled sentences based on our best methods BERT_{LARGE} and GPT-2_{XL} are shown in Figure 4 and Figure 5 respectively. We also randomly sample sentences in which BERT_{LARGE} differs from GPT-2_{XL} in the resulting triples for comparison, which are illustrated in Figure 6. In each table, “**ID**” represents the document ID of a sampled sentence in the TAC KBP-OIE corpus. “**Sentence**” indicates the sampled sentence. “**Triples to gold triples**” column contains the extraction triples (on the left side of

¹⁵<https://github.com/knowitall/ollie>

¹⁶<https://github.com/gabrielStanovsky/props>

¹⁷<https://github.com/allenai/openie-standalone>

¹⁸<https://github.com/dair-iitd/OpenIE-standalone>

¹⁹<https://nlp.stanford.edu/software/openie.html>

“→”) and their corresponding gold triples (on the right side of “→”).

B.2 Wikidata-OIE

OIE Extractions and Gold Extractions Similar to TAC KBP-OIE, we randomly sample 100 documents from the Wikidata-OIE corpus (i.e., English Wikipedia), then sample sentences from those documents. Similar to TAC KBP-OIE, Figure 7 and Figure 8 show the uncurated triples and the corresponding gold triples of the sampled sentences based on our zero-shot systems $BERT_{LARGE}$ and $GPT-2_{XL}$ respectively. Figure 9 illustrates the randomly sampled sentences in which $BERT_{LARGE}$ extracts different triples compared to that from $GPT-2_{XL}$. In each table, “**ID**” represents the Wikipedia page’s title of the sampled sentence. “**Sentence**” indicates the sampled sentence. “**Triples to gold triples**” column contains the triples (on the left side of “→”) and their corresponding gold triples (on the right side of “→”).

