

AnnIE: An Annotation Platform for Constructing Complete Open Information Extraction Benchmark

Niklas Friedrich¹, Kiril Gashteovski², Mingying Yu^{1,2}, Bhushan Kotnis²,
Carolyn Lawrence², Mathias Niepert^{2,3,4}, Goran Glavas^{1,4}

¹ University of Mannheim, Mannheim, Germany

² NEC Laboratories Europe, Heidelberg, Germany

³ University of Stuttgart, Stuttgart, Germany, ⁴ LMU Munich, Munich, Germany

{nfriedri,minyu,gglavas}@mail.uni-mannheim.de

{firstname.lastname}@neclab.eu

Abstract

Open Information Extraction (OIE) is the task of extracting facts from sentences in the form of relations and their corresponding arguments in schema-free manner. Intrinsic performance of OIE systems is difficult to measure due to the *incompleteness* of existing OIE benchmarks: ground truth extractions do not group all acceptable surface realizations of the same fact that can be extracted from a sentence. To measure performance of OIE systems more realistically, it is necessary to manually annotate *complete facts* (i.e., clusters of all acceptable surface realizations of the same fact) from input sentences. We propose AnnIE: an interactive annotation platform that facilitates such challenging annotation tasks and supports creation of *complete fact-oriented OIE evaluation benchmarks*. AnnIE is modular and flexible in order to support different use case scenarios (i.e., benchmarks covering different types of facts) and different languages. We use AnnIE to build two complete OIE benchmarks: one with verb-mediated facts and another with facts encompassing named entities. We evaluate several OIE systems on our complete benchmarks created with AnnIE. We publicly release AnnIE under non-restrictive license.¹

1 Introduction

Open Information Extraction (OIE) is the task of extracting relations and their arguments from natural language text in schema-free manner (Banko et al., 2007). Consider the input sentence "Edmund Barton, who was born in Australia, was a judge". Without the use of a pre-specified schema, an OIE system should extract the triples ("Edmund Barton"; "was born in"; "Australia") and ("Edmund Barton"; "was"; "judge"). The output of OIE systems is used in many downstream tasks, including open link prediction (Broscheit et al., 2020), automated knowledge base construction (Gashteovski

et al., 2020), question answering (Khot et al., 2017) and text summarization (Xu and Lapata, 2021).

Intrinsic evaluation of OIE systems is done either manually (Mausam et al., 2012; Pal et al., 2016) or with the use of evaluation benchmarks (Stanovsky and Dagan, 2016; Bhardwaj et al., 2019). While manual evaluations are usually of higher quality, they are expensive and time consuming. Automated benchmark evaluations are faster and more economic than manual OIE evaluations (Hohe-necker et al., 2020), but are less reliable than human judgments of extraction correctness (Zhan and Zhao, 2020), because they are based on approximate token-level matching of system extractions against ground truth extractions. The main shortcoming of existing OIE benchmarks is their *incompleteness*: they do not exhaustively list all acceptable surface realizations of the same piece of information (i.e., same fact) and, because of this, resort to unreliable scoring functions based on token-level matching between system and gold extractions (Schneider et al., 2017).

Obtaining complete manual OIE annotations is, however, very difficult and time-consuming. Annotating a *complete* OIE benchmark requires human annotators to write *all* possible combinations of extractions expressing the same fact (i.e., exhaustively list all acceptable surface realizations of the same fact; see Section 3). To facilitate and speed up this process, we introduce AnnIE, a dedicated annotation tool for constructing complete fact-oriented OIE benchmark. AnnIE facilitates the annotation process by (1) highlighting the tokens of interest (e.g., for verb-mediated extractions, it highlights the verbs, which are candidates for head words of predicates); (2) providing web-based interface for annotating triples and grouping them into *fact synsets*, i.e., groups of informationally equivalent extractions (Section 3). To the best of our knowledge, AnnIE is the first publicly-available annotation platform for constructing OIE benchmarks.

¹<https://github.com/nfriedri/annie-annotation-platform>

We showcase AnnIE² by creating two complete fact-based OIE benchmarks: (1) benchmark of verb-mediated facts on English, German, Chinese, Galician, Arabic and Japanese, making this gold data the first such OIE resource on languages other than English; (2) benchmark for facts associating named entities (for English only). We then benchmark several state-of-the-art OIE systems on these fact-based benchmarks and demonstrate that they are significantly less effective than indicated by existing OIE benchmarks that use token-level scoring. We hope that AnnIE motivates the creation of many more fact-based (as opposed to token-level) OIE evaluation benchmarks.

