

LSOIE: A Large-Scale Dataset for Supervised Open Information Extraction

Jacob Solawetz
Roboflow, Inc.
Minneapolis, MN, USA
jacob@roboflow.ai

Stefan Larson
Rosegold AI
Ann Arbor, MI, USA
stefan@rosegold.ai

Abstract

Open Information Extraction (OIE) systems seek to compress the factual propositions of a sentence into a series of n -ary tuples. These tuples are useful for downstream tasks in natural language processing like knowledge base creation, textual entailment, and natural language understanding. However, current OIE datasets are limited in both size and diversity. We introduce a new dataset by converting the QA-SRL 2.0 dataset to a large-scale OIE dataset (LSOIE). Our LSOIE dataset is 20 times larger than the next largest human-annotated OIE dataset. We construct and evaluate several benchmark OIE models on LSOIE, providing baselines for future improvements on the task. Our LSOIE data, models, and code are made publicly available.¹

1 Introduction

Open Information Extraction (OIE) (Banko et al., 2007) aims to automatically extract all factual propositions of a sentence into a series of n -ary tuples. For example, the sentence “the cook baked and ate the cake” would produce two extractions representing the two basic propositions of the sentence: (the cook, **ate**, the cake) and (the cook, **baked**, the cake). In OIE, extraction arguments are required to be contiguous spans from the sentence and the resulting tuple should be intelligible as natural text when read in order. The schema-free nature of OIE provides a flexible framework in which to capture semantic relations between entities in natural language text. Open Information Extraction tuples are useful to a variety of downstream tasks including knowledge base creation (Zhang et al., 2019), textual entailment (Levy et al., 2014), and other natural language understanding tasks (Mausam, 2016).

¹Our LSOIE dataset, models, and code can be found at <https://github.com/Jacobsolawetz/large-scale-oie>.

	Domains	#Sent.	#Ext.
OIE2016	Wiki, Newswire	3,180	8,477
AW-OIE	Wiki, Wikinews	3,300	17,165
LSOIE-wiki	Wiki, Wikinews	24,296	56,662
LSOIE-sci	Science	47,998	97,550
	Ext. / Sent.	Vocab	Ordered
OIE2016	2.7	13,863	
AW-OIE	5.2	15,853	
LSOIE-wiki	2.3	46,617	✓
LSOIE-sci	2.0	51,668	✓

Table 1: OIE dataset metrics. Our new dataset LSOIE has substantially more text available than prior work, and includes a new science domain. Our dataset conversion process leverages the scope of the QA-SRL 2.0 bank and improves upon previous methodology.

Open Information Extraction relations may be explicitly stated by verbal predicates, or implicitly stated through nominalizations. In this paper, we focus only on explicit extractions. With the original goal of OIE as web scale information extraction (Banko et al., 2007), an OIE system can focus solely on explicit extractions because the redundancy of language will inevitably display implicit information elsewhere.

The interest in OIE has grown: both in terms of the types of models that can be applied to tackle OIE (Cui et al., 2018; Stanovsky et al., 2018; Jiang et al., 2019), and in terms of the downstream applications to which OIE can be applied (Mausam, 2016; Zhang et al., 2019). As the interest in OIE grows, however, so too should the scale of the corpora available for training and evaluating OIE models.

In this paper, we expand the reach and quality of OIE data by developing a new dataset, LSOIE, which is built by converting the QA-SRL BANK 2.0 dataset (FitzGerald et al., 2018) to the task of OIE. Our new dataset contains almost ten times as many extractions and about 20 times as many sentences as previous OIE datasets built from human anno-




Figure 1: An example annotated sentence from QA-SRL 2.0 (FitzGerald et al., 2018). In this case, the annotations are derived from the question and answers: - Where does someone **provide** something? **In Asian countries**. Who **provides** something? **physicians**. What is being **provided**? **drugs**. The extracted tuple in our new LSOIE dataset is (physicians, **provide**, drugs, in Asian countries).

tations (see Table 1). We benchmark LSOIE with several models, providing baseline results for future research. Our LSOIE dataset, models, and code are publicly available.

2 Background

2.1 OIE Datasets

Available OIE corpora fall into three categories: (1) converted from crowdsourcing, (2) model-derived, and (3) directly crowdsourced.