ID	Sentence	Triplets to gold triples
SP13_ENG_001	Bashardost left Pakistan for France in 1981	(Bashardost, left, Pakistan) → (Ramazan Bashardost, per:countries_of_residence, Pakistan)
SP13_ENG_004	"Douglas Flint will succeed Stephen Green as Group Chairman and Stuart Galliver will be appointed Group Chief Executive, a statement Michael Govegan said in a statement."	(Douglas Flint, as, Group Chairman) → (Douglas Flint, per:title, Chairman)
SP13_ENG_007	Mohammed Tantai, head of Al-Azhar University, told a schoolgirl to remove her niqab when he spotted her during a tour of an Al-Azhar affiliated school, the independent Al-Masry al-Youn newspaper reported his work.	(Mohammed Tantai, head, of, Al-Azhar University) → (Mohammed Sayed Tantai, per:employee_or_member_of, Al-Azhar University)
SP13_ENG_009	He took office in 2006 by defeating longtime incumbent Arwar Chowdhry from Pakistan, who was later barred for alleged financial corruption.	(Logimic Incumbent Arwar Chowdhry, from, Pakistan) → (Anwar Chowdhry, per:origin, Pakistan)
SP13_ENG_010	In addition to his wife, Wendy, Dho is survived by son Daniel, grandchildren Julie and Joey, and father Pat.	(Dho, is, survived, by, son, Daniel) → (Ronnie James Dho, per:children, Daniel)
SP13_ENG_011	Marshall is charged with grand larceny and fraud and faces up to 25 years in prison if convicted.	(Marshall, is, charged, with, grand larceny and fraud) → (Anthony Marshall, per:charges, fraud)
SP13_ENG_012	Marshall, a Tony Award-winning Broadway producer and former U.S. diplomat, sat nonetheless as the jury forewoman read each verdict about the word "guilty" clearly resonating in the otherwise silent courtroom.	(Marshall, a, Tony Award-winning Broadway Producer) → (Anthony Marshall, per:title, producer)
SP13_ENG_011	The blueblooded scion of one of America's most illustrious families appeared to listen impassively as verdicts finding him guilty on 14 counts of grand larceny, conspiracy and fraud were read to a packed courtroom.	(Him, guilty, Grand Larceny) → (Anthony Marshall, per:charges, grand larceny)
SP13_ENG_011	But Marshall's son, Philip, told a different story.	(Marshall's Son, Philip) → (Anthony Marshall, per:children, Philip)
SP13_ENG_012	Al-Qaeda's American spokesman Adam Gadahn, also known as Azam the American, called on Muslims in the West on Sunday to carry out more attacks like the deadly shooting at the US base in Fort Hood, Texas.	(Adam Gadahn, also known, as, Azam) → (Adam Gadahn, per:alternate_names, Azam)
SP13_ENG_012	Gadahn, also known as Azam the American, was born in 1978.	(Gadahn, also known as, Azam) → (Adam Gadahn, per:alternate_names, Azam)
SP13_ENG_012	Gadahn grew up in California and converted to Islam before he moved to Pakistan in 1998 and attended an al-Qaeda training camp six years later, according to media reports.	(Gadahn, moved to, Pakistan) → (Adam Gadahn, per:countries_of_residence, Pakistan)
SP13_ENG_012	Gadahn moved to Pakistan in 1998, according to the FBI, and is said to have attended an al-Qaeda training camp six years later, serving as a translator and consultant for the group.	(He, lived in, Lebanon) → (Mohammed Oudch, per:countries_of_residence, Lebanon)
SP13_ENG_014	After the Munich attack, he lived in Lebanon, Jordan and several Eastern European countries, where he had close ties to Communist bloc intelligence agencies.	(Clifton, attended, Howard University) → (Lacelle Clifton, per:schools_attended, Howard University)
SP13_ENG_015	Clifton attended Howard University but left before graduating to pursue poetry.	(Clifton, attended, Howard University) → (Lacelle Clifton, per:schools_attended, Howard University)
SP13_ENG_017	Alexander Haig devoted his career to serving our country, both as a soldier and as a diplomat, Albright said.	(Haig, and retired from, The Army) → (Alexander Haig, per:employee_or_member_of, Army)
SP13_ENG_019	McGregor is survived by his wife, Lori, and four children, daughters Taylor and Landri, and a son, Logan.	(McGregor, his, wife, Lori) → (Kell McGregor, per:spouse, Lori)
SP13_ENG_020	"Mike was a first-rate journalist, a valued member of our staff for 25 years and we will miss him," Times Editor Russ Stanton said.	(Mike, was, a, First-Rate Journalist) → (Mike Penner, per:title, first-rate journalist)
SP13_ENG_020	Penner is survived by his brother, John, a copy editor at the Times, and his former wife, Times sportswriter Lisa Dillman.	(Penner, his, brother, John) → (Mike Penner, per:siblings, John)
SP13_ENG_024	She was charged with theft in Beaumont, Texas, for allegedly failing to pay for \$10,000 worth of dental work in 2006.	(She, was, charged, with) → (Crystal Taylor, per:charges, theft)
SP13_ENG_025	Mexican native Nancy Kissel was convicted of murder and sentenced to 10 years in prison in Hong Kong, a High Court in September 2005.	(Mexican Native Nancy Kissel, was, convicted, Murder) → (Nancy Kissel, per:charges, murder)
SP13_ENG_026	Neal returned to New York and concentrated on stage work.	(Neal, returned to, New York) → (Patricia Neal, per:states/provinces_of_residence, New York)
SP13_ENG_026	In 1953, she married Ronald Dahl, the British writer famed for "Charlie and the Chocolate Factory" and "James and the Giant Peach" and other tales for children.	(She, married, Ronald Dahl) → (Patricia Neal, per:spouse, Ronald Dahl)
SP13_ENG_026	Oscar-winning actress Patricia Neal had died of lung cancer at her home on Martha's Vineyard, Massachusetts, on Sunday.	(Patricia, died, of, Lung Cancer) → (Patricia Neal, per:cause_of_death, lung cancer)
SP13_ENG_027	Al-Hakim's son, Ammar al-Hakim, has been groomed for months to take his father's place.	(Al-Hakim's Son, Ammar Al-Hakim) → (Abdul Aziz Al-Hakim, per:children, Ammar al-Hakim)
SP13_ENG_027	Al-Hakim, the head of Supreme Iraqi Islamic Council (SIC), the largest Shiite party in Iraq.	(Al-Hakim, is, Supreme Iraqi Islamic Council) → (Abdul Aziz Al-Hakim, per:employee_or_member_of, Supreme Iraqi Islamic Council)
SP13_ENG_027	His former Shiite partners have gathered again to form their own group, the Iraqi National Alliance (INA), which includes the influential Supreme Iraqi Islamic Council (SIC) of Ammar al-Hakim, who succeeded his father Abdul Aziz al-Hakim, who died in a hospital in Iran last month after a long battle with cancer.	(Al-Hakim, who, died in, Iran) → (Abdul Aziz Al-Hakim, per:country_of_death, Iran)
SP13_ENG_028	"I rather have Sally doing this than some stranger, or some foolbot trying to be the next Billy Mays," said the guy who actually is the next Billy Mays, his son Billy Mays III.	(The Next Billy Mays, his son, Billy Mays III) → (Billy Mays, per:children, Billy Mays III)
SP13_ENG_029	Fignon continued cycling during and after a stint in the Army, and drew attention in the early 1980s when he managed to keep up with Hinault during a race in which amateurs and professionals rode together.	(Fignon, during and after a stint in, The Army) → (Laurent Fignon, per:employee_or_member_of, The Army)
SP13_ENG_029	Laurent Patrick Fignon was born in Paris on Aug. 12, 1960.	(Laurent Patrick Fignon, was, born in, Paris) → (Laurent Fignon, per:city_of_birth, Paris)
SP13_ENG_030	Anderson became the Tigers' manager in June 1979 and built on a foundation that included Alan Trammell at shortstop, Lou Whitaker at second and the outfield and Jack Morris on the pitching staff.	(Anderson, became, The Tigers' Manager) → (Sparky Anderson, per:title, manager)
SP13_ENG_030	In addition to his wife, Carol, Anderson is survived by his sons, Lee and Albert; his daughter, Shirlee Englebrecht; and many grandchildren.	(Anderson, is, survived by, his sons, Lee) → (Sparky Anderson, per:children, Albert)
SP13_ENG_031	Blake Edwards, a writer and director who was hailed as a Hollywood master of screwball farces and rade comedies like "Victor/Victoria" and the "Pink Panther" movies, died Wednesday night in Santa Monica, Calif. He was 88.	(Blake Edwards, writer) → (Blake Edwards, per:title, writer)
SP13_ENG_032	Hwang, who lives in Seoul under tight police security, has written books and delivered speeches condemning Kim's regime as authoritarian and dictatorial.	(Hwang, who, lives in, Seoul) → (Hwang Kang-Yop, per:city_of_residence, Seoul)
SP13_ENG_033	Kaczynska, who was 67, married Kaczynski in 1978 after meeting him in the northern Polish city of Gdansk, where they were both academics.	(Kaczynska, married, Kaczynski) → (Maria Kaczynska, per:spouse, Kaczynski)
SP13_ENG_036	Upon his release he went into exile in India, where he masterminded the 1973 hijacking of a Royal Nepal Airlines plane known to be carrying hundreds of thousands of dollars in cash to fund his banned Nepali Congress party.	(His Release, exile, India) → (Girija Prasad Koirala, per:countries_of_residence, India)
SP13_ENG_036	Koirala began his political career as a union organizer and was imprisoned for seven years in 1960 after a failed uprising against the monarchy.	(Koirala, began his, A Union Organizer) → (Girija Prasad Koirala, per:title, union organizer)
SP13_ENG_036	Koirala was born in 1925 in Bihar of India where his father Krishna Prasad Koirala and his family were living in exile.	(Koirala, was, born in, Bihar) → (Girija Prasad Koirala, per:city_of_birth, Bihar)
SP13_ENG_036	Koirala was born in 1925 in Bihar of India at the time when his father Krishna Prasad Koirala along with his family was exiled by British rulers.	(Koirala, was, born in, Bihar) → (Girija Prasad Koirala, per:city_of_birth, Bihar)