2 Related Work

2.1 Evaluation of OIE Systems

OIE systems are evaluated either manually (Mausam et al., 2012; Pal et al., 2016; Gashteovski et al., 2019), w.r.t. a downstream task (Mausam, 2016; Lin et al., 2020), or with the use of evaluation benchmarks (Stanovsky and Dagan, 2016; Bhardwaj et al., 2019). Manual evaluations are usually of higher quality because they are performed by one or more expert annotators (Del Corro and Gemulla, 2013). They are, however, expensive and time consuming, which makes the development of OIE systems very slow. On the other hand, downstream evaluation of OIE systems is faster, but provides insights only about their performance w.r.t. particular tasks and does not provide insights on the intrinsic (i.e., task-agnostic) correctness of the extractions. Finally, using evaluation benchmarks is both task-agnostic and fast, though current benchmarks might contain noise (Zhan and Zhao, 2020). Moreover, current benchmarks suffer from incompleteness; i.e., they are not designed in a manner that aims to contain all possible extractions from an input sentence. Therefore, they rely on lenient token-overlap based evaluation, which could result in misleading results (Lechelle et al., 2019). To address this, we move away from such token-based evaluations and move towards fact-based evaluation (Section 3).

2.2 Annotation Tools

To facilitate the annotation process of NLP tasks, many interactive annotation tools have been designed. Such work covers tasks like sequence labelling (Lin et al., 2019; Lee et al., 2020), coreference resolution (Bornstein et al., 2020) and

²Video demo: <https://youtu.be/2wn75U8Lc5w>

treebank projection across languages (Akbik and Vollgraf, 2017). For annotating OIE extractions, however, there are no publicly available tools. The two commonly used benchmarks—OIE2016 (Stanovsky and Dagan, 2016) and CaRB (Bhardwaj et al., 2019)—only provide annotated data and no dedicated annotation tool. OIE2016 uses a dataset from a similar task (QA-SRL), which is then automatically ported to OIE. This approach does not require an annotation tool, but the quality of the benchmark (i.e., ground truth extractions) decreases due to the automatic label projection (Zhan and Zhao, 2020). CaRB addresses this issue by sampling from the same input sentences used by OIE2016, and then crowdsourcing manual extractions. However, their annotation OIE interface has four major limitations: (1) it cannot be used to create complete fact-based OIE benchmarks (Section 3), i.e., it does not allow for different extractions (e.g., triples) that correspond to the same fact; this leads to incomplete annotations and unreliably lenient token-overlap-based evaluation measures; (2) it focuses only on one type of OIE (verb-mediated extractions); (3) it is not publicly available; (4) it does not support annotations for languages other than English.

3 Fact-Based OIE Evaluation

Due to their incompleteness, previous benchmarks lack clarity about whether an extraction indeed represents a correct fact or not. In particular, given a system extraction, they do not assign a binary score (correct/incorrect), but rather calculate per-slot token overlap scores. Consider, for example, the input sentence from Table 1 and the scores that the recent OIE benchmark CaRB (Bhardwaj et al., 2019) assigns to extractions t_1 to t_3 . Because all tokens for each slot for $t_1 - t_3$ are also present in the gold extraction, CaRB credits these extractions with a perfect precision score, even though the extractions clearly state incorrect facts. In similar vein, the CaRB recall score of the extraction t_4 is lower than the recall score of t_3 , even though t_4 captures the correct core fact and t_3 does not.

To address these issues, we propose moving away from such lenient token-overlap scoring and going towards fact-level exact matching. To this end, we propose an evaluation framework, dubbed BenchIE (Gashteovski et al., 2022), for OIE evaluation based on facts, not tokens. Here, the annotator is instructed to *exhaustively* list all possible surface

Input sentence: "Sen. Mitchell is confident he has sufficient votes to block such a measure with procedural actions."						
CaRB gold extraction: ("Sen. Mitchell"; "is confident he has"; "sufficient votes to block ...procedural actions")						
	Input OIE extraction			CaRB (P / R)		Fact-based
t_1	("Sen. Mitchell"; "is confident he has"; "sufficient")			1.00	0.44	0
t_2	("Sen. Mitchell"; "is confident he has"; "sufficient actions")			1.00	0.50	0
t_3	("Sen. Mitchell"; "is confident he has"; "sufficient procedural actions")			1.00	0.56	0
t_4	("Sen. Mitchell"; "is confident he has"; "sufficient votes")			1.00	0.50	1

Table 1: Difference in scores between CaRB and fact-based evaluation. For the input sentence, CaRB provides only one extraction which covers all the words in the sentence. Then, for each input OIE extraction (from t_1 to t_4) it calculates token-wise precision and recall scores w.r.t. the golden annotation. Fact-based evaluation (with all acceptable extractions of the fact exhaustively listed) allows for exact matching against OIE extractions.

realizations of the same fact, allowing for a binary judgment (correct/incorrect) of correctness of each extraction (it either exactly matches some of the acceptable gold realizations of some fact or it does not match any). The example in Table 2 illustrates the concept of a *fact synset*: a collection of all acceptable extractions for the same fact (i.e., same piece of knowledge).

Because benchmarks based on fact synsets are supposed to be complete, a system OIE extraction is considered correct if and only if it *exactly* matches any of the gold extractions from any of the fact synsets. The number of *true positives* (TPs) is the number of fact synsets (i.e., different facts) “covered” by at least one system extraction. This way, a system that extracts N different triples of the same fact, will be rewarded only once for the correct extraction of the fact. False negatives (FNs) are then fact synsets not covered by any of the system extractions. Finally, each system extraction that does not exactly match any gold triple (from any synset) is counted as a false positive (FP). We then compute *Precision*, *Recall*, and F_1 score from TP, FP, and FN in the standard fashion. For more details on the evaluation framework, see (Gashteovski et al., 2022).

