

Open Information Extraction on Scientific Text: An Evaluation

Paul Groth Elsevier Labs 1600 John F. Kennedy Blvd. Suite 1800 Philadelphia, PA p.groth@elsevier.com	Mike Lauruhn Elsevier Labs 1600 John F. Kennedy Blvd. Suite 1800 Philadelphia, PA m.lauruhn@elsevier.com	Antony Scerri Elsevier Labs 1600 John F. Kennedy Blvd. Suite 1800 Philadelphia, PA a.scerri@elsevier.com	Ron Daniel, Jr. Elsevier Labs 1600 John F. Kennedy Blvd. Suite 1800 Philadelphia, PA r.daniel@elsevier.com
---	---	---	---

Abstract

Open Information Extraction (OIE) is the task of the unsupervised creation of structured information from text. OIE is often used as a starting point for a number of downstream tasks including knowledge base construction, relation extraction, and question answering. While OIE methods are targeted at being domain independent, they have been evaluated primarily on newspaper, encyclopedic or general web text. In this article, we evaluate the performance of OIE on scientific texts originating from 10 different disciplines. To do so, we use two state-of-the-art OIE systems using a crowd-sourcing approach. We find that OIE systems perform significantly worse on scientific text than encyclopedic text. We also provide an error analysis and suggest areas of work to reduce errors. Our corpus of sentences and judgments are made available.

1 Introduction

The scientific literature is growing at a rapid rate (Bornmann and Mutz, 2015). To make sense of this flood of literature, for example, to extract cancer pathways (Poon et al., 2014) or find geological features (Leveling, 2015), increasingly requires the application of natural language processing. Given the diversity of information and its constant flux, the use of unsupervised or distantly supervised techniques are of interest (Quirk and Poon, 2017). In this paper, we investigate one such unsupervised method, namely, Open Information Extraction (OIE) (Banko et al., 2007). OIE is the task of the unsupervised creation of structured information from text. OIE is often used as a starting point for a number of downstream tasks including knowledge base construction, relation extraction, and question answering (Mausam, 2016).

While OIE has been applied to the scientific literature before (Groth et al., 2016), we have not found a systematic evaluation of OIE as applied to scientific publications. The most recent evaluations of OIE extraction tools (Gashteovski et al., 2017; Schneider et al., 2017) have instead looked at the performance of these tools on traditional NLP information sources (i.e. encyclopedic and news-wire text). Indeed, as (Schneider et al., 2017) noted, there is little work on the evaluation of OIE systems. Thus, the goal of this paper is to evaluate the performance of the state of the art in OIE systems on scientific text.

Specifically, we aim to test two hypotheses:

1. H1: There is no significant difference in the accuracy of OIE systems on scientific vs. general audience content.
2. H2: There is no significant difference in performance of current state-of-the-art OIE systems on scientific and medical content.

Additionally, we seek to gain insight into the value of unsupervised approaches to information extraction and also provide information useful to implementors of these systems. We note that our evaluation differs from existing OIE evaluations in that we use crowd-sourcing annotations instead of expert annotators. This allows for a larger number of annotators to be used. All of our data, annotations and analyses are made openly available.¹

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>

¹<http://dx.doi.org/10.17632/6m5dyx4b58.2>

The rest of the paper is organized as follows. We begin with a discussion of existing evaluation approaches and then describe the OIE systems that we evaluated. We then proceed to describe the datasets used in the evaluation and the annotation process that was employed. This is followed by the results of the evaluation including an error analysis. Finally, we conclude.

2 Existing Evaluation Approaches

OIE systems analyze sentences and emit relations between one predicate and two or more arguments (e.g. Washington :: was :: president). The arguments and predicates are not fixed to a given domain. (Note, that throughout this paper we use the word ‘triple’ to refer interchangeably to binary relations.) Existing evaluation approaches for OIE systems have primarily taken a ground truth-based approach. Human annotators analyze sentences and determine correct relations to be extracted. Systems are then evaluated with respect to the overlap or similarity of their extractions to the ground truth annotations, allowing the standard metrics of precision and recall to be reported.