SP13_ENG_037	Chabrol's survivors also include his third wife, Aureo Pajot, who acted as his script supervisor on nearly all of his movies from 1968 on and whom he married in 1981; and Pajot's daughter, Cecile Maistre, who was an assistant director on his films and wrote the script with him for "The Girl on the Train" (2002).	(Chabrol's Survivors, third wife, Aureo Pajot) → (Claude Chabrol, per:spouse, Aureo Pajot)
SP13_ENG_038	The joint statement said Cunningham was "an inspiring performer and dancer into his 80s, and a visionary choreographer and dedicated teacher throughout his life; he led quietly and by example," the statement said.	(Cunningham, was, An Inspiring Performer) → (Merce Cunningham, per:title, performer)
SP13_ENG_038	Merce Cunningham, the nonagenarian choreographer, is planning for a world without him.	(Merce Cunningham, the Nonagenarian Choreographer) → (Merce Cunningham, per:title, choreographer)
SP13_ENG_039	Account on Monday cleared the widower of British reality television star Jade Goody, who died of cancer last year, of rape.	(British Reality Television Star Jade, who, died, Cancer) → (Jade Goody, per:cause_of_death, cancer)
SP13_ENG_041	Don Hewitt, the CBS newsmen who invented the highly popular TV newsmagazine "60 Minutes" and produced it for 36 years, died Wednesday.	(Don Hewitt, the, CBS Newsmen) → (Don Hewitt, per:title, newsmen)
SP13_ENG_041	He was the consummate television newsmen. Don Hewitt, a longtime CBS News executive and creator of the long-running "60 Minutes" news program, told Reuters.	(Don Hewitt, executive) → (Don Hewitt, per:title, executive)
SP13_ENG_041	Hewitt was already a highly respected TV newsmen.	(Hewitt, was, A Highly Respected TV Newsmen) → (Don Hewitt, per:title, TV newsmen)
SP13_ENG_041	Donald Shepard Hewitt was born in New York on Dec. 14, 1922, and grew up in the suburb of New Rochelle.	(Donald Shepard Hewitt, born, New York) → (Don Hewitt, per:state/province_of_birth, New York)
SP13_ENG_043	Elizane Louise Greenweck was born in Brooklyn on Oct. 23, 1940.	(Elizane Louise Greenweck, was, born, Oct.) → (Lillie Greenweck, per:date_of_birth, 1940-10-23)
SP13_ENG_044	A little more than a year after Dumez died from bladder cancer, the colorful remains of his estate have been consigned by his family to Star Galleries in Hudson, N.Y., which will auction them Nov. 20.	(Dumez, died, Bladder Cancer) → (Dominick Dumez, per:cause_of_death, bladder cancer)
SP13_ENG_048	Charles Gwathmey, an architect known for his influential modernist home designs and famous clients like director Steven Spielberg, has died.	(Charles Gwathmey, architect) → (Charles Gwathmey, per:title, architect)
SP13_ENG_049	Besides his wife, Mandelbrot is survived by two sons, Laurent, of Paris, and Didier, of Newton, Mass., and three grandchildren.	(Mandelbrot, is, survived by, two sons, Laurent) → (Benoit Mandelbrot, per:children, Laurent)
SP13_ENG_049	For years, he worked for IBM in New York.	(He, for, IBM) → (Benoit Mandelbrot, per:employee_or_member_of, IBM)
SP13_ENG_049	After several years spent largely at the Centre National de la Recherche Scientifique in Paris, Mandelbrot was hired by IBM in 1958 to work at the Thomas J. Watson Research Center in Yorktown Heights, N.Y. Although he worked frequently with academic researchers and served as a visiting professor at Harvard and the Massachusetts Institute of Technology, it was not until 1987 that he began to teach at Yale, where he earned tenure in 1999.	(Mandelbrot, was, hired by, IBM) → (Benoit Mandelbrot, per:employee_or_member_of, IBM)
SP13_ENG_056	"It's an issue for everybody in the state because peanuts are a big part of our economy," said Don Koehler, executive director of the Georgia Peanut Commission.	(The Georgia Peanut Commission, director, Don Koehler) → (Georgia Peanut Commission, org:top_members_employees, Don Koehler)
SP13_ENG_060	"We'll be meeting with scientists, university and science policy officials to explore practical opportunities for exchange and collaboration," Agre, the AAAS president, was quoted as saying.	(The Aas President, Agre) → (American Association for the Advancement of Science, org:top_members_employees, Peter C. Agre)
SP13_ENG_060	However, Alan Leshner, chief executive officer of the American Association for the Advancement of Science, noted that Nobels are generally given for work that's a decade old or more, and that the U.S. mustn't become complacent.	(The American Association, chief executive, Alan Leshner) → (American Association for the Advancement of Science, org:top_members_employees, Alan Leshner)
SP13_ENG_062	"Fits of rail, they never have enough funding," said Andy Kunz, president of the U.S. High Speed Rail Association, a nonprofit that advocates a national high-speed rail network.	(The U.S. High Speed Rail Association, Andy Kunz) → (U.S. High Speed Rail Association, org:top_members_employees, Andy Kunz)
SP13_ENG_064	China's shock at NATO's military campaign in the former Yugoslavia helped prod Beijing into playing a bigger role in U.N. peacekeeping, said Bates Gill, director of the Stockholm International Peace Research Institute and co-author of a recent report on China's peacekeeping activities.	(The Stockholm International Peace Research Institute, director, Bates Gill) → (Stockholm International Peace Research Institute, org:top_members_employees, Bates Gill)
SP13_ENG_064	"Non-state actors, for example, a small group of pirates off the coast of Somalia, Al Qaeda and Taliban who operate across borders and have more and more sophisticated means of violence, are becoming bigger and bigger challenges to the international system," said Bates Gill, director of the Stockholm International Peace Research Institute.	(The Stockholm International Peace Research Institute, director, Bates Gill) → (Stockholm International Peace Research Institute, org:top_members_employees, Bates Gill)
SP13_ENG_068	"It is not surprising that one primarily cosmetic business is trying to know another under the bon by transferring a tax from rich doctors and their wealthy customers to struggling small businesses," John Overstreet, director of the Indoor Tanning Association, said in a statement Saturday.	(The Indoor Tanning Association, director, John Overstreet) → (Indoor Tanning Association, org:top_members_employees, John Overstreet)
SP13_ENG_076	The chairman of the Swiss Bankers Association, Patrick Olier, said weekly NZZ am Sonntag that Italy and France have shown interest in deals like ones Switzerland signed this week with Germany and Britain.	(The Swiss Bankers Association, Patrick Olier) → (Swiss Bankers Association, org:top_members_employees, Patrick Olier)
SP13_ENG_076	The majority of voters in Switzerland, which manages more than 75 percent of the world's foreign-held private wealth, support banking secrecy, according to a survey published last month by the Swiss Bankers Association in Basel.	(The Swiss Bankers Association in, Basel) → (Swiss Bankers Association, org:city_of_headquarters, Basel)
SP13_ENG_078	"Americans have a right to know the truth - Islam is a religion of intolerance and violence," said Richard Thompson, legal director of the Thomas More Law Center in Ann Arbor.	(The Thomas More Law Center in, Ann Arbor) → (Thomas More Law Center, org:city_of_headquarters, Ann Arbor)
SP13_ENG_083	New solutions may be enacted for these orphan, though, said Mary Robinson, CEO of the National Council for Adoption.	(The National Council, Mary Robinson) → (National Council for Adoption, org:top_members_employees, Mary Robinson)
SP13_ENG_082	"When you close a company, you end up causing more problems than you prevented," said Chuck Johnson, CEO of the National Council for Adoption.	(Adoption, Council for, Chuck Johnson) → (National Council for Adoption, org:top_members_employees, Chuck Johnson)
SP13_ENG_084	"This is definitely a Goldilocks problem," said Jason Grunnet, president of the Bipartisan Policy Center and an energy advisor to the Obama campaign last year.	(The Bipartisan Policy Center, Jason Grunnet) → (Bipartisan Policy Center, org:top_members_employees, Jason Grunnet)
SP13_ENG_085	"Banks are in strong need for the capital markets (to raise funds)," Li Fuan, a director at the China Banking Regulatory Commission, was quoted as saying at a forum over the weekend.	(The China Banking Regulatory Commission, director at, Li Fuan) → (China Banking Regulatory Commission, org:top_members_employees, Li Fuan)
SP13_ENG_089	RITA Novosti and interfax cite Anatoly Iaukin, head of Rosoboronexport, as saying Thursday "nothing is blocking the continuation of military-technical cooperation" with Iran.	(Rosoboronexport, head of, Anatoly Iaukin) → (Rosoboronexport, org:top_members_employees, Anatoly Iaukin)
SP13_ENG_089	Rosoboronexport is the only company in Russia that is allowed to export arms, dual-use products and military-related services.	(Rosoboronexport, is, the only company in, Russia) → (Rosoboronexport, org:country_of_headquarters, Russia)
SP13_ENG_089	With his wife, Cornelia, Middelhoff invested money in 2000 and 2001 with Esch in funds that were formed to buy five properties from KarstadtQuelle, as Arcandor was then known, and leased back to the department store chain before Middelhoff joined the company, according to Middelhoff's spokesman.	(Arcandor, KarstadtQuelle) → (Arcandor, org:alternate_names, KarstadtQuelle)
SP13_ENG_091	Arcandor's main-order subsidiary Quelle is in worse shape however, Klaus Hubert Goertz told a press conference in the western city of Essen, where Arcandor is based.	(Arcandor, Essen) → (Arcandor, org:city_of_headquarters, Essen)