4 AnnIE: Platform Description

AnnIE is a web-based platform that facilitates manual annotations of fact-based OIE benchmarks. In this section, we discuss: (1) the functionality of highlighting tokens of interest; (2) how AnnIE facilitates creation of complete fact-based OIE benchmarks; (3) AnnIE’s software architecture; and (4) AnnIE’s web interface and its multilingual support.

4.1 Tokens of Interest

One of the key functionalities of AnnIE is its ability to highlight *tokens of interest* – tokens that com-

monly constitute parts of extractions of interest. For example, most OIE systems focus on extracting verb-mediated triples (Angeli et al., 2015; Koluru et al., 2020b). In such case, *verbs* clearly represent tokens of interest and are candidates for head words of fact predicates. Other example of tokens of interest may be named entities, which could be useful for extracting information from domain-specific text. There has been prior work on extracting open information from specific domains, including the biomedical (Wang et al., 2018), legal (Siragusa et al., 2018) and scientific domain (Lauscher et al., 2019). In this work, it is important to extract open relations between named entities. Accordingly, highlighting mentions of named entities then facilitates manual extraction of the type of facts that the benchmark is supposed to cover (i.e., relations between named entities). AnnIE allows the user to define a custom function that yields the tokens of interest from the input sentence and then highlights these tokens for the annotator with a background color (Figure 2).

To further facilitate the manual annotations, future versions of AnnIE could also include recommendations for whole OIE triples (e.g., by recommending high-confidence extractions from already existing OIE systems) or for slots (e.g., given a subject, recommend a potential relation; or given a relation, recommend candidates for the arguments). We leave such improvements for future work.

4.2 Annotating Fact Synsets

Given a sentence with highlighted tokens of interest, the annotator can start constructing fact synsets. Fact synsets are clusters of fact-equivalent extractions. AnnIE currently supports only the annotation of triples: for each extraction/triple the user first selects which slot she wants to annotate (subject, predicate, or object) and then selects the tokens

Input sentence: "Sen. Mitchell is confident he has sufficient votes to block such a measure with procedural actions."			
f_1	("Sen. Mitchell" "he";	"is";	"confident [he has sufficient ... actions]")
f_2	("Sen. Mitchell" "he"; ("Sen. Mitchell" "he";	"is confident he has"; "is confident he has";	"sufficient votes"; "suff. votes to block [such] [a] measure")
f_3	("Sen. Mitchell" "he"; ("Sen. Mitchell" "he"; ("Sen. Mitchell" "he";	"is conf. he has sufficient votes to block"; "is confident he has ... to block [such]"; "is confident he has ... to block [such] [a]";	"[such] [a] measure"; "[a] measure"; "measure")
f_4	("Sen. Mitchell" "he"; ("Sen. Mitchell" "he";	"is conf. he has ... [such] [a] measure with"; "is confident he has ... [such] [a] measure";	"procedural actions"; "with procedural actions")

Table 2: Example sentence with four *fact synsets* (f_1 – f_4). We account for entity coreference and accept both "Sen. Mitchell" and "he" as subjects: the delimiter "|" is a shorthand notation for different extractions. In the same vein, square brackets ([]) are a shorthand notation for multiple extractions: triples both with and without the expression(s) in the brackets are considered correct.

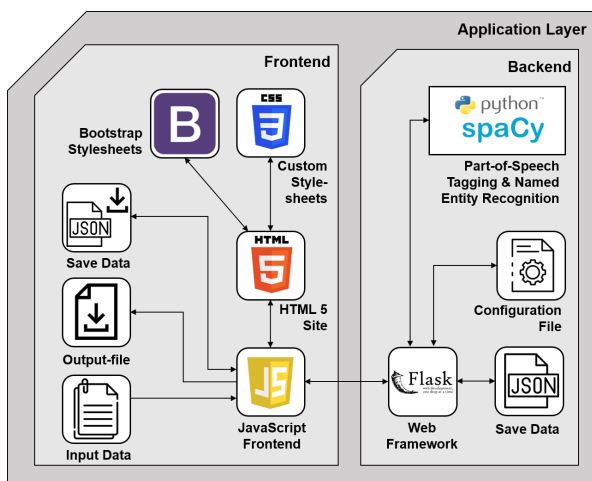


Figure 1: Software architecture of AnnIE.

that constitute that slot. Each token (part of one of the three slots) can additionally be marked as “optional”, which means that the tool will create variants of that extraction both with and without those tokens. Once a triple is fully denoted (i.e., tokens for all three slots selected), the annotator chooses whether (1) the triple is a different variant of an already existing fact (i.e., existing fact synset), in which case the triple is added to an existing cluster or (2) a first variant of a new fact, in which case a new cluster (i.e., fact synset) is created and the triple added to it. To facilitate this decision, the annotator can at any time review the already existing fact synsets. Figure 3 shows the interface for creating triples and adding them to fact synsets.