This seems sensible but is actually problematic because of different but equivalent representations of the information in an article. For example, consider the sentence “The patient was treated with Emtricitabine, Etravirine, and Darunavir”. One possible extraction is:

(The patient :: was treated with :: Emtricitabine, Etravirine, and Darunavir)

Another possible extraction is:

(The patient :: was treated with :: Emtricitabine)

(The patient :: was treated with :: Etravirine)

(The patient :: was treated with :: Darunavir)

Neither of these is wrong, but by choosing one approach or the other a pre-constructed gold set will falsely penalize a system that uses the other approach.

From such evaluations and their own cross dataset evaluation, (Schneider et al., 2017) list the following common errors committed by OIE systems:

- wrong boundaries around the arguments or predicate of a relation;
- generation of redundant relations from the same sentence;
- wrong extractions (e.g. omitting large parts of a sentence as an argument);
- missing extractions - extractions that should have been extracted from a sentence;
- uninformative extractions that omit critical information from a sentence;

In our evaluation, we take a different approach. We do not define ground truth relation extractions from the sentences in advance. Instead, we manually judge the correctness of each extraction after the fact. We feel that this is the crux of the information extraction challenge. Is what is being extracted correct or not? This approach enables us to consider many more relations through the use of a crowd-sourced annotation process. Our evaluation approach is similar to the qualitative analysis performed in (Schneider et al., 2017) and the evaluation performed in (Gashtevski et al., 2017). However, our evaluation is able to use more judges (5 instead of 2) because we apply crowd sourcing.² For our labelling instructions, we adapted those used by (Gashtevski et al., 2017) to the crowd sourcing setting.³

As previously noted existing evaluations have also only looked at encyclopedic or newspaper corpora. Several systems (e.g. (Banko et al., 2007; Del Corro and Gemulla, 2013)) have looked at text from the web as well, however, as far as we know, none have specifically looked at evaluation for scientific and medical text.

²(Schneider et al., 2017) also perform a ground truth based evaluation. This evaluation had many more sentences than ours, but used predefined gold extractions from multiple corpora annotated according to different criteria; illustrating the kinds of issues we mention with such an evaluation method.

³Labelling instructions are included in the associated dataset. - <http://dx.doi.org/10.17632/6m5dyx4b58.2#file-7edb4b86-c0e6-4169-aea0-4862d39461d3>

3 Systems

We evaluate two OIE systems (i.e. extractors). The first, OpenIE 4 (Mausam, 2016), descends from two popular OIE systems OLLIE (Fader et al., 2011) and Reverb (Fader et al., 2011). We view this as a baseline system. The second was MiniIE (Gashteovski et al., 2017), which is reported as performing better than OLLIE, ClauseIE (Del Corro and Gemulla, 2013) and Stanford OIE (Del Corro and Gemulla, 2013). MiniIE focuses on the notion of minimization - producing compact extractions from sentences. In our experience using OIE on scientific text, we have found that these systems often produce overly specific extractions that do not provide the redundancy useful for downstream tasks. Hence, we thought this was a useful package to explore.

We note that both OpenIE 4 and MiniIE support relation extractions that go beyond binary tuples, supporting the extraction of n-ary relations. We note that the most recent version of Open IE (version 5) is focused on n-ary relations. For ease of judgement, we focused on binary relations. Additionally, both systems support the detection of negative relations.

In terms of settings, we used the off the shelf settings for OpenIE 4. For MiniIE, we used their “safe mode” option, which uses slightly more aggressive minimization than the standard setting. In the recent evaluation of MiniIE, this setting performed roughly on par with the default options (Gashteovski et al., 2017). Driver code showing how we ran each system is available.⁴

4 Datasets

We used two different data sources in our evaluation. The first dataset (WIKI) was the same set of 200 sentences from Wikipedia used in (Gashteovski et al., 2017). These sentences were randomly selected by the creators of the dataset. This choice allows for a rough comparison between our results and theirs.

The second dataset (SCI) was a set of 220 sentences from the scientific literature. We sourced the sentences from the OA-STM corpus.⁵ This corpus is derived from the 10 most published in disciplines. It includes 11 articles each from the following domains: agriculture, astronomy, biology, chemistry, computer science, earth science, engineering, materials science, math, and medicine. The article text is made freely available and the corpus provides both an XML and a simple text version of each article.

We randomly selected 2 sentences with more than two words from each paper using the simple text version of the paper. We maintained the id of the source article and the line number for each sentence.