Converted from crowdsourcing: Stanovsky and Dagan (2016) created the OIE2016 dataset by converting the crowd-annotated QA-SRL (He et al., 2015) dataset’s question-answer pairs to OIE extraction relations. Similarly, Stanovsky et al. (2018) generated the AW-OIE dataset by converting the crowd-annotated Question Answer Meaning Representation (QAMR) dataset’s question-answer pairs.

The OIE2016 and AW-OIE datasets were the first datasets used for supervised OIE. These datasets provided the basis for supervised approaches in NLP, but they are small and extractions lack accuracy, as they are converted in the order that question answer pairs appear in the base dataset.

Model-derived: Cui et al. (2018) and Jia et al. (2018) generate large derivative training datasets by running rules-based models and keeping high confidence extractions for downstream tasks. Similarly, Gashteovski et al. (2019) introduce the largest OIE dataset to date (over 340M triples) by deriving extractions from `MinIE` Gashteovski et al. (2017) with the goal of automatically constructing a knowledge base. While model-derived datasets are useful for knowledge base construction, using them for downstream tasks teaches the new model to replicate the behavior of the original, often noisy, base model.

Directly crowdsourced: Bhardwaj et al. (2019) point out that the evaluation framework used in Stanovsky and Dagan (2016) is rather noisy and the tuple matching algorithm is overly lenient because it only looks at lexical overlap for the whole extraction, ignoring the ordering of arguments. Bhardwaj et al. (2019) provide an alternative evaluation set

that has been crowdsourced specifically for OIE, annotating 1,282 sentences. While this dataset is useful for the evaluation of OIE systems, its format differs from other work in OIE - the predicate entry in CARB (Bhardwaj et al., 2019) tuples contains context that is often broken into separate tuples by other OIE systems.

2.2 The QA-SRL Bank 2.0

In QA-SRL, each predicate-argument relationship in a sentence is labeled manually with a question-answer pair. FitzGerald et al. (2018) design a large-scale crowdsourcing annotation pipeline to incentivize extensive and accurate coverage. Relative to the original QA-SRL annotations (He et al., 2015), which were collected from 10 hired freelance workers, the new QA-SRL dataset achieves similar precision (95.7% versus 97.5%) and lower recall (72.4% versus 86.6%). Relative to Propbank (Palmer et al., 2005), an expert annotation system designed to capture all semantic roles in a sentence, the QA-SRL 2.0 authors find that their work 95% precision and 85% recall. FitzGerald et al. (2018) then build a supervised QA-SRL parser and extend the reach of their dataset by over-generating new candidate question-answer pairs and passing them through their validation process.

The QA-SRL paradigm is well-suited to be a precursor to OIE extractions, as it captures predicate-argument relations in a schema-free way.

3 The LSOIE Dataset

Our work expands upon and addresses the shortcomings present in Stanovsky and Dagan (2016) and Stanovsky et al. (2018). We apply a similar conversion processes used for OIE2016 on the QA-SRL BANK 2.0 dataset. In addition, we implement novel conversion heuristics to ensure data quality and order arguments. The result is LSOIE, an OIE dataset that is much larger and diverse than prior work.




Figure 2: The distribution of token level tags (listed clockwise) in the LSOIE dataset. P denotes the extraction’s predicate, A_0 - A_N denote the extraction’s arguments, and O denotes that a given token does not belong to the extraction.

3.1 LSOIE Conversion Process

We produce LSOIE via conversion from QA-SRL in the same manner as Stanovsky and Dagan (2016), with several important changes to adapt their method to the QA-SRL BANK 2.0.

A QA-SRL annotation for a predicate p consists of a list of questions $Q = \{q_0, \dots, q_n\}$, and a set of answer spans $A_i = \{a_{i0}, \dots, a_{in_i}\}$ for each question q_i . For each tuple (a_0, \dots, a_k) in the Cartesian product $\times_i^n A_i$, we produce the extraction tuple $(a_0, p, a_1, \dots, a_k)$.

In our example extraction in Figure 1, the target predicate p is *provide*. The list of questions Q is [*Where does someone provide something?*, *Who provides something?*, *What is being provided?*] The list of arguments A is [*In Asian countries*, *physicians*, *drugs*]. The converted extraction tuple is (*physicians*, *provide*, *drugs*, *in Asian countries*).