4.3 Platform Architecture

AnnIE is a simple local executable web application (Figure 1) that consists of a backend layer and a frontend layer. It starts a local server that provides a user interface accessible from any browser that supports JavaScript.

Backend. AnnIE’s backend server is based on

Flask³, a popular light-weight Python web framework. We implemented HTTP endpoints for receiving requests and sending responses. Additionally, Flask hosts the frontend (HTML and CSS) files as a web server. The Flask server interacts with (1) the NLP library SpaCy⁴ (which we employ for POS-tagging and NER, in service of highlighting tokens of interest); (2) a configuration file; (3) data files on the local hard drive and (4) the frontend (i.e., the web interface). AnnIE’s backend is highly modularized, so that any component may easily be further customized or replaced with a different module. For example, the SpaCy-based NLP module (for POS-tagging and NER) can easily be replaced with any other NLP toolkit, e.g., Stanza⁵ (Qi et al., 2020). AnnIE is also easily customizable through a configuration file, where the user may specify the types of tokens to be highlighted or select colors for highlighting. The I/O module expects the input (a collection of sentences for OIE annotation) to be in JSON format and saves the annotated fact synsets in JSON as well.

Frontend. The application frontend is implemented in JavaScript and based on the Bootstrap library and custom CSS stylesheets. We adopt model-view-controller (MVC) architecture for the frontend: it uses a data structure (i.e., *model*) capturing the entire annotation process (i.e., information about the annotation, loaded text file, current triple in annotation, etc.; you can find more details in the Appendix, Section A.2). Based on the current state of the model, the frontend renders the interface (i.e., *view*) by enabling and disabling particular annotation functionality. The controller connects the two: it renders the view based on the current state of the model. We implemented

³<https://github.com/pallets/flask>

⁴<https://spacy.io/>

⁵<https://stanfordnlp.github.io/stanza/>




Figure 2: Highlighting tokens of interest. In this example, tokens of interest are verbs and named entities.

additional I/O scripts that complement the core functionality of the main controller. These scripts handle the formatting of the output file as well as the loading of the data from the input files. Saving or loading data is customizable: one merely needs to overwrite the `load()` and `save()` methods of the controller.

Data Flow. Upon selection of the input file, the data is sent from the frontend to the backend via an HTTP request. In the backend, the sentence is then tokenized, POS-tagged, and processed for named entities with the NLP module (we rely on `SpaCy`). The tokenized and labeled sentence is then sent as a JSON object to the frontend, where each token is displayed as one button (Figure 2). The default version of the tool allows the user to choose (in the configuration file) between four coloring schemes for highlighting token buttons.

4.4 Web Interface and Language Support

The user can start from scratch by uploading a text file that contains unannotated sentences, or load previously saved work (JSON file). For each sentence, the user starts from a full set of sentence tokens with highlighted tokens of interest (Figure 2). The user then constructs triples and places them into fact synsets (Figure 3). At any point during the annotation, the user can generate human-readable output of the annotated extractions and download it as a text file in a tab-separated format (Figure 4). Alternatively, the user can save the annotation progress as a JSON file that can later be loaded in order to continue annotating.

AnnIE supports OIE annotations for sentences in any language supported by its NLP module (i.e., available POS-tagging and NER models). By default, AnnIE relies on `SpaCy` and can therefore support creation of OIE benchmarks for all languages for which `SpaCy` provides POS-tagging and NER models. Section A.3 from the appendix provides details about how this module can be adjusted to the user’s preference.



Figure 3: Manual labeling of OIE triples. The user selects tokens from the tokenized input sentence and places them into the correct slot: **subject** (green), **predicate** (yellow) or **object** (blue). Then, the user adds the extracted triple either to an active fact cluster (i.e., fact synset) or to a new one. The user can also select which tokens are optional by clicking the "Optional" button on an active token selection. For larger version of the same figure, see Figure 11 in Appendix A.6.

5 Demonstration Study

To showcase our tool’s suitability for different OIE scenarios, we generated two complete fact-based OIE benchmarks using AnnIE: (1) a benchmark for verb-mediated facts; (2) a benchmark with facts involving named entities (NEs). We then evaluated several OIE systems and compared their fact-based scores with the token-overlap lenient scores of the existing CaRB benchmark (Bhardwaj et al., 2019).