5 Annotation Process

We employed the following annotation process. Each OIE extractor was applied to both datasets with the settings described above. This resulted in the generation of triples for 199 of the 200 WIKI sentences and 206 of the 220 SCI sentences. That is there were some sentences in which no triples were extracted. We discuss later the sentences in which no triples were extracted. In total 2247 triples were extracted.

The sentences and their corresponding triples were then divided. Each task contained 10 sentences and all of their unique corresponding triples from a particular OIE systems. Half of the ten sentences were randomly selected from SCI and the other half were randomly selected from WIKI. Crowd workers were asked to mark whether a triple was correct, namely, did the triple reflect the consequence of the sentence. Examples of correct and incorrect triples were provided. Complete labelling instructions and the presentation of the HITS can be found with the dataset. All triples were labelled by at least 5 workers.

Note, to ensure the every HIT had 10 sentences, some sentences were duplicated. Furthermore, we did not mandate that all workers complete all HITS.

We followed recommended practices for the use of crowd sourcing in linguistics (Erlewine and Kotek, 2016). We used Amazon Mechanical Turk as a means to present the sentences and their corresponding triples to a crowd for annotation. Within Mechanical Turk tasks are called Human Intelligence Tasks (HITs). To begin, we collected a small set of sentences and triples with known correct answers. We did

⁴See directory “Code for applying information extraction tools” in the associated data - <http://dx.doi.org/10.17632/6m5dyx4b58.2#folder-39ce9705-ccf3-4f95-890a-508ce155ece4>

⁵<http://elsevierlabs.github.io/OA-STM-Corpus/>

this by creating a series of internal HITs and loaded them the Mechanical Turk development environment called the Mechanical Turk Sandbox. The HITs were visible to a trusted group of colleagues who were asked to complete the HITs.

Having an internal team of workers attempt HITs provides us with two valuable aspects of the eventual production HITs. First, internal users are able to provide feedback related to usability and clarity of the task. They were asked to read the instructions and let us know if there was anything that was unclear. After taking the HITs, they are able to ask questions about anomalies or confusing situations they encounter and allow us to determine if specific types of HITs are either not appropriate for the task or might need further explanation in the instructions. In addition to the internal users direct feedback, we were also able to use the Mechanical Turk Requester functionality to monitor how long (in minutes and seconds) it took each worker to complete each HIT. This would come into factor how we decided on how much to pay each Worker per HIT after they were made available to the public.

The second significant outcome from the internal annotations was the generation of a set of ‘expected’ correct triples. Having a this set of annotations is an integral part of two aspects of our crowdsourcing process. First, it allows us to create a qualification HIT. A qualification HIT is a HIT that is made available to the public with the understanding the Workers will be evaluated based on how closely they matched the annotations of the internal annotators. Based upon this, the Workers with the most matches would be invited to work on additional tasks. Second, we are able to add the internal set of triples randomly amongst the other relations we were seeking to have annotated. This allows us to monitor quality of the individual Workers over the course of the project. Note, none of this data was used in the actual evaluation. It was only for the purposes of qualifying Workers.

We are sensitive to issues that other researchers have in regards to Mechanical Turk Workers earning fair payment in exchange for their contributions to the HITs (Fort et al., 2011) . We used the time estimates from our internal annotation to price the task in order to be above US minimum wage. All workers were qualified before being issued tasks. Overall, we employed 10 crowd workers. On average it took 30 minutes for a worker to complete a HIT. In line with (Crump et al., 2013), we monitored for potential non-performance or spam by looking for long response times and consecutive submitted results. We saw no indicators of low quality responses.

6 Judgement Data and Inter-Annotator Agreement

In total, 11262 judgements were obtained after running the annotation process. Every triple had at least 5 judgements from different annotators. All judgement data is made available.⁶ The proportion of overall agreement between annotators is 0.76 with a standard deviation of 0.25 on whether a triple is consequence of the given sentence. We also calculated inter-annotator agreement statistics. Using Krippendorff’s alpha inter-annotator agreement was 0.44. This calculation was performed over all data and annotators as Krippendorff’s alpha is designed to account for missing data and work across more than two annotators. Additionally, Fleiss’ Kappa and Scott’s pi were calculated pairwise between all annotators where there were overlapping ratings (i.e. raters had rated at least one triple in common). The average Fleiss’s Kappa was 0.41 and the average of Scott’s pi was 0.37. Using (Artstein and Poesio, 2008) as a guide, we interpret these statistics as suggesting there is moderate agreement between annotators and that agreement is above random chance. This moderate level of agreement is to be expected as the task itself can be difficult and requires judgement from the annotators at the margin.