To ensure data quality and as a result of differences between the original QA-SRL dataset and the QA-SRL BANK 2.0, we had to make two important changes to the algorithm:

Answer Filtering: The original QA-SRL dataset has a single set of mutually-exclusive answer spans for each question, written by a single annotator. In contrast, the QA-SRL BANK 2.0 has answer judgments from three annotators for each question, some providing answer sets and others marking the questions as invalid. To consolidate these, we only include questions marked as valid by all three annotators. Then, for each question, we iteratively draw the longest remaining answer

<i>Bats are the only mammals that can truly fly.</i> (Bats, fly)
<i>Greece moved up three to be ranked tenth.</i> (Greece, ranked , tenth)
<i>A popular student, in 1915 Mao was elected secretary of the Students Society.</i> (Mao, elected , secretary of the Students Society, in 1915)
<i>The proposed amendment already passed both houses in 2011.</i> (The proposed amendment, passed , both houses, in 2011)
<i>In polygynous species, males try to monopolize and mate with multiple females.</i> (males, monopolize , multiple females)
<i>Animals adapted to live in the desert are called xerocoles.</i> (Animals, adapted , to live in the desert)

Table 2: Example sentences with example extractions. Note that only one example extraction is shown here, though a sentence can yield multiple extractions.

span that does not overlap with a previously drawn answer span, until there are none left.

In answer filtering, our primary motivation was to clean the raw version of crowd workers’ answer responses in the QA-SRL 2.0 dataset, where questions can be posed that are not valid or the answer to them is ambiguous. We found it advantageous for dataset quality to require a strict agreement between all annotators. In choosing the longest answer span, we were motivated to not miss relevant portions of the argument, as individual crowd workers occasionally annotated a limited portion of the answer span that did not encapsulate the whole semantic meaning of the derived argument.

Argument Ordering: Stanovsky and Dagan (2016)’s original algorithm relies on the original, annotator-written order of QA-SRL questions, which may or may not produce a sensible argument ordering. Furthermore, in the QA-SRL BANK 2.0, the original order in which the questions were written is unavailable.

So, to determine argument order, we use a heuristic based on the relative order between answer spans for each question in their source text. We consider the abstract form of questions, which includes verb tense without information about its lemma. For a given question q_i in an extraction, let $q_{i,x}$ represent the percentage of predicates in the QA-SRL BANK 2.0 where the answer span to the generalized version of q_i appears in the x^{th} place relative to other answer spans, according to the natural order of the sentence. For each argument slot in the derived extraction, the answer to the question with the highest probability $q_{i,x}$ of naturally occurring in that slot is chosen as the argument.




Figure 3: Top: performance of Supervised OIE systems on the LSOIE-wiki test set. Bottom: `ls_oie` estimated confidence at each extraction.

In our example extraction in Figure 1, the question *Who* [predicate] *something*? precedes *What is being* [predicate]? which precedes *Where does someone* [predicate] *something*?, enabling our algorithm to accurately extract argument ordering, which is not available from the natural ordering of the sentence or the ordering of crowd annotations in FitzGerald et al. (2018).

3.2 Dataset Statistics

We run our updated dataset conversion process over the directly crowdsourced portion of the train, development, and test partitions of the QA-SRL BANK 2.0. Stratifying the resulting data by domain, we present the new LSOIE corpus in two sections, LSOIE-wiki and LSOIE-sci. Dataset statistics are shown in Table 1. Example extractions are shown in Table 2. We provide the distribution of argument, predicate, and null tag labels in Figure 2. The LSOIE corpus expands the scope of OIE2016 and AW-OIE in size, textual diversity, and domain.

4 Benchmark Evaluation

Models: We evaluate several models on our new LSOIE dataset. Following Stanovsky et al. (2018), we model OIE as a supervised learning problem and format it as BIO tagging with tunable thresholding on extractions. We benchmark several model variants:

- **rnnoie** is a replication of the model in Stanovsky et al. (2018), based on a bidirectional LSTM transducer over GloVe embed-

	LSOIE-wiki		LSOIE-sci	
Model	F_1	AUC	F_1	AUC
rnnoie	.22	.07	.26	.10
ls_oie	.28	.13	.33	.18
ls_oie_crf	.29	.14	.33	.19
srl_bert_oie2016	.23	.08	.29	.13
srl_bert_ls	.31	.16	.37	.21
ls_oie_sci	-	-	.34	.19
ls_oie_crf_sci	-	-	.35	.20
srl_bert_ls_sci	-	-	.38	.22

Table 3: Modeling results on the LSOIE test sets.

dings (Pennington et al., 2014) and learned part-of-speech embedding features.