5.1 Experimental Setup

OIE Systems. We comparatively evaluated several state-of-the-art OIE systems against the gold fact synsets annotated with AnnIE. For OIE on English, we used ClausIE (Del Corro and Gemulla, 2013), Stanford (Angeli et al., 2015), MinIE (Gash-teovski et al., 2017), ROIE (Stanovsky et al., 2018) and OpenIE6 (Kolluru et al., 2020a). For Chinese, German, Galician, Japanese and Arabic, we used the supervised M²OIE (Ro et al., 2020) model, which is based on multilingual BERT (Devlin et al., 2019), trained on large English dataset (Zhan

		EN					ZH	DE	AR	GL	JA
		ClausIE	MinIE	Stanford	ROIE	OpenIE6	M ² OIE	M ² OIE	M ² OIE	M ² OIE	M ² OIE
P	CaRB	0.58	0.45	0.17	0.44	0.48	/	/	/	/	/
	Fact-based	0.50	0.43	0.11	0.20	0.31	0.18	0.09	0.16	0.15	0.00
	Δ	+0.08	+0.02	+0.06	+0.24	+0.17	/	/	/	/	/
R	CaRB	0.53	0.44	0.29	0.60	0.67	/	/	/	/	/
	Fact-based	0.26	0.28	0.16	0.09	0.21	0.10	0.03	0.03	0.06	0.00
	Δ	+0.27	+0.16	+0.13	+0.51	+0.46	/	/	/	/	/
F_1	CaRB	0.56	0.44	0.22	0.51	0.56	/	/	/	/	/
	Fact-based	0.34	0.34	0.13	0.13	0.25	0.13	0.04	0.05	0.09	0.00
	Δ	+0.22	+0.10	+0.09	+0.38	+0.31	/	/	/	/	/

Table 3: Comparison of performance of OIE systems on fact-based v.s. CaRB benchmarks for English (EN), Chinese (ZH), German (DE), Arabic (AR), Galician (GL) and Japanese (JA). Metrics used: precision (P), recall (R) and F_1 score (F_1). Δ is the difference between the CaRB scores and the fact-based scores. **Bold numbers** indicate highest score for English per row (i.e., highest score for P / R / F_1 per benchmark) or highest score difference per row (i.e., highest Δ for P / R / F_1 per benchmark). The fact-based benchmark on English reveals that CaRB overestimates the performance of OIE systems, but with the help of AnnIE it is easily possible to create benchmarks which provide more reliable performance estimates. Creating such fact-based benchmarks for a series of other languages highlights the need for future OIE research to focus on languages other than English.




Figure 4: Human-readable representation of the annotated extractions. Annotations can be downloaded as a human-readable file or as a JSON file (loadable for further annotation with AnnIE).

and Zhao, 2020) and transferred to the target language by means of its multilingual encoder. We trained M²OIE using the implementation and recommended hyperparameter setup from the original work (Ro et al., 2020).

Verb-Mediated Triples. We first evaluate the OIE systems in the most common setup: for verb-mediated facts. In this scenario, OIE system extractions are triples with verb-phrase predicates. We randomly sampled 300 sentences from the CaRB benchmark, and two experts independently annotated them manually for fact synsets with AnnIE (we provide the annotation guidelines in the Appendix, Section A.4). Then, the annotators merged the annotations by resolving the disagreements through a discussion. The annotation effort was approximately two working weeks per annotator. To show that AnnIE is in principle language agnostic, native speakers of German, Chinese, Japanese and Galician translated these 300 sentences to their respective languages. For Arabic, a native speaker managed to translate the first 100 sentences. only

due to limited resources. Then, the native speakers annotated fact synsets in these languages with AnnIE. Due to limited resources, the sentences for these languages were translated and then annotated with OIE extractions by one annotator per language. Finally, we evaluated one OIE system for each language on this benchmark.

NE-centric Triples. We used AnnIE to build a benchmark consisting of facts connecting named entities (NEs): triples in which both subjects and objects are named entities. Since NEs are frequently mentioned in news stories, we selected the sentences for annotation from the NYT10k dataset (Gashteovski et al., 2017), a random sample of 10k sentences from the New York Times corpus (Sandhaus, 2008). We split the NE-centric benchmark in two parts: (1) NE-2: 150 sentences from NYT10k with exactly 2 NEs (as detected by *SpaCy*); (2) NE-3⁺: we sample 150 sentences from NYT10k such that they have 3 or more NE mentions. The annotation guidelines, while similar to those for the verb-mediated triples, differ in two important aspects: (1) the annotator should extract *only* the facts in which both arguments are named entities; (2) besides verb-mediated relations, the annotator was allowed to extract noun-mediated relations too; e.g., ("*Sundar Pichai*"; "*CEO*"; "*Google*").

5.2 Results and Discussion

English OIE. We score the OIE systems against the gold fact synsets produced with AnnIE, using the fact-based evaluation protocol (Section 3). For the verb-mediated extractions, we compare our fact-based evaluation scores against the token-overlap

		ClausIE (4 / 10)	MinIE (26 / 48)	Stanford (12 / 22)	ROIE (2 / 7)	OIE6 (8 / 20)
P	NE-2	0.75	0.58	0.45	0.05	0.38
	NE-3+	0.78	0.54	0.63	0.05	0.32
R	NE-2	0.05	0.23	0.08	0.02	0.05
	NE-3+	0.04	0.13	0.06	0.02	0.03
F_1	NE-2	0.09	0.33	0.13	0.03	0.08
	NE-3+	0.07	0.21	0.11	0.02	0.06