Table 1 shows examples of triples that were associated with higher disagreement between annotators. One can see for example, in the third example, that annotators might be confused by the use of a pronoun (him). Another example is in the last sentence of the table, where one can see that there might be disagreement on whether the subsequent prepositional phrase behind light microscope analysis should be included as part of the extracted triple.

We take the variability of judgements into account when using this data to compute the performance of the two extraction tools. Hence, to make assessments as to whether a triple correctly reflects the content

⁶‘aggregated_results_anon.csv’ in the associated dataset. - <http://dx.doi.org/10.17632/6m5dyx4b58.2#file-03de3c93-8a01-4cd2-bfe5-c5bfeaa4f492>

SYSTEM	SOURCE	TRIPLE	PAIRWISE AGREEMENT	SENTENCE
MinIE	SCI	additional QUANT_S_1 on scalp face :: were digitized :: QUANT_O_1 anatomical landmarks	40% False	To coregister MEG and sMRI data, three anatomical landmarks (nasion and right and left preauriculars) as well as an additional 150+ points on the scalp and face were digitized for each subject using the Probe Position Identification (PPI) System (Polhemus, Colchester, VT).
OpenIE 4	WIKI	Dawlish :: to 2nd :: place L:in the Western League	40% True	The previous season had seen him lead Dawlish to 2nd place in the Western League , their highest ever league finishing position .
MiniIE	WIKI	him :: lead Dawlish to 2nd place in :: Western League'	40% False	
MinIE	SCI	> 250 μ m fractions :: be found through :: light microscope analysis to be dominated by amorphous organic matter	40% True	'The > 250 μ m fractions, found through light microscope analysis to be dominated (> 90%) by amorphous organic matter (AOM), were also analysed for $\delta^{13}C$.

Table 1: Examples of difficult to judge triples and their associated sentences.

from which it is extracted, we rely on the unanimous positive agreement between crowd workers. That is to say that if we have 100% inter-annotator agreement that a triple was correctly extracted we label it as correct.

7 Experimental Results

Table 2 show the results for the combinations of systems and data sources. The CORRECT TRIPLES column contains the number of triples that are labelled as being correct by all annotators. TOTAL TRIPLES are the total number of triples extracted by the given systems over the specified data. Precision is calculated as typical where Correct Triples are treated as true positives. On average, 3.1 triples were extracted per sentence.

Figure 1 shows the performance of extractors in terms of precision as inter-annotator agreement de-

SYSTEM	SOURCE	TOTAL TRIPLES	CORRECT TRIPLES	PRECISION
OpenIE 4 + MinIE	SCI + WIKI	2247	985	0.44
OpenIE 4 + MinIE	WIKI	1101	590	0.54
OpenIE 4 + MinIE	SCI	1146	395	0.34
MinIE	SCI + WIKI	1591	617	0.39
MinIE	WIKI	785	382	0.49
MinIE	SCI	806	235	0.29
OpenIE 4	SCI + WIKI	656	368	0.56
OpenIE 4	WIKI	316	208	0.66
OpenIE 4	SCI	340	160	0.47

Table 2: Results of triples extracted from the SCI and WIKI corpora using the Open IE and MinIE tools.

creases. In this figure, we look only at agreement on triples where the majority agree that the triple is correct. Furthermore, to ease comparison, we only consider triples with 5 judgements this excludes 9 triples. We indicate not only the pair-wise inter-annotator agreement but also the number of annotators who have judged a triple to be correct. For example, at the 40% agreement level at least 3 annotators have agreed that a triple is true. The figure separates the results by extractor and by data source.




Figure 1: Precision at various agreement levels. Agreement levels are shown as the proportion of overall agreement. In addition, we indicate the the minimum number of annotators who considered relations correct out of the total number of annotators.