- **ls_oie** is a replication of `rnnoie` trained on LSOIE.
- **ls_oie_crf** is the same as `ls_oie`, but trained end-to-end with a Conditional Random Field on top to capture BIO transition constraints and trained to maximize the likelihood of the gold BIO sequence.
- **srl_bert_ls** is based on `ls_oie`, but uses BERT (Devlin et al., 2019) as the bidirectional encoder and the Sentence A / Sentence B embedding feature as the predicate indicator, inspired by Shi and Lin (2019).
- **srl_bert_oie2016** is the same architecture as `srl_bert_ls` but applied to the OIE2016 data.
- ***_sci** models were trained with the same architectures applied only to the LSOIE-sci training set.

Experiments and Evaluation: We use the AllenNLP framework (Gardner et al., 2018) built on PyTorch (Paszke et al., 2019) to implement, train, and test our models. We train `rnnoie` and `srl_bert_oie2016` on OIE2016 and `ls_oie` and `srl_bert_ls` on LSOIE-wiki. We also focus the series of models by only training on LSOIE-sci. We do not evaluate *_sci models on LSOIE-wiki. We limit our evaluation to supervised OIE systems.

We evaluate our system’s performance against the gold test data in LSOIE-wiki and LSOIE-sci by considering extractions to be a match if they contain the same predicate as the gold extraction and contain the syntactic head of each gold argument. Syntactic heads are extracted with the Stanford CoreNLP dependency parser (Chen and Manning, 2014). Although it would be ideal to have the gold syntactic head, this method is preferable to taking the lexical overlap of the entire extraction

Stanovsky and Dagan (2016), ignoring argument tags and ordering as pointed out in Lechelle et al. (2019).

We then assign a confidence score to each extraction to allow for tuning the precision-recall tradeoff. For the non-CRF models, we use the mean log probability assigned to the tag labels in the extraction as the confidence score. For the CRF model, we use the log probability assigned to the entire sequence. We differ from Stanovsky et al. (2018) where confidence was calculated as the product of the inverse of the model’s estimate probability for each tag label, preferring longer extractions which were more likely to get a 50% lexical match, outweighing the deficit of swimming upstream against the model’s estimated confidence and still producing a downward sloping precision recall curve.

We use Viterbi decoding to extract the most likely valid BIO tagging sequence given the model’s probability output for each BIO tag. We import the Viterbi algorithm functionality from the AllenNLP library (Gardner et al., 2018).

5 Discussion

Figure 3 shows precision and recall curves on the LSOIE-wiki test set, accompanied by the `ls_oie` model’s estimated confidence. Table 3 shows F_1 and AUC scores for the benchmark models on the LSOIE-wiki and LSOIE-sci test sets.

The OIE modeling task is difficult. Results on both evaluation sets show that the BERT model and the CRF output layer improve over the baseline model. Training with the LSOIE improves model performance. When science is the target domain, the `*_sci` models are preferable, as they have slightly higher in-domain performance, showing the value of the domain split in LSOIE.

5.1 Error Analysis

We conduct a manual error analysis of the `ls_oie` model, where we find that our baseline models could benefit from more careful extractions.

Incorrect predicate: At minimum confidence, 53% of the model’s precision errors come from verbs that are not present in the gold dataset. Half of these are legitimate predicates that are missing from the gold dataset and the other half are auxiliary verbs, that should not be present in the gold dataset. Depending on the deployment environment, the model could be improved with predicate filtering heuristics at prediction time.

Argument Concatenation: We examined 500 incorrect extractions by `ls_oie`. We found that 36% of unmatched extractions were semantically similar to the gold extraction. These extractions either concatenated arguments A_1-A_N into A_1 while gold did not, split these arguments apart while gold did, or dropped a non-material argument. For future modeling, this is an argument to drop A_2 and beyond from the dataset and only model OIE with extraction triples.