Table 4: Performance of OIE systems on fact-based evaluation on NE-centric triples. NE-2 / NE-3+: results on sentences that contain 2 / 3 or more NEs (labelled with SpaCy). Numbers in brackets below an OIE system name indicate the number of OIE triples on which the evaluation was done for NE-2 / NE-3+. **Bold numbers** indicate highest score per row.

scores of CaRB (Bhardwaj et al., 2019): the results are shown in Table 3. Comparison of Fact-based and CaRB scores indicates that: (1) CaRB largely overestimates the performance of OIE systems; (2) current OIE systems predominantly fail to extract correct facts, which strongly points to the need for creating complete fact-oriented benchmarks, a task that AnnIE facilitates. For more detailed discussion, multi-faceted evaluation and error analysis, see Gashteovski et al. (2022).

Multilingual OIE. Finally, Table 3 shows that the results for OIE systems in languages other than English are significantly worse, which shows that more research is needed in this direction. The results for Chinese OIE seem to be particularly encouraging, as the difference of the F_1 score between some OIE systems in English and M²OIE in Chinese is not too large as it is between English and the other languages. For example, M²OIE in Chinese has the same F_1 score as Stanford’s OIE system and ROIE. We applied the same training strategy for Japanese, but the F_1 score of this OIE system is 0. This indicates that more research is needed for Japanese in defining the problem well and proposing methods for solving it. Nevertheless, we release the gold datasets for OIE in all investigated languages: English, German, Galician, Chinese, Japanese and Arabic; which we believe to be an important resource for research for subsequent multilingual OIE. Figure 12 and Figure 13 from the Appendix show an example sentence and its corresponding OIE annotations in different languages. For a more detailed discussion on the results of multilingual OIE evaluation, see Kotnis et al. (2022).

NE-centric OIE. Table 4 shows the performance

of OIE systems on the NE-centric benchmark. In both subsets of 150 sentences—NE-2 and NE-3+—only a fraction of them contain actual knowledge facts that connect a pair of NEs (59/150 and 97/150 respectively). Because the OIE systems used in the evaluation are not specifically designed to extract NE-centric facts, we make the evaluation fairer by pruning the system extractions before fact-based evaluation: we keep only the triples that contain in both subject and object NEs found among subjects and objects of gold extractions. In other words, we primarily test whether the OIE systems extract acceptable predicates between NEs between which there is a predicate in the gold standard. The results show that the current OIE systems extract very few NE-centric triples (e.g., ClausIE extracts only 4 triples for the NE-2 dataset and 10 for the NE-3+ dataset, see Table 4). Because of this, one should interpret the results in Table 4 with caution. This evaluation, nonetheless shows that the current OIE systems are not ill-suited for a NE-centric OIE task, warranting more research efforts in this direction.

6 Conclusions

Exhaustively annotating all acceptable OIE triples is a tedious task, but important for realistic intrinsic evaluation of OIE systems. To support annotators, we introduced AnnIE: annotation tool for constructing comprehensive evaluation benchmarks for OIE. AnnIE allows custom specification of tokens of interests (e.g., verbs) and is designed for creating fact-oriented benchmarks in which the fact-equivalent—yet superficially differing—extractions are grouped into fact synset. AnnIE’s lightweight architecture, easy installation and customizable components make it a practical solution for future OIE annotation.

Acknowledgments

We thank the anonymous reviewers for their invaluable feedback and support.

References

- Alan Akbik and Roland Vollgraf. 2017. *The Projector: An Interactive Annotation Projection Visualization Tool*. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP): System Demonstrations*, pages 43–48.
- Gabor Angeli, Melvin Jose Johnson Premkumar, and Christopher D. Manning. 2015. *Leveraging Linguistic Structure For Open Domain Information Extrac-*