We see that as expected the amount of triples agreed to as correct grows larger as we relax the requirement for agreement. For example, analyzing Open IE’s results, at the 100% agreement level we see a precision of 0.56 whereas at the 40% agreement level we see a precision of 0.78. Table 3 shows the total number of correct extractions at the three agreement levels.

AGREEMENT	OPEN IE 4 SCI	MINIE SCI	OPEN IE 4 WIKI	MINIE WIKI	OPEN IE 4 TOTAL	MINIE TOTAL
40% (3+/5)	253	462	256	553	509	1015
60% (4+/5)	218	361	241	488	459	849
100% (5/5)	160	235	207	377	367	612

Table 3: Correct triples at different levels of agreement subsetted by system and data source . Agreement levels follow from Figure 1

7.1 Testing H1: Comparing the Performance of OIE on Scientific vs. Encyclopedic Text

From the data, we see that extractors perform better on sentences from Wikipedia (0.54 P) than scientific text (0.34 P). Additionally, we see that there is higher annotator agreement on whether triples extracted from Wikipedia and scientific text are correct or incorrect: 0.80 - SD 0.24 (WIKI) vs. 0.72 - SD 0.25 (SCI). A similar difference in agreement is observed when only looking at triples that are considered to be correct by the majority of annotators: 0.87 - SD 0.21 (WIKI) vs. 0.78 - SD 0.25 (SCI) . In both cases, the difference is significant with p-values < 0.01 using Welch’s t-test. The differences between

data sources are also seen when looking at the individual extraction tools. For instance, for Open IE 4 the precision is 0.19 higher for wikipedia extractions over those from scientific text. *With this evidence, we reject our first hypothesis that the performance of these extractors are similar across data sources.*

7.2 Testing H2: Comparing the Performance of Systems

We also compare the output of the two extractors. In terms precision, Open IE 4 performs much better across the two datasets (0.56P vs 0.39P). Looking at triples considered to be correct by the majority of annotators, we see that Open IE 4 has higher inter-annotator agreement 0.87 - SD 0.22 (Open IE) vs 0.81 - SD 0.24 (MinIE). Focusing on scientific and medical text (SCI), again where the triples are majority annotated as being correct, Open IE has higher inter-annotator agreement (Open IE: 0.83 - SD 0.24 vs MiniIE: 0.76 - SD 0.25). In both cases, the difference is significant with p-values < 0.01 using Welch's t-test. This leads us to conclude that Open IE produces triples that annotators are more likely to agree as being correct.

MinIE provides many more correct extractions than OpenIE 4 (935 more across both datasets). The true recall numbers of the two systems can not be calculated with the data available, but the 40% difference in the numbers of correct extractions is strong evidence that the two systems do not have equivalent behavior.

A third indication of differences in their outputs comes from examining the complexity of the extracted relations. Open IE 4 generates longer triples on average (11.5 words) vs. 8.5 words for MinIE across all argument positions. However, Open IE 4 generates shorter relation types than MinIE (Open IE - 3.7 words; MiniIE 6.27 words) and the standard deviation in terms of word length is much more compact for Open IE 4 - 1 word vs 3 words for MinIE. *Overall, our conclusion is that Open IE 4 performs better than MinIE both in terms of precision and compactness of relation types, while not matching MinIE's recall, and thus we reject our second hypothesis.*

7.3 Other Observations

The amount of triples extracted from the scientific text is slightly larger than that extracted from the Wikipedia text. This follows from the fact that the scientific sentences are on average roughly 7 words longer than encyclopedic text.

The results of our experiment also confirm the notion that an unsupervised approach to extracting relations is important. We have identified 698 unique relation types that are part of triples agreed to be correct by all annotators. This number of relation types is derived from only 400 sentences. While not every relation type is essential for downstream tasks, it is clear that building specific extractors for each relation type in a supervised setting would be difficult.

8 Error Analysis

We now look more closely at the various errors that were generated by the two extractors.

Table 4 shows the sentences in which neither extractor produced triples. We see 3 distinct groups. The first are phrases that are incomplete sentences usually originating from headings (e.g. Materials and methods). The next group are descriptive headings potentially coming from paper titles or figure captions. We also see a group with more complex prepositional phrases. In general, these errors could be avoided by being more selective of the sentences used for random selection. Additionally, these systems could look at potentially just extracting noun phrases with variable relation types, hence, expressing a cooccurrence relation.