Figure 4: BERT_{LARGE} on TAC KBP-OIE.

ID	Sentence	Triplets to gold triples
SF13_ENG_001	In a sign of just how disenfranchised some Afghans have become with their government, Baharadost, a doctoral scholar who lived in France for two decades, is widely believed to be at least fourth in popularity among 42 candidates in the August elections.	(Baharadost, A Doctoral Scholar) → (Ramazan Baharadost, per:title, scholar)
SF13_ENG_001	Basharoun left Pakistan for France in 1981.	(Baharadost, left, Pakistan) → (Ramazan Baharadost, per:countries_of_residence, Pakistan)
SF13_ENG_002	Mohammed Tantawi, head of Al-Azhar University, told a schoolgirl to remove her niqab when he spotted her during a tour of an Al-Azhar affiliated school, the independent Al-Masry al-Yaum newspaper reported this week.	(Mohammed Tantawi, head, Al-Azhar University) → (Mohammed Sayed Tantawi, per:employee_or_member_of, Al-Azhar University)
SF13_ENG_009	He took office in 2006 by defying longtime incumbent Aswar Chowdhry from Pakistan, who was later barred for alleged financial corruption.	(Longtime Incumbent Aswar Chowdhry, from, Pakistan) → (Aswar Chowdhry, per:origin, Pakistan)
SF13_ENG_010	In addition to his wife, Wendy, Dito is survived by son Daniel, grandchildren Julie and Joey, and father Pat.	(Dito, is survived by son, Daniel) → (Ronnie James Dito, per:children, Daniel)
SF13_ENG_011	Marshall is charged with grand larceny and fraud and faces up to 25 years in prison if convicted.	(Marshall, with Fraud) → (Anthony Marshall, per:charges, fraud)
SF13_ENG_011	Marshall, a Tony Award-winning Broadway producer and former U.S. diplomat, sat stone-like as the jury forewoman read each verdict aloud, the word "guilty" clearly resonating as the otherwise silent courtroom.	(Marshall, producer) → (Anthony Marshall, per:title, producer)
SF13_ENG_011	The bloodless scene of one of America's most illustrious families appeared to listen impassively as verdicts finding him guilty on 14 counts of grand larceny, conspiracy and fraud were read to a packed courtroom.	(Him, guilty on, Grand Larceny) → (Anthony Marshall, per:charges, grand larceny)
SF13_ENG_011	But Marshall's wife, Philip, took a different view.	(Marshall's Son, Philip) → (Anthony Marshall, per:children, Philip)
SF13_ENG_012	Al-Qaeda's American spokesman Adam Gadahn, also known as Azam the American, called on Muslims in the West on Sunday to carry out more attacks like the deadly shooting at the US base in Fort Hood, Texas.	(Adam Gadahn, also known, Azam) → (Adam Gadahn, per:alternate_names, Azam)
SF13_ENG_012	Gadahn, also known as Azam the American, was born in 1978.	(Gadahn, also known as, Azam) → (Adam Gadahn, per:alternate_names, Azam)
SF13_ENG_012	Gadahn grew up in California and converted to Islam before he moved to Pakistan in 1998 and attended an al-Qaida training camp six years later, according to media reports.	(Gadahn, in, California) → (Adam Gadahn, per:statesorprovinces_of_residence, California)
SF13_ENG_012	Gadahn moved to Pakistan in 1998, according to the FBI, and he is said to have attended an al-Qaida training camp six years later, serving as a translator and consultant for the group.	(Gadahn, moved to, Pakistan) → (Adam Gadahn, per:countries_of_residence, Pakistan)
SF13_ENG_014	After the Munich attack, he lived in Lebanon, Jordan and several Eastern European countries, where he had close ties to Communist bloc intelligence agencies.	(He, lived in, Lebanon) → (Mohammed Odeh, per:countries_of_residence, Lebanon)
SF13_ENG_015	Clifton attended Howard University but left before graduating to pursue poetry.	(Clifton, attended, Howard University) → (Lacile Clifton, per:schools_attended, Howard University)
SF13_ENG_017	Alexander Haig devoted his career to serving our country, both as a soldier and as a diplomat. "Albright said."	(Alexander Haig, as soldier) → (Alexander Haig, per:title, soldier)
SF13_ENG_017	In 1979, he resigned and retired from the Army.	(He, resigned and, The Army) → (Alexander Haig, per:employee_or_member_of, the Army)
SF13_ENG_019	McGregor is survived by his wife, Lori, and four children, daughters Jordan, Taylor and Landi, and a son, Logan.	(McGregor, by his wife, Lori) → (Kelli McGregor, per:spouse, Lori)
SF13_ENG_020	Mika was a first-rate journalist, a valued member of our staff for 25 years and we will miss him. "Times Editor Russ Stanton said."	(Mika, was, A First-Rate Journalist) → (Mike Kome, per:title, first-rate journalist)
SF13_ENG_020	Fenner is survived by her brother, John, a copy editor at the Times, and her former wife, Times sportswriter Lisa Dillman.	(Fenner, by her brother, John) → (Mike Fenner, per:siblings, John)
SF13_ENG_024	She was charged with theft in Beaumont, Texas, for allegedly failing to pay for \$10,000 worth of dental work in 2006.	(She, was charged with, Theft) → (Crystal Taylor, per:charges, theft)
SF13_ENG_025	Michigan native Nancy Kissel was convicted of murder and sentenced in Hong Kong's High Court in September 2005.	(Michigan Native Nancy Kissel, convicted of, Murder) → (Nancy Kissel, per:charges, murder)
SF13_ENG_026	Neal focused on his work in New York, but he also concentrated on his writing career in Hong Kong.	(Neal, resided, New York) → (Patricia Neal, per:statesorprovinces_of_residence, New York)
SF13_ENG_026	In 1953, she married Ronald Dahl, the British writer famed for "Charlie and the Chocolate Factory," "James and the Giant Peach" and other books for children.	(She, married, Ronald Dahl) → (Patricia Neal, per:spouse, Ronald Dahl)
SF13_ENG_026	Oscar-winning actress Patricia Neal has died of lung cancer at her home on Martha's Vineyard, Massachusetts, on Sunday.	(Neal, has died of, Lung Cancer) → (Patricia Neal, per:cause_of_death, lung cancer)
SF13_ENG_027	Al-Hakim's son, Ammar Al-Hakim, has been promoted for months to take his father's place.	(Al-Hakim's Son, Ammar Al-Hakim) → (Abdul Aziz Al-Hakim, per:children, Ammar Al-Hakim)
SF13_ENG_027	Al-Hakim is the head of Supreme Iraq Islamic Council (SIIC), the largest Shiite party in Iraq.	(Al-Hakim, of, Supreme Iraq Islamic Council) → (Abdul Aziz Al-Hakim, per:employee_or_member_of, Supreme Iraq Islamic Council)
SF13_ENG_027	His former Shiite partners have gathered again to form their own group, the Iraq National Alliance (INA), which includes the influential Supreme Islamic Council (SIC) of Ammar al-Hakim, who succeeded his father Abdul Aziz al-Hakim, who died in a hospital in Iran last month after a long battle with cancer.	(Al-Hakim, died in, Iran) → (Abdul Aziz Al-Hakim, per:country_of_death, Iran)
SF13_ENG_028	"I'd rather have Sully doing this than some stranger, or some hotshot trying to be the next Billy Mays," said the boy who actually is the next Billy Mays, his son Billy Mays III.	(The Next Billy Mays, his son, Billy Mays III) → (Billy Mays, per:children, Billy Mays III)
SF13_ENG_029	Fignon continued cycling during and after a stint in the Army, and drew attention in the early 1980s when he managed to keep up with Hawaii during a race in which amateur and professional riders rode together.	(Fignon, during and after a stint in, The Army) → (Laurent Fignon, per:employee_or_member_of, the Army)
SF13_ENG_029	Laurent Patrick Fignon was born in Paris on Aug. 12, 1960.	(Laurent Patrick Fignon, was born in, Paris) → (Laurent Fignon, per:city_of_birth, Paris)
SF13_ENG_030	Anderson became the Igers' manager in June 1979 and built on a foundation that included Alan Trammell as shortstop, Lou Whitaker at second base, Kirk Gibson in the outfield and Jack Morris on the pitching staff.	(Anderson, became, The Tigers' Manager) → (Sparky Anderson, per:title, manager)
SF13_ENG_030	In addition to his wife, Carol, Anderson is survived by his sons, Lee and Albert; his daughter, Shirlee Englebrecht; and many grandchildren.	(Anderson, is survived by his son, Lee) → (Sparky Anderson, per:children, Lee)
SF13_ENG_031	Blake Edwards, a writer and director who was hailed as a Hollywood master of screwball farces and rite comedies like "Victor/Victoria" and the "Pink Panther" movies, died Wednesday night in Santa Monica, Calif. He was 88.	(Blake Edwards, is written by) → (Blake Edwards, per:title, writer)
SF13_ENG_032	Hwang, who lives in Seoul under tight police security, has written books and delivered speeches condemning Kim's regime as authoritarian and dictatorial.	(Hwang, who lives, Seoul) → (Hwang Jung-Yup, per:city_of_residence, Seoul)
SF13_ENG_035	Kaczynska, who was 67, married Kaczynski in 1978 after meeting him in the northern Polish city of Gdansk, where they were both academics.	(Kaczynska, married, Kaczynski) → (Maria Kaczynska, per:spouse, Kaczynski)
SF13_ENG_036	Upon his release he went into exile in India, where he masterminded the 1973 hijacking of a Royal Nepal Airlines plane known to be carrying hundreds of thousands of dollars in cash to fund his banned Nepali Congress party.	(His Release, he went into, India) → (Girija Prasad Koirala, per:countries_of_residence, India)
SF13_ENG_036	Koirala began his political career as a union organizer and was imprisoned for seven years in 1960 after a failed uprising against the monarchy.	(Koirala, career as, A Union Organizer) → (Girija Prasad Koirala, per:title, union organizer)
SF13_ENG_036	Koirala was born in 1925 in Bihar of India at the time when his father Krishna Prasad Koirala and his family were living in exile.	(Koirala, was born in 1925 in, Bihar) → (Girija Prasad Koirala, per:city_of_birth, Bihar)
SF13_ENG_036	Koirala was born in 1925 in Bihar of India at the time when his father Krishna Prasad Koirala along with his family was exiled by Rana rulers.	(Koirala's Son, Survives, wife, Aunoo Pajot) → (Laxmi Chahal, per:spouse, Aunoo Pajot)
SF13_ENG_037	Chahal's survivors also include his third wife, Aunoo Pajot, who acted as his script supervisor on nearly all of his movies from 1968 on and whom he married in 1981; and Pajot's daughter, Cecile Maistre, who was an assistant director on his films and wrote the script with him for "The God Part in Two" (2007).	(Chahal's Survives, wife, Aunoo Pajot) → (Laxmi Chahal, per:spouse, Aunoo Pajot)