- tion. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 344–354.
- Michele Banko, Michael J Cafarella, Stephen Soderland, Matthew Broadhead, and Oren Etzioni. 2007. **Open Information Extraction from the Web**. In *Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI)*, pages 2670–2676.
- Sangnie Bhardwaj, Samarth Aggarwal, and Mausam Mausam. 2019. **CaRB: A Crowdsourced Benchmark for Open IE**. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6263–6268.
- Aaron Bornstein, Arie Cattan, and Ido Dagan. 2020. **CoRefi: A Crowd Sourcing Suite for Coreference Annotation**. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP): System Demonstrations*.
- Samuel Broscheit, Kiril Gashteovski, Yanjie Wang, and Rainer Gemulla. 2020. **Can We Predict New Facts with Open Knowledge Graph Embeddings? A Benchmark for Open Link Prediction**. In *Proceedings of the Annual Meeting on Association for Computational Linguistics (ACL)*, pages 2296–2308.
- Luciano Del Corro and Rainer Gemulla. 2013. **ClauseIE: Clause-Based Open Information Extraction**. In *Proceedings of the International World Wide Web Conferences (WWW)*, pages 355–366.
- 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 Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT)*, pages 8554–8565.
- Kiril Gashteovski, Rainer Gemulla, and Luciano Del Corro. 2017. **MinIE: Minimizing Facts in Open Information Extraction**. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2630–2640.
- Kiril Gashteovski, Rainer Gemulla, Bhushan Kotnis, Sven Hertling, and Christian Meilicke. 2020. **On Aligning Openie Extractions with Knowledge Bases: A Case Study**. In *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems*, pages 143–154.
- Kiril Gashteovski, Sebastian Wanner, Sven Hertling, Samuel Broscheit, and Rainer Gemulla. 2019. **OPIEC: An Open Information Extraction Corpus**. In *In Proceedings of the Conference on Automated Knowledge Base Construction (AKBC)*.
- Kiril Gashteovski, Mingying Yu, Bhushan Kotnis, Carolin Lawrence, Mathias Niepert, and Goran Glavaš. 2022. **BenchIE: A Framework for Multi-Faceted Fact-Based Open Information Extraction Evaluation**. In *Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Patrick Hohenecker, Frank Mtumbuka, Vid Kocijan, and Thomas Lukasiewicz. 2020. **Systematic Comparison of Neural Architectures and Training Approaches for Open Information Extraction**. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8554–8565.
- Tushar Khot, Ashish Sabharwal, and Peter Clark. 2017. **Answering Complex Questions Using Open Information Extraction**. In *Proceedings of the Annual Meeting on Association for Computational Linguistics (ACL)*, pages 311–316.
- Keshav Kolluru, Vaibhav Adlakha, Samarth Aggarwal, Soumen Chakrabarti, et al. 2020a. **OpenIE6: Iterative Grid Labeling and Coordination Analysis for Open Information Extraction**. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3748–3761.
- Keshav Kolluru, Samarth Aggarwal, Vipul Rathore, Mausam, and Soumen Chakrabarti. 2020b. **IMoJIE: Iterative Memory-Based Joint Open Information Extraction**. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 5871–5886.
- Bhushan Kotnis, Kiril Gashteovski, Daniel Oñoro Rubio, Ammar Shaker, Vanesa Rodriguez-Tembras, Makoto Takamoto, Mathias Niepert, and Carolin Lawrence. 2022. **MILLIE: Modular & Iterative Multilingual Open Information Extraction**. In *Proceedings of the Annual Meeting on Association for Computational Linguistics (ACL)*.
- Anne Lauscher, Yide Song, and Kiril Gashteovski. 2019. **MinSciE: Citation-centered Open Information Extraction**. In *2019 ACM/IEEE Joint Conference on Digital Libraries (JCDL)*, pages 386–387. IEEE.
- William Lechelle, Fabrizio Gotti, and Phillippe Langlais. 2019. **WiRe57: A Fine-Grained Benchmark for Open Information Extraction**. In *Proceedings of the Linguistic Annotation Workshop (LAW@ACL)*, pages 6–15.
- Dong-Ho Lee, Rahul Khanna, Bill Yuchen Lin, Jamin Chen, Seyeon Lee, Qinyuan Ye, Elizabeth Boschee, Leonardo Neves, and Xiang Ren. 2020. **LEAN-LIFE: A Label-Efficient Annotation Framework Towards Learning from Explanation**. In *Proceedings of the Annual Meeting on Association for Computational Linguistics (ACL Demo)*, pages 372–379.
- Bill Yuchen Lin, Dong-Ho Lee, Frank F Xu, Ouyuan Lan, and Xiang Ren. 2019. **AlpacaTag: An Active Learning-based Crowd Annotation Framework for Sequence Tagging**. In *Proceedings of the Annual*