We also looked at where there was complete agreement by all annotators that a triple extraction was incorrect. In total there were 138 of these triples originating from 76 unique sentences. There were several patterns that appeared in these sentences.

- Long complex sentences led to incorrect extractions. For example, ``Most strikingly, there was a mutant gene-dose-dependent increase in caspase-cleaved TAU fragments, as determined by the caspase-cleaved TAU-specific antibody C3 (Gamblin et al., 2003; Guillozet-Bongaarts et al., 2005), in neurons derived from the isogenic

SOURCE	SENTENCE
SCI	<p>Note that Eq. Materials and methods Site and experimental carbonate chemistry Soil aggregate characteristics An equation analogous to Eq.</p> <hr/> <p>An Experimental Platform for the Efficient Generation of Human Cranial Placodes In Vitro Analysis 1: Manifest Vascular and Nonvascular Disease as a Predictor of Depressive Symptoms Autografts Elicit Only a Minimal Immune Response in the Primate Brain Production of wild-type TTR and ATTR Val30Met by differentiated hepatocyte-like cells Cryopreservation and recovery of hiPSCs in suspension culture</p> <hr/> <p>For a detailed description of analysis methods and precision see Lee et al. (1997). Thus there are, up to associativity, only finitely many such q. In the absence of heavy elements, $H3+$ forms near the base of the model and subsequent infrared cooling balances the EUV heating rates.</p> <hr/> <p>Assume that L is a line bundle. The set of the representative points is called the non-dominated front (or Pareto front).</p>
WIKI	Simultaneously won the “ Hope of the World Ballet ” Prize .

Table 4: Sentences in which no triples were extracted

TAU-A152T-iPSCs (Figures 4J and 4K).''') led to the following triples that were deemed incorrect: isogenic TAU-A152T-iPSCs :: is :: Figures and Gamblin et :: is :: caspase-cleaved TAU-specific antibody C3

- The use of pronouns as subjects led to triples that were deemed to be incorrect. (e.g. ``So he was forced to requisition not only the public treasury of Gades but also the wealth from its temples. led to the following incorrect triples: he :: was forced :: to requisition not only the public treasury of Gades but also the wealth from its temples and he :: to requisition :: not only the public treasury of Gades but also the wealth from its temples.
- Complex mathematical formula often generated incorrect triples
- MinIE as part of its extraction process substitutes quantities with variables (e.g. QUANT_1). We included this potential in our labelling instructions but this often led to triples that were labelled as incorrect by the annotators. We believe this source of error could stem from the given instructions.
- Reuse of abbreviations as adjectives also led to incorrect triples. For example, “BMP antagonists” and “BMP pathway” in the following sentence: We also observed significant transcriptional changes in WNT and BMP pathway components such as an increase in the WNT pathway inhibitor DKK-1 and BMP antagonists, such as GREMLIN-1 and BAMBI (Figures 2B-2D), which are known transcriptional targets of BMP signaling (Grotewold et al., 2001). led to the following triples that were labelled incorrect: BMP antagonists :: are known :: and increase :: is :: bmp pathway component.

We also see similar errors to those pointed out by (Schneider et al., 2017), namely, uninformative extractions, the difficulty in handling n-ary relations that are latent in the text, difficulties handling negations, and very large argument lengths. In general, these errors together point to several areas for further improvement including:

- deeper co-reference resolution either for variables in mathematical formula or for pronouns;
- improved handling of prepositional phrases;
- relaxing requirements for correct grammar within sentences;
- better handling of abbreviations.

9 Conclusion

The pace of change in the scientific literature means that interconnections and facts in the form of relations between entities are constantly being created. Open information extraction provides an important tool to keep up with that pace of change. We have provided evidence that unsupervised techniques are needed to be able to deal with the variety of relations present in text. The work presented here provides an independent evaluation of these tools in their use on scientific text. Past evaluations have focused on encyclopedic or news corpora which often have simpler structures. We have shown that existing OIE systems perform worse on scientific and medical content than on general audience content.