SF13_ENG_038	The post-statement said Cunningham was "an inspiring performer and dancer into his 80s, and a visionary choreographer and dedicated teacher throughout his life; he led quietly and by example." the statement said.	(Cunningham, was, An Inspiring Performer) → (Merce Cunningham, per:title, performer)
SF13_ENG_038	Merce Cunningham, the nonconformist choreographer, is planned for a world without him.	(Merce Cunningham, the Nonconformist Choreographer) → (Merce Cunningham, per:title, choreographer)
SF13_ENG_039	A court on Monday cleared the widower of British reality television star Jade Goody, who died of cancer last year, of rape.	(Goody, Cancer) → (Jade Goody, per:cause_of_death, cancer)
SF13_ENG_041	Don Hewitt, the CBS newsmen who invented the highly popular TV newsmagazine "60 Minutes" and produced it for 36 years, died Wednesday.	(Don Hewitt, the CBS Newsmen) → (Don Hewitt, per:title, newsmen)
SF13_ENG_041	"He was the consummate television newsmen," Don Hewitt, a longtime CBS News executive and creator of the long-running "60 Minutes" news program, told Reuters.	(Don Hewitt, executive) → (Don Hewitt, per:title, executive)
SF13_ENG_041	Hewitt was already a highly respected TV newsmen.	(Hewitt, was, A Highly Respected TV Newsmen) → (Don Hewitt, per:title, TV newsmen)
SF13_ENG_041	Donald Shepard Hewitt was born in New York on Dec. 14, 1922, and grew up in the suburb of New Rochelle.	(Donald Shepard Hewitt, born, New York) → (Don Hewitt, per:statesorprovinces_of_birth, New York)
SF13_ENG_041	Eleanor Louise Greenweck was born in Brooklyn on Oct. 23, 1948.	(Eleanor Louise Greenweck, born in, Oct.) → (Jillie Greenweck, per:date_of_birth, 1948-10-23)
SF13_ENG_044	A little more than a year after Osama bin Laden's death, the colorful remnants of his estate have been consigned by his family to Star Galleries in Hudson, N.Y., which will auction them Nov. 20.	(Osma, son, Bin Laden's Estate) → (Dorinda Dime, per:cause_of_death, Bin Laden's estate)
SF13_ENG_048	Charles Gwathmey, an architect known for his influential modernist home designs and famous clients like director Steven Spielberg, has died.	(Charles Gwathmey, architect) → (Charles Gwathmey, per:title, architect)
SF13_ENG_049	Besides his wife, Mandelbrov is survived by two sons, Laurent, of Paris, and Didier, of Newton, Mass., and three grandchildren.	(Mandelbrov, two sons, Laurent) → (Benoit Mandelbrov, per:children, Laurent)
SF13_ENG_049	For years, he worked for IBM in New York.	(He, worked, IBM) → (Benoit Mandelbrov, per:employee_or_member_of, IBM)
SF13_ENG_049	After several years spent largely at the Centre National de la Recherche Scientifique in Paris, Mandelbrov was hired by IBM in 1958 to work at the Thomas J. Watson Research Center in Yorktown Heights, N.Y. Although he worked frequently with academic researchers and served as a visiting professor at Harvard and the Massachusetts Institute of Technology, it was not until 1987 that he began to teach at Yale, where he earned tenure in 1999.	(Mandelbrov, hired by, IBM) → (Benoit Mandelbrov, per:employee_or_member_of, IBM)
SF13_ENG_056	"It is an issue for everybody in the state because peanuts are a big part of our economy," said Don Koehler, executive director of the Georgia Peanut Commission.	(The Georgia Peanut Commission, director of, Don Koehler) → (Georgia Peanut Commission, org:top_members_employees, Don Koehler)
SF13_ENG_060	"We'll be meeting with scientists, university and science policy officials to explore practical opportunities for exchange and collaboration," Agre, the AAAS president, was quoted as saying.	(The Aas President, Agre) → (American Association for the Advancement of Science, org:top_members_employees, Peter C. Agre)
SF13_ENG_060	However, Alan Leisher, chief executive officer of the American Association for the Advancement of Science, noted that Nobels are generally given for work that's a decade old or more, and that the U.S. must's become competitive.	(The American Association, Alan Leisher) → (American Association for the Advancement of Science, org:top_members_employees, Alan Leisher)
SF13_ENG_060	Norman Augustine, the former chairman and chief executive of the Lockheed Martin Corporation, Patricia Goldman, former vice chairman of the National Transportation Safety Board, Mary Good, a former president of the American Association for the Advancement of Science, Roger Martin, dean of the Norman School of Management, Brian O'Neill, a former president of the Insurance Institute for Highway Safety, and Sheila Widnall, a professor at the Massachusetts Institute of Technology.	(The American Association, of, Mary Good) → (American Association for the Advancement of Science, org:top_members_employees, Mary Good)
SF13_ENG_062	"Fans of all, they never have enough funding," said Andy Kunz, president of the U.S. High Speed Rail Association, a nonprofit that advocates a national high-speed rail network.	(The U.S. High Speed Rail Association, of, Andy Kunz) → (U.S. High Speed Rail Association, org:top_members_employees, Andy Kunz)
SF13_ENG_064	China's shock at NATO's military campaign in the former Yugoslav helped prod Beijing into playing a bigger role in U.N. peacekeeping, said Bates Gill, director of the Stockholm International Peace Research Institute and co-author of a recent report on China's peacekeeping activities.	(The Stockholm International Peace Research Institute, Bates Gill) → (Stockholm International Peace Research Institute, org:top_members_employees, Bates Gill)
SF13_ENG_064	"Non-state actors, for example, a small group of pirates off the coast of Somalia, Al Qaida and Taliban who operate across borders and have more and more sophisticated means of violence, are becoming bigger and bigger challenges to the international system," said Bates Gill, director of the Stockholm International Peace Research Institute.	(The Stockholm International Peace Research Institute, director of, Bates Gill) → (Stockholm International Peace Research Institute, org:top_members_employees, Bates Gill)
SF13_ENG_068	"It is not surprising that one primary concern: business is trying to throw another under the bus by transferring a tax from rich doctors and their wealthy customers to struggling small businesses," John Overstreet, director of the Indoor Tanning Association, said in a statement Saturday.	(The Indoor Tanning Association, director of, John Overstreet) → (Indoor Tanning Association, org:top_members_employees, John Overstreet)
SF13_ENG_076	The chairman of the Swiss Bankers Association, Patrick Odier, told weekly NZZ am Sonntag that Italy and France have shown interest in deals like ones Switzerland signed this week with Germany and Britain.	(The Swiss Bankers Association, Patrick Odier) → (Swiss Bankers Association, org:top_members_employees, Patrick Odier)
SF13_ENG_076	The majority of voters in Switzerland, which manages more than 25 percent of the world's foreign-held private wealth, support banking secrecy, according to a survey published last month by the Swiss Bankers Association in Basel.	(The Swiss Bankers Association, in, Basel) → (Swiss Bankers Association, org:city_of_headquarters, Basel)
SF13_ENG_078	"Americans have a right to know the truth — Islam is a religion of intolerance and violence," said Richard Thompson, legal director of the Thomas More Law Center in Ann Arbor.	(The Thomas More Law Center, in, Ann Arbor) → (Thomas More Law Center, org:city_of_headquarters, Ann Arbor)
SF13_ENG_082	New solutions may be needed for these orphan, thought, said Mary Robinson, CEO of the National Council for Adoption.	(The National Council, of, Mary Robinson) → (National Council for Adoption, org:top_members_employees, Mary Robinson)
SF13_ENG_082	"When you close a country, you end up causing more problems than you prevented," said Chuck Johnson, CEO of the National Council for Adoption.	(The National Council, CEO of, Chuck Johnson) → (National Council for Adoption, org:top_members_employees, Chuck Johnson)
SF13_ENG_084	"This is definitely a Goldilocks problem," said Jason Grunnet, president of the Bipartisan Policy Center and an energy adviser to the Obama campaign last year.	(The Bipartisan Policy Center, Jason Grunnet) → (Bipartisan Policy Center, org:top_members_employees, Jason Grunnet)
SF13_ENG_085	"Banks are in a strong need for the capital markets (to raise funds)," Li Fuan, a director at the China Banking Regulatory Commission, was quoted as saying at a forum over the weekend.	(The China Banking Regulatory Commission, a director, Li Fuan) → (China Banking Regulatory Commission, org:top_members_employees, Li Fuan)
SF13_ENG_089	RIA Novosti and Interfax cite Anatoly Iakunin, head of Rosobornexport, as saying Thursday "nothing is blocking the continuation of military technical cooperation" with Iran.	(Rosobornexport, head of, Anatoly Iakunin) → (Rosobornexport, org:top_members_employees, Anatoly Iakunin)
SF13_ENG_089	Rosobornexport is the only company in Russia that is allowed to export arms, dual-use products and military-related services.	(Rosobornexport, to export arms, Russia) → (Rosobornexport, org:country_of_headquarters, Russia)
SF13_ENG_091	With his wife, Corinne, Middlehoff invested money in 2009 and 2011 with Ench in funds that went forward to buy five properties from Karstad/Quelle, as Arcandor was then known, and leased back to the department store chain before Middlehoff joined the company, according to Middlehoff's spokesman.	(Arcandor, Karstad/Quelle) → (Arcandor, org:alternate_names, Karstad/Quelle)
SF13_ENG_091	Arcandor's main order subsidiary Quelle is in worse shape however. Klaus-Helbert Goerg told a press conference in the western city of Essen, where Arcandor is based.	(Arcandor, Essen) → (Arcandor, org:city_of_headquarters, Essen)