True Errors: Among the extraction errors, 2/3 involve errors in argument ordering, often following the natural order of the sentence. The other 1/3 of errors involved the model making nonsensical extractions or not making extracting arguments beyond A_0 , presumably because of lack of confidence and defaulting to the O label.

LSOIE Modeling Improvements: We also manually examined 100 extractions where `ls_oie` chose the right extraction over `rnnoie`. In these cases, we found improved argument ordering, increased confidence on relevant A_1 objects, and better accuracy identifying subjects that are distant from the predicate.

6 Conclusion

In this paper, we introduced the LSOIE dataset as a resource for supervised OIE. We have algorithmically re-purposed the QA-SRL BANK 2.0 into a new OIE dataset, LSOIE, which contains over 70,000 sentences and over 150,000 extraction tuples. To benchmark the new dataset, we trained and evaluated a series of supervised OIE models, providing baselines for future research on the OIE modeling task.

The code and datasets introduced in this paper can be found at <https://github.com/Jacobsolawetz/large-scale-oie>.

Acknowledgments

A special thanks to Julian Michael and Gabriel Stanovsky for providing much needed guidance and feedback for our research. This project would not have been possible without their insight and previous progress on the topic. We also thank the anonymous reviewers for their feedback.

References

Michele Banko, Michael J. Cafarella, Stephen Soderland, Matt Broadhead, and Oren Etzioni. 2007.

- Open information extraction from the web. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Sangnie Bhardwaj, Samarth Aggarwal, and Mausam Mausam. 2019. **CaRB: A crowdsourced benchmark for open IE**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Danqi Chen and Christopher Manning. 2014. **A fast and accurate dependency parser using neural networks**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Lei Cui, Furu Wei, and Ming Zhou. 2018. **Neural open information extraction**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **Bert: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*.
- Nicholas FitzGerald, Julian Michael, Luheng He, and Luke Zettlemoyer. 2018. **Large-scale QA-SRL parsing**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. **AllenNLP: A deep semantic natural language processing platform**. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*.
- Kiril Gashteovski, Rainer Gemulla, and Luciano del Corro. 2017. **MinIE: Minimizing facts in open information extraction**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Kiril Gashteovski, Sebastian Wanner, Sven Hertling, Samuel Broscheit, and Rainer Gemulla. 2019. **Opiec: An open information extraction corpus**. In *Proceedings of the Conference on Automatic Knowledge Base Construction (AKBC)*.
- Luheng He, Mike Lewis, and Luke Zettlemoyer. 2015. **Question-answer driven semantic role labeling: Using natural language to annotate natural language**. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Shengbin Jia, Yang Xiang, and Xiaojun Chen. 2018. **Supervised neural models revitalize the open relation extraction**. *CoRR*, abs/1809.09408.
- Zhengbao Jiang, Pengcheng Yin, and Graham Neubig. 2019. **Improving open information extraction via iterative rank-aware learning**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- William Lechelle, Fabrizio Gotti, and Phillippe Langlais. 2019. **WiRe57 : A fine-grained benchmark for open information extraction**. In *Proceedings of the 13th Linguistic Annotation Workshop (LAW)*.
- Omer Levy, Ido Dagan, and Jacob Goldberger. 2014. **Focused entailment graphs for open IE propositions**. In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning (CoNLL)*.
- Mausam. 2016. **Open information extraction systems and downstream applications**. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI)*.
- Martha Palmer, Daniel Gildea, and Paul Kingsbury. 2005. **The proposition bank: An annotated corpus of semantic roles**. *Comput. Linguist.*, 31(1):71–106.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. **Pytorch: An imperative style, high-performance deep learning library**. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. **GloVe: Global vectors for word representation**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Peng Shi and Jimmy Lin. 2019. **Simple BERT models for relation extraction and semantic role labeling**. *CoRR*, abs/1904.05255.
- Gabriel Stanovsky and Ido Dagan. 2016. **Creating a large benchmark for open information extraction**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Gabriel Stanovsky, Julian Michael, Luke Zettlemoyer, and Ido Dagan. 2018. **Supervised open information extraction**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*.
- Dongxu Zhang, Subhabrata Mukherjee, Colin Lockard, Luna Dong, and Andrew McCallum. 2019. **OpenKI: Integrating Open Information Extraction and Knowledge Bases with Relation Inference**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*.