There are a range of avenues for future work. First, the application of Crowd Truth framework (Aroyo and Welty, 2013) in the analysis of these results might prove to be useful as we believe that the use of unanimous agreement tends to negatively impact the perceived performance of the OIE tools. Second, we think the application to n-ary relations and a deeper analysis of negative relations would be of interest. To do this kind of evaluation, an important area of future work is the development of guidelines and tasks for more complex analysis of sentences in a crowd sourcing environment. The ability, for example, to indicate argument boundaries or correct sentences can be expected of expert annotators but needs to be implemented in a manner that is efficient and easy for the general crowd worker. Third, we would like to expand the evaluation dataset to an even larger numbers of sentences. Lastly, there are a number of core natural language processing components that might be useful for OIE in this setting, for example, the use of syntactic features as suggested by (Christensen et al., 2011). Furthermore, we think that coreference is a crucial missing component and we are actively investigating improved coreference resolution for scientific texts.

To conclude, we hope that this evaluation provides further insights for implementors of these extraction tools to deal with the complexity of scientific and medical text.

References

- Lora Aroyo and Chris Welty. 2013. Crowd truth: Harnessing disagreement in crowdsourcing a relation extraction gold standard. *WebSci2013. ACM*, 2013.
- Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.
- 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'07*, pages 2670–2676, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Lutz Bornmann and Rüdiger Mutz. 2015. Growth rates of modern science: A bibliometric analysis based on the number of publications and cited references. *Journal of the Association for Information Science and Technology*, 66(11):2215–2222.
- Janara Christensen, Mausam, Stephen Soderland, and Oren Etzioni. 2011. An analysis of open information extraction based on semantic role labeling. In *Proceedings of the Sixth International Conference on Knowledge Capture, K-CAP '11*, pages 113–120, New York, NY, USA. ACM.

- Matthew J. C. Crump, John V. McDonnell, and Todd M. Gureckis. 2013. Evaluating amazon’s mechanical turk as a tool for experimental behavioral research. *PLOS ONE*, 8(3):1–18, 03.
- Luciano Del Corro and Rainer Gemulla. 2013. Clauseie: Clause-based open information extraction. In *Proceedings of the 22Nd International Conference on World Wide Web, WWW ’13*, pages 355–366, New York, NY, USA. ACM.
- Michael Yoshitaka Erlewine and Hadas Kotek. 2016. A streamlined approach to online linguistic surveys. *Natural Language & Linguistic Theory*, 34(2):481–495, may.
- Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011. Identifying relations for open information extraction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP ’11*, pages 1535–1545, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Karën Fort, Gilles Adda, and K. Bretonnel Cohen. 2011. Amazon Mechanical Turk: Gold Mine or Coal Mine? *Computational Linguistics*, 37(2):413–420, jun.
- Kiril Gashteovski, Rainer Gemulla, and Luciano Del Corro. 2017. MinIE: Minimizing Facts in Open Information Extraction. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2620–2630.
- Paul T. Groth, Sujit Pal, Darin McBeath, Brad Allen, and Ron Daniel. 2016. Applying universal schemas for domain specific ontology expansion. In Jay Pujara, Tim Rocktäschel, Danqi Chen, and Sameer Singh, editors, *Proceedings of the 5th Workshop on Automated Knowledge Base Construction, AKBC@NAACL-HLT 2016, San Diego, CA, USA, June 17, 2016*, pages 81–85. The Association for Computer Linguistics.
- Johannes Leveling. 2015. Tagging of temporal expressions and geological features in scientific articles. In *Proceedings of the 9th Workshop on Geographic Information Retrieval, GIR ’15*, pages 6:1–6:10, New York, NY, USA. ACM.
- Mausam Mausam. 2016. Open information extraction systems and downstream applications. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI’16*, pages 4074–4077. AAAI Press.
- Hoifung Poon, Kristina Toutanova, and Chris Quirk. 2014. Distant supervision for cancer pathway extraction from text. In *Pacific Symposium on Biocomputing Co-Chairs*, pages 120–131. World Scientific.
- Chris Quirk and Hoifung Poon. 2017. Distant supervision for relation extraction beyond the sentence boundary. In Mirella Lapata, Phil Blunsom, and Alexander Koller, editors, *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 1: Long Papers*, pages 1171–1182. Association for Computational Linguistics.
- Rudolf Schneider, Tom Oberhauser, Tobias Klatt, Felix A. Gers, and Alexander Löser. 2017. Analysing errors of open information extraction systems. In *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*, pages 11–18. Association for Computational Linguistics.