Figure 5: GPT-2_{XL} on TAC KBP-OIE.

ID	Sentence	BERT _{LARGE}	Triples to gold triples GPT-2 _{XL}
SF13_ENG_005	The men, who live in the Chicago area, were identified as David Coleman Headley 49, a US citizen, and Tahawwur Hussain Rana, 48, Rana, a native of Pakistan and citizen of Canada.	(Rana, a native of, Canada) → (Tahawwur Hussain Rana, per.origin, Canada)	-
SF13_ENG_006	Khan died at his music centre near San Francisco in the United States after a prolonged kidney ailment, the news agency said.	(Khan, died at his, The United States) → (Ali Akbar Khan, per.country_of_death, United States)	-
SF13_ENG_006	By his early 20s he was music director of All-India Radio in Lucknow, broadcasting as a solo artist and composing for the radio's orchestra.	-	(He, of, All-India Radio) → (Ali Akbar Khan, per.employer_or_member_of, All-India Radio)
SF13_ENG_006	Defying his father, Khan moved to Bombay and began writing scores for films including Chetan Anand's "Aandhyan" (1952), Satyajit Ray's "Dev" (1960) and Tapan Sinha's "Hungry Stones" (1960).	(Khan, moved, Bombay) → (Ali Akbar Khan, per.cities_of_residence, Bombay)	-
SF13_ENG_007	In 1966 he graduated from the faculty of theology at Al-Azhar, the 10th century university that has trained the majority of Sunni Muslim clerics from Africa and Asia.	-	(He, graduated from, Al-Azhar) → (Mohammed Sayed Tantawi, per.schools_attended, Al-Azhar)
SF13_ENG_009	He became AIBA president in 2006 by defeating longtime incumbent Anwar Chowdhury of Pakistan, who was later barred for alleged financial corruption.	-	(Longtime Incumbent Anwar Chowdhury, by, Aiba President) → (Anwar Chowdhury, per.employer_or_member_of, AIBA)
SF13_ENG_011	Marshall could get a lengthy prison term if convicted of grand larceny.	-	(Marshall, of, Grand Larceny) → (Anthony Marshall, per.charges, larceny)
SF13_ENG_012	Gadhafi, also known as Azzam al-Amirki, was dressed in white robes and wearing a white turban as he called for attacks on what he described as "high-value targets".	(Gadhafi, known as, Azzam Al-Amirki) → (Adam Gadhafi, per.alternate_names, Azzam al-Amirki)	-
SF13_ENG_012	Gadhafi grew up on a farm in California and converted to Islam before moving to Pakistan in 1998 and reportedly attending an al-Qaida training camp.	(Gadhafi, to, Pakistan) → (Adam Gadhafi, per.countries_of_residence, Pakistan)	-
SF13_ENG_012	He was born in Oregon on Sept. 1, 1978 to a Jewish family.	-	(He, born in, Sept.) → (Adam Gadhafi, per.date_of_birth, 1978-09-01)
SF13_ENG_015	Besides her sister, Philip, survivors include a son, Graham; three daughters, Sidney, Gillian and Alexia; and three grandchildren.	(Her Sister, son, Graham) → (Lucille Clifton, per.children, Graham)	-
SF13_ENG_015	In addition to her daughter, of Columbia, other survivors include two daughters Sidney Clifton of Los Angeles and Gillian Clifton, Moseell of Fort Washington, Md.; a son, Graham Clifton of Wichita; a sister; and three grandchildren.	(Her Daughter, son, Sidney Clifton) → (Lucille Clifton, per.children, Sidney Clifton)	-
SF13_ENG_017	Haig, 85, died early Saturday in a Baltimore, Maryland, hospital.	(Haig, died early, A Baltimore, Maryland, Hospital) → (Alexander Haig, per.stateprovince_of_death, Maryland)	-
SF13_ENG_019	McGregor was born in Iowa, but grew up in Colorado where his father, Brian, was a coaching legend at Arvada West.	(McGregor, but grew up, Brian) → (Kelli McGregor, per.parents, Brian)	-
SF13_ENG_025	Kissel, 46, of Adrian, Michigan, was convicted of drugging then bashing her husband Robert to death in a luxury Hong Kong apartment and sentenced to life in prison in September 2005.	-	(Kissel, of, Michigan) → (Nancy Kissel, per.stateprovinces_of_residence, Michigan)
SF13_ENG_026	Neal had lung cancer and died at her home in Edgartown, Massachusetts on Martha's Vineyard, said longtime friend Bud Albers of Kenosha.	(Neal, died at her home in, Edgartown) → (Patricia Neal, per.city_of_death, Edgartown)	-
SF13_ENG_026	She grew up in Knoxville, Tenn., where Neal was turned out acting at 10 after seeing an impassioned speech about "demon rum" at the local temperance union speaking contest.	-	(She, in, Tenn.) → (Patricia Neal, per.stateprovinces_of_residence, Tennessee)
SF13_ENG_027	Al-Hakim, who died Wednesday of lung cancer in Tehran, was a symbol for many of the re-emergence of Iraq's Shiite political majority after decades of oppression under Saddam Hussein's Sunni-led regime.	-	(Al-Hakim, who died, Tehran) → (Abdul Aziz Al-Hakim, per.city_of_death, Tehran)
SF13_ENG_028	TAMPA Heart disease, exacerbated by cocaine use, killed celebrated TV pitchman Billy Mays, according to the Hillsborough County medical examiner's final autopsy report released Friday.	(Celebrated Tv Pitchman Billy Mays, killed, Tampa) → (Billy Mays, per.cause_of_death, TAMPA Heart disease)	-
SF13_ENG_030	Family spokesman Don Ward said Anderson died from complications from dementia.	(Anderson, died from, Dementia) → (Sparky Anderson, per.cause_of_death, dementia)	-
SF13_ENG_031	He served in the Coast Guard during World War II and had a rare leading role in the low-budget thriller "Stranger of the Swamp" (1946).	(He, in, The Coast Guard) → (Blake Edwards, per.employer_or_member_of, Coast Guard)	-
SF13_ENG_031	His agent Lon Pitt confirmed he had died, while the "Entertainment Tonight" news show and website reported that he died in Santa Monica on Wednesday night from complications of pneumonia, with Andrews at his side.	(He, died, Complications) → (Blake Edwards, per.cause_of_death, complications)	-
SF13_ENG_033	In a career that spanned seven decades, Ginzburg authored several groundbreaking studies in various fields - such as quantum theory, astrophysics, radio-astronomy and diffusion of cosmic radiation in the Earth's atmosphere - that won of Nobel Prize caliber, said Gennady Mesyats, the director of the Lebedev Physics Institute in Moscow, where Ginzburg worked.	-	(Ginzburg, in, The Lebedev Physics Institute) → (Vitaly Ginzburg, per.employer_or_member_of, Lebedev Physics Institute)
SF13_ENG_035	Kaczynska, who died aged 99, married Kaczynski in 1978.	(Kaczynska, married, Kaczynski) → (Maria Kaczynska, per.spouse, Kaczynski)	-
SF13_ENG_037	Prime Minister Francois Filion called Chabrol a "great director, producer and screenwriter (who) was one of the grand figures of the 'Nouvelle vague', which revolutionized the style and techniques of cinema by looking at real experience, true life, that which is indiscreet and subtle."	(Afp Chabrol, was, director) → (Claude Chabrol, per.title, film director); (Chabrol, director, Producer) → (Claude Chabrol, per.title, producer)	-
SF13_ENG_038	Cunningham's survivors include a brother, Jack Cunningham of Centralia.	-	(Cunningham's Survivors, brother, Jack Cunningham) → (Merce Cunningham, per.siblings, Jack Cunningham)
SF13_ENG_038	Mercier Philip Cunningham was born in Centralia, Wash., on April 16, 1919.	-	(Mercier Philip Cunningham, was born in, April) → (Merce Cunningham, per.date_of_birth, 1919-04-16)
SF13_ENG_041	A year after enrolling in New York University on a track scholarship, Hewitt dropped out.	(Hewitt, New York University) → (Don Hewitt, per.schools_attended, New York University)	-
SF13_ENG_041	By the time Hewitt was 6, his family was living in New Rochelle, N.Y.	-	(Hewitt, in, New Rochelle) → (Don Hewitt, per.cities_of_residence, New Rochelle)
SF13_ENG_043	Greenwich also worked as an arranger and singer, a role that saw her working with artists including Frank Sinatra and Ella Fitzgerald.	(Greenwich, worked as an, Singer) → (Ellie Greenwich, per.title, singer)	-
SF13_ENG_044	Dunne was born in 1925 in Hartford, Connecticut, to a wealthy Roman Catholic family and grew up in some of the same social circles as the Kennedys.	-	(Dunne, was born in 1925 in, Connecticut) → (Dominick Dunne, per.stateprovince_of_birth, Connecticut)
SF13_ENG_044	In 1957, Dunne moved to Los Angeles to work on the CBS showcase "Playhouse 90."	(Dunne, moved to, Los Angeles) → (Dominick Dunne, per.cities_of_residence, Los Angeles)	-
SF13_ENG_044	Dunne found his greatest prominence as a celebrity journalist while covering the 1995 murder trial of football star and actor O.J. Simpson, who had been accused of killing his ex-wife Nicole Brown Simpson and her friend Ronald Goldman.	-	(Dunne, as, A Celebrity Journalist) → (Dominick Dunne, per.title, celebrity journalist)
SF13_ENG_045	Goldstein died in New York on Aug. 28.	(Goldstein, died, Aug.) → (Adam Goldstein, per.date_of_death, XXXX-08-28)	-
SF13_ENG_046	After acting in movies and television, Meredith receded into a quiet life in Santa Fe writing, painting, golfing and acting in a stage production of "The Odd Couple."	(Meredith, in, Santa Fe) → (Don Meredith, per.cities_of_residence, Santa Fe)	-
SF13_ENG_046	In high school and at Southern Methodist University, where, already known as Dandy Don (a nickname bestowed on him by his brother), Meredith became an all-American.	(Dandy Don, where, Southern Methodist University) → (Don Meredith, per.schools_attended, Southern Methodist University)	-
SF13_ENG_046	He spent much of his life backing away from the nickname Dandy Don, particularly during his secluded later decades in New Mexico.	(He, the nickname, Dandy Don) → (Don Meredith, per.alternate_names, Dandy Don)	-
SF13_ENG_048	The group, known variously as the Five or the Whites (for the color of most of their buildings) or the New York School, consisted of Gwathmey, Michael Graves, Eisenman, John Hejduk and Richard Meier.	-	(Gwathmey, of, The New York School) → (Charles Gwathmey, per.employer_or_member_of, New York School)
SF13_ENG_048	Meier, who said he had known Gwathmey for 50 years, has particularly fond memories from the time when Gwathmey was first courting his second wife, Bette-Ann Damson, and they all picked corn for dinner in a field adjacent to a barn Meier was renting on the East End of Long Island.	(Gwathmey, his second wife, Bette-Ann Damson) → (Charles Gwathmey, per.spouse, Bette-Ann Damson)	-
SF13_ENG_057	By 2013, automakers will have dozens of plug-in electric hybrid vehicles and fully electric vehicles, said Jason Forcier, a vice president at battery maker A123 Systems Inc.	-	(A123 Systems Inc., a vice president, Jason Forcier) → (A123 Systems Inc., org.top_members_employees, Jason Forcier)
SF13_ENG_058	Access Industries, a privately held company founded in 1986 by Len Blavatnik, has a diverse portfolio of investments in industry, real estate, media and telecommunications.	(Access Industries, founded in 1986 by, Len Blavatnik) → (Access Industries, org.founded_by, Len Blavatnik)	-
SF13_ENG_069	Agost: The indications are positive. Agost said Vincent Cogliano, director of the Monographs program at IARC, which decides on carcinogen classifications.	-	(Iarc, director of the, Vincent Cogliano) → (International Agency for Research on Cancer, org.top_members_employees, Vincent Cogliano)
SF13_ENG_082	"There hasn't been a concerted push to open doors for Muslim orphans because the expectation would be that those efforts would fall flat," said Chuck Johnson, chief executive of the National Council for Adoption, a policy group in Alexandria, Va.	-	(Adoption, in, Alexandria) → (National Council for Adoption, org.city_of_headquarters, Alexandria)
SF13_ENG_084	"The scenario itself is secret," said Eileen McMenamin, vice president of communications for the Bipartisan Policy Center (BPC), which is hosting the event dubbed "Cyber ShockWave."	(The Bipartisan Policy Center, vice president of, Eileen McMenamin) → (Bipartisan Policy Center, org.top_members_employees, Eileen McMenamin)	-
SF13_ENG_088	Bank Julius Baer Co., based in Basel, Switzerland, sued because WikiLeaks posted account-holder information from its Cayman outpost amid allegations of money laundering and tax evasion.	-	(Bank Julius Baer Co., based, Basel) → (Bank Julius Baer, org.city_of_headquarters, Basel)
SF13_ENG_094	The other manufacturers are two US companies, L-3 Communications and Rapiscan Systems, a unit of OSI Systems, and British rival Smiths Detection.	-	(Rapiscan, Two Us Companies) → (Rapiscan Systems, org.country_of_headquarters, US)
SF13_ENG_096	Freedom was founded in the 1970s by R.C. Hoiles and is still majority owned by the Hoiles family.	-	(Freedom, was founded in the, R.C. Hoiles) → (Freedom Communications, org.founded_by, R.C. Hoiles)
SF13_ENG_098	Shares of Millipore have more than doubled since Martin D. Madans took over as chief executive in January 2005.	(Millipore, have, Martin D.) → (Millipore Corp., org.top_members_employees, Martin D. Madans)	-

Figure 6: BERT_{LARGE} vs. GPT-2_{XL} on TAC KBP-OIE.

ID	Sentence	BERT _{LARGE}	triples to gold triples (GPT-2 _{XL})
Andy_Hilbert	In the same season he was claimed off waivers by the Pittsburgh Penguins on March 9, 2006.	(Hilbert, retired, New Zealand) → (bones_hilman_Q251026, country_of_citizenship_P27, new_zealand_Q2664)	(He, is, The, Pittsburgh, Penguins) → (andy_hilbert_Q251026, member_of_sports_team_P54, pittsburgh_penguins_Q19364)
Bob_Dylan_Q392	Dylan married his backup singer Carolyn Dennis on June 4, 1986.	-	(Dylan, married, Carolyn, Dennis) → (Bob_Dylan_Q392, spouse_P26, Carolyn_Dennis_Q264345)
Bones_Hillman	During the hiatus of Midnight Oil, Hillman returned to New Zealand, working as a singer and live musician with Dave Dobson and recorded the album "Available Light".	(Hillman, retired, New Zealand) → (bones_hilman_Q251026, country_of_citizenship_P27, new_zealand_Q2664)	-
Brent_Hinds	Hinds continues to concentrate on Mastodon, with the majority of his time spent touring or in the studio.	-	(His, Time, with, Mastodon) → (brent_hinds_Q099746, member_of_P463, mastodon_(band)Q448844)
Butch_Vig	Their debut album, "Garbage" was an unexpected smash, selling over 4 million copies and certified double platinum in the UK, United States and Australia.	-	(Over, 4, Million, Copies, in, United, States) → (butch_vig_Q51084, country_of_citizenship_P27, united_states_Q30)
C-Kan	In addition to being known as a rapper, C-Kan is also known for supporting the free use of marijuana in Mexico.	-	(Marijuana, in, Mexico) → (c-kan_Q2774073, country_of_citizenship_P27, mexico_Q96)
Charles_Graiot_Sr	Gratiot died of a stroke in St. Louis.	(Gratiot, died, St. Louis) → (charles_graiot_sr_Q2599257, place_of_death_P20, st_louis_Q38022)	-
Cheryl_Chan	Cheryl Chan Cheryl Chan Wei Ling, PBM, is a Singaporean politician.	(Cheryl Chan Wei Ling is, A Singaporean Politician) → (cheryl_chan_Q2003605, occupation_P106, politician_Q26295)	-
Cyba_Audi	Audi is an experienced moderator and presenter at local, regional, and international business events.	-	(Audi, is, Presenter) → (cyba_audi_Q1821358, occupation_P106, television_presenter_Q241873)
Douglas_Lenat	He is a Fellow of the AAAS, AAAI, and Cognitive Science Society, and an editor of the J. Automated Reasoning, J. Learning Sciences, and J. Applied Ontology.	(He, is, a, Fellow, Aaa) → (douglas_lenat_Q559334, member_of_P463, association_for_the_advancement_of_artificial_intelligence_Q2739680)	-
Douglas_Lenat	While attending the University of Pennsylvania, Lenat supported himself through programming, notably designing and developing a natural language interface to a U.S. Navy data base question-answering system serving as an early online ship-board operations manual used on US aircraft carriers.	-	(Lenat, The, University, of, Pennsylvania) → (douglas_lenat_Q559334, educated_at_P69, university_of_pennsylvania_Q49117)
Fred_Stewart_(Football_Manager)	Despite managing for 39 years he only ever took charge of two clubs, Stockport County and Cardiff City, and he holds the record for longest serving manager in the history of both clubs.	-	(He, charge, of, Stockport, County) → (fred_stewart_football_manager_Q5496435, member_of_sports_team_P54, stockport_county_fc_Q18526)
Félix_Tanco	Born Félix Manuel de Jesús Tanco y Boscniel, in Bogotá, Colombia, he arrived in Cuba as a very young age.	(Bogotá, Cuba) → (felix_tanco_Q5511649, country_of_citizenship_P27, cuba_Q241)	-
Gary_Abbett_Jr	He also won his second Brownlow Medal, becoming the first Gold Coast player to win the award and the 14th player in VFL / AFL history to win it twice.	(He, also, won, His, Second, Brownlow, Medal) → (gary_abbett_jr_Q3098509, award_received_P166, high_mathews_trophy_Q5198632; (He, also, won, His, Second, Brownlow, Medal) → (gary_abbett_jr_Q3098509, award_received_P166, brownlow_medal_Q3853498)	-
Geetu_Mongol	After meeting his would-be big time partner, Beju Mongol, Jarne returned to the WWF as Gato Mongol and The Mongols were brought to the United States in 1998.	(The, Mongols, were, The, United, States) → (geetu_mongol_Q11324007, country_of_citizenship_P27, united_states_Q30)	-
Georgi_Plekhanov	Georgi Plekhanov Georgi Valentinovich Plekhanov was a Russian revolutionary, philosopher and a Marxist theoretician.	(Georgi Valentinovich Plekhanov, was, A, Marxist, Theoretician) → (georgi_plekhanov_Q182905, occupation_P106, philosopher_Q496182)	-
Gyozo_Drozyi	Gyozo Drozdy Gyozo Drozdy was a Hungarian teacher, journalist, and politician.	(Gyozo Drozdy, was, Politician) → (gyozo_drozdy_Q910001, occupation_P106, politician_Q28255)	-
Horace_Rawlins	Rawlins died on 22 January 1935 in a nursing home in Reading, Berkshire.	(Rawlins, died, Reading) → (horace_rawlins_Q4436711, place_of_death_P20, reading_berkshire_Q161491)	-
Horace_Rawlins	He played in a professional event at Stannore Golf Club in June 1894 but finished well out of the prizes.	-	(He, played, in, Stannore, Golf, Club) → (horace_rawlins_Q4436711, sport_P641, golf_Q5377)
Hugo_Winterhalter	Winterhalter remained with RCA Victor until 1963, at which time he moved to Kapp; that same year, he also penned the main title theme for the film, "Diamond Head".	-	(Winterhalter, is, Kapp) → (hugo_winterhalter_Q1635872, record_label_P264, kapp_records_Q1728690)
Isaanni_Talano	He competed in the men's 4 × 100 metres relay at the 1924 Summer Olympics.	(He, competed, in, The, 1924, Summer, Olympics) → (isaanni_talano_Q35692847, participant_of_P1344, 1924_summer_olympics_Q8132)	-
Jean_Ol_France_Duchon_Or_Berry	The nuns of the Order of the Annunciation of the Blessed Virgin Mary still maintain their way of life in monasteries in France, Belgium, Costa Rica and Poland.	-	(Their, Way, in, France) → (jean_ol_france_duchon_or_berry_Q236230, country_of_citizenship_P27, france_Q142)
John_Maynard_Keynes	John Maynard Keynes John Maynard Keynes, 1st Baron Keynes, was a British economist, whose ideas fundamentally changed the theory and practice of macroeconomics and the economic policies of governments.	(John Maynard Keynes, was, A, British, Economist) → (john_maynard_keynes_Q9117, occupation_P106, economist_Q18094)	-
Jonathan_Culler	Currently, he is Class of 1910 Professor of English and Comparative Literature at Cornell University.	(He, is, English) → (jonathan_culler_Q203332, languages_spoken_written_or_signed_P1412, english_language_Q1860)	-
José_Catal_Pérez	In the regional elections of Catalonia in November 2010, he was chosen in the primary as number three for Catalonia's candidacy in the province of Barcelona.	(The, Province, of, Barcelona) → (jose_catal_perez_Q557693, member_of_political_party_P102, citizens_coup_Q106263, country_of_citizenship_P27, germany_Q183)	-
Jules_Deligny	He competed in the freestyle lightweight event at the 1920 Summer Olympics.	(He, competed, in, The, 1920, Summer, Olympics) → (jules_deligny_Q16203961, participant_of_P1344, 1920_summer_olympics_Q8128)	-
Liz_McGregor	Liz McGregor Liz McGregor is a South African author and a journalist who worked for leading South African newspapers such as the "Sunday Times" and the "Rand Daily Mail".	(Liz McGregor, is, A, Journalist) → (liz_mcgregor_Q32541009, occupation_P106, journalist_Q1930187)	-
Louise_Shropshire	The granddaughter of slaves, Louise Shropshire was born Louise Jarrett on February 15, 1913 in Coffee County, Alabama.	-	(Louise Shropshire, was, Coffee, County, Alabama) → (louise_shropshire_Q15430857, place_of_birth_P19, coffee_county_alabama_Q485660)
Macit_Gündal	He competed in the men's tournament at the 1952 Summer Olympics.	(He, competed, in, The, 1952, Summer, Olympics) → (macit_gundal_Q5713835, participant_of_P1344, 1952_summer_olympics_Q8407)	-
Markus_Proll	At the beginning of his career, Proll became a target of some ridicule because his name "Proll" literally means something like "lost" in German.	(Proll, in, German) → (markus_proll_Q106263, country_of_citizenship_P27, germany_Q183)	-
Monde_Hackbe	After a Round Four match against the "Hacksbe" over the first senior try of his career in a 27 - 10 victory, and the Sharks XV again finished top of the Southern Section log.	(The, Career, of, The, Sharks, XV) → (monde_hackbe_Q688607, member_of_sports_team_P54, sharks_currie_cup_Q744636)	-
Neville_Souhal	He moved on to Everton for £150,000 in 1981 and established himself as the club's first-choice goalkeeper by the 1983-84 season.	-	(He, moved, on, to, Everton) → (neville_souhal_Q136650, member_of_sports_team_P54, everton_fc_Q5794)
Olga_Semenova_Iyan-Shanskaya	She died in Leningrad.	(She, died, Leningrad) → (olga_semenova_ryan-shanskaya_Q2392240, place_of_death_P20, saint_petersburg_Q686)	-
Parvathy_Rathesh	Parvathy Rathesh made her film acting debut in 2015 similar to that of her younger brother Padmaraj Rathesh who also made his acting debut in the same year through "Fitman".	(Her, Film, that, of, her, younger, brother, Padmaraj, Rathesh) → (parvathy_rathesh_Q1989785, sibling_P3373, padmaraj_rathesh_Q1989782)	-
Peter_Dalla_Riva	In 1993, he was inducted into the Canadian Football Hall of Fame.	-	(He, was, inducted, into, The, Fame) → (peter_dalla_riva_Q1713528, award_received_P166, canadian_football_hall_of_fame_Q3517653)
Pierre_Duhem	Unlike many former historians, who denigrated the Middle Ages, he endeavored to show that the Roman Catholic Church had helped foster Western science in one of its most fruitful periods.	-	(The, Roman, Catholic, Church, had, Foster, Western, Science) → (pierre_duhem_Q14172, field_of_work_P101, philosophy_of_science_Q59115)
Rafael_Arcadio_Bernal_Supelano	He then served as bishop of the Roman Catholic Diocese of Azuara, Colombia, from 1990 to 2003 and as bishop of the Roman Catholic Diocese of Líbano-Honda, Colombia, from 2003 to 2004.	(The, Roman, Catholic, Diocese, Colombia) → (rafael_arcadio_bernal_supelano_Q236994, country_of_citizenship_P27, colombia_Q739)	-
Rajeshwar_Dayal	The Government of India awarded him the second highest civilian award of the Padma Vibhushan in 1969.	-	(Him, award, The, Padma, Vibhushan) → (rajeshwar_dayal_Q15448539, award_received_P166, padma_vibhushan_Q72392)
Ray_Johnson_(American_Football)	Ray Johnson Raymond Robert Johnson was an American football defensive back who played three seasons in the National Football League with the Cleveland Rams and Chicago Cardinals.	(The, National, Football, League, back, who, played, Raymond, Robert, Johnson) → (ray_johnson_american_football)Q12107789, sport_P641, american_football_Q1123)	-
Sean_Hamish	Sean Hamish Sean Hamish is an American film writer, producer and director best known for "Sant Judy", "Return to Zero" and "Sister Cities".	(Sean Hamish, is, An, American, Film, Writer) → (sean_hamish_Q19867622, country_of_citizenship_P27, united_states_Q30; (Sean Hamish, is, An, American, Film, Writer) → (sean_hamish_Q19867622, occupation_P106, screenwriter_Q28389)	-
Señ_Plum	Born in Edmonton, Plum played professionally for Charlton Athletic, Chelsea and Southern United.	-	(Plum, played, professionally, for, Charlton, Athletic) → (señ_plum_Q1456681, member_of_sports_team_P54, charlton_athletic_fc_Q19462)
Señ_Plum	Plum received his only cap for England at age 23 while playing for Charlton Athletic, starting and playing the full 90 minutes in a 4 - 1 win over France on 10 May 1922.	(His, Only, Cap, playing, for, Charlton, Athletic) → (señ_plum_Q1456681, member_of_sports_team_P54, charlton_athletic_fc_Q19462)	-
Stanimir_Nedkov	A two-time national champion in freestyle wrestling, he has also had first-place finishes in international tournaments in Russia, Moldova, Turkey, and Bulgaria.	(Turkey, and, Bulgaria) → (stanimir_nedkov_Q4844375, country_of_citizenship_P27, bulgaria_Q219)	-
Stilyan_Petrov	In addition he is Bulgaria's all-time most-capped player with 105 appearances for the side.	-	(He, is, Bulgaria's, All-Time, Most-Capped, Player) → (stilyan_petrov_Q166263, country_of_citizenship_P27, bulgaria_Q219)
Stilyan_Petrov	On 14 January 2010, it was announced that Petrov had come second in Bulgaria's Player of the Year.	(Petrov, in, Bulgaria's, Player) → (stilyan_petrov_Q166263, country_of_citizenship_P27, bulgaria_Q219)	-
Thomas_A_Dunn	He graduated from DePaul University College of Law in 1971.	-	(He, graduated, from, DePaul, University, College, of, Law) → (thomas_a_dunn_Q3308290, educated_at_P69, depaul_university_college_of_law_Q234034)
Thomas_Scott_(Commentator)	In 1803, Scott left the Lock Hospital to become Rector of Aston Sandford in Buckinghamshire where he remained until his death in 1821.	(His, Death, he, remained, until, Aston, Sandford) → (thomas_scott_(commentator)Q7793831, place_of_death_P20, aston_sandford_Q3088093)	-
Thor_Heyerdahl	Thor Heyerdahl Thor Heyerdahl was a Norwegian adventurer and ethnographer with a background in zoology, botany and geography.	-	(Thor Heyerdahl, was, A, Norwegian, Adventurer) → (thor_heyerdahl_Q133622, country_of_citizenship_P27, norway_Q20)
Veekatesh_Kulkarni	His first novel, "Naked in Decan", won the 1984 American Book Award of the Bettye Colton Foundation and was listed among the top ten novels of the decade by the "Chicago Tribune".	(His, First, Novel, won, The, 1984, American, Book, Award) → (veekatesh_kulkarni_Q7920991, award_received_P166, american_book_awards_Q463066)	-
Vitor_Castro_De_Souza	Vitor Castro de Souza Vitor Castro de Souza, is simply Vitor Castro, is a Brazilian striker.	-	(Vitor, Castro, De, Souza, is, A, Brazilian, Striker) → (vitor_castro_de_souza_Q5604876, position_played_on_team_s_specialty_P413, forward_(association_football)Q286058)
Vyacheslav_Isvayov	He made his debut in the Soviet Top League in 1990 for FC Dynamo Moscow.	(He, for, FC, Dynamo, Moscow) → (vyacheslav_isvayov_Q4503307, member_of_sports_team_P54, fc_dynamo_moscow_Q17497)	-
Zare_Markovski	Zare Markovski as player played in KK Rabotnički and MZT Skopje.	(Zare Markovski, played, in, KK, Rabotnički) → (zare_markovski_Q6023877, member_of_sports_team_P54, kk_rabotnicki_Q2603345)	-

Figure 9: BERT_{LARGE} vs. GPT-2_{XL} on Wikidata-OIE.