- Meeting of the Association for Computational Linguistics (ACL): System Demonstrations*, pages 58–63.
- Xueling Lin, Haoyang Li, Hao Xin, Zijian Li, and Lei Chen. 2020. [KB Pearl: A Knowledge Base Population System Supported by Joint Entity and Relation Linking](#). In *Proceedings of the Very Large Data Base Endowment (PVLDB)*, pages 1035–1049.
- Mausam. 2016. [Open Information Extraction Systems and Downstream Applications](#). In *Proceedings of the twenty-fifth international joint conference on artificial intelligence*, pages 4074–4077.
- Mausam, Michael Schmitz, Robert Bart, Stephen Soderland, and Oren Etzioni. 2012. [Open Language Learning for Information Extraction](#). In *Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 523–534.
- Harinder Pal et al. 2016. [Demonyms and Compound Relational Nouns in Nominal Open IE](#). In *Proceedings of the 5th Workshop on Automated Knowledge Base Construction*, pages 35–39.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D Manning. 2020. [Stanza: A Python Natural Language Processing Toolkit for Many Human Languages](#). In *Association for Computational Linguistics (ACL)*, page 101–108.
- Youngbin Ro, Yukyung Lee, and Pilsung Kang. 2020. [Multi-IOIE: Multilingual Open Information Extraction Based on Multi-Head Attention with BERT](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1107–1117, Online. Association for Computational Linguistics.
- Evan Sandhaus. 2008. [The new york times annotated corpus](#). *Linguistic Data Consortium, Philadelphia*, 6(12):e26752.
- 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, Copenhagen, Denmark. Association for Computational Linguistics.
- Giovanni Siragusa, Rohan Nanda, Valeria De Paiva, and Luigi Di Caro. 2018. [Relating Legal Entities via Open Information Extraction](#). In *Research Conference on Metadata and Semantics Research*, pages 181–187. Springer.
- Gabriel Stanovsky and Ido Dagan. 2016. [Creating a Large Benchmark for Open Information Extraction](#). In *Proceedings of the International Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2300–2305.
- Gabriel Stanovsky, Julian Michael, Luke Zettlemoyer, and Ido Dagan. 2018. [Supervised Open Information Extraction](#). In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 885–895.
- Xuan Wang, Yu Zhang, Qi Li, Yinyin Chen, and Jiawei Han. 2018. [Open Information Extraction with Meta-pattern Discovery in Biomedical Literature](#). In *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics (BCB)*, pages 291–300.
- Yumo Xu and Mirella Lapata. 2021. [Generating Query Focused Summaries from Query-free Resources](#). In *Annual Meeting of the Association for Computational Linguistics (ACL)*, page 6096–6109.
- Junlang Zhan and Hai Zhao. 2020. [Span Model for Open Information Extraction on Accurate Corpus](#). In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 9523–9530.

```

/*Add YOURLABEL class to css file */
.btn-YOURLABEL {
  color: white;
  background-color: #87FF00;
  border-color: #87FF00;
  opacity: 0.5;
}
.btn-YOURLABEL:hover,
.btn-YOURLABEL:focus {
  color: white;
  background-color: #75DE00;
  border-color: #75DE00;
}

```

Figure 5: Add new class to css

A Appendix

A.1 Highlighting Functions

The tool is designed to make customizations rather easily. For adjusting the color scheme, the color “hex”-values inside the style.css files need to be set to the desired color codes. This needs to be done to the button itself with the desired color as well as to the “hover” property of the button, where usually a darker version of the same color is used. If complete new labels are introduced to the tool, the css needs to include an appropriate class to handle these as shown in Figure 5.

In case any new coloring schemes are required, these can either be entered additionally or exchanged against the standard functions implementing the scheme above. The different colorings are applied using the functions `fullColoring()`, `verbColoring()`, `namedEntitiesColoring()`, and `noneColoring()` inside `GraphicalInterface.js`. These functions can be adjusted by changing the switch statements handling which tokens need to be colored depending on their label. There, new cases can simply be added or superfluous colored labels can be removed. Another option is to add new coloring functions. These can rely on the provided ones. They simply need to be registered to the tool by being added to the first switch statement inside the function `createTaggedContent()`. An example of such a function is given in Figure 6, while the “register” procedure is shown in Figure 7.

In both cases, additionally, the function `downgrade()` needs to be adjusted accordingly to the above-mentioned changes to ensure that the buttons can be selected and deselected properly. This step is shown in Figure 8.

A.2 Data model structure used in the frontend

The data model structure used in the frontend is shown on Figure 9.

A.3 Multilinguality

By extending the definition of the function `read_config_file()` inside the backend file `tokenizer.py`, further languages can be included. Therefore, SpaCy simply needs to be forced to load the appropriate language model. For details, see code snippet showing how to adapt the `tokenizer.py` script to accept further language models in Figure 10.

A.4 Annotation Guidelines (English)

A.4.1 General Principle

The annotator should manually extract verb-mediated triples from a natural language sentence. Each triple should represent two entities or concepts, and the verb-mediated relation between them. For example, from the input sentence “*Michael Jordan, who is a former basketball player, was born in Brooklyn.*”, there are three entities and concepts—*Michael Jordan*, *former basketball player* and *Brooklyn*—which are related as follows: (“*Michael Jordan*”; “*is*”; “*former basketball player*”) and (“*Michael Jordan*”; “*was born in*”; “*Brooklyn*”).

Once the triple is manually extracted, it should be placed into the correct fact synset (see Section A.4.2).

A.4.2 Fact Synsets

Once a triple is manually extracted, the annotator should place the triple into its corresponding fact synset (details about fact synsets in Section 3). In case there is no existing fact synset for the manually extracted triple, the annotator should create one and place the triple in that synset.

Coreference. The annotator should place extractions that refer to the same entity or concept under the same fact synset. Consider the following input sentence: “*His son, John Crozie, was an aviation pioneer.*”; The following triples should be placed in the same fact synset:

- (“*His son*”; “*was*”; “[*an*]⁶ *aviation pioneer*”)
- (“*J. Crozie*”; “*was*”; “[*an*] *aviation pioneer*”)

⁶words in square brackets indicate optional tokens (see Section A.4.3)

```

function yourLabelColoring(labelText, labelPos, index) {
  let output = '';
  switch (labelPos) {
    //Remove cases not necessary:
    case 'NOUN': output += `<button class="btn btn-noun ml-1 mb-1" id="posLabel-${index}">`; break;
    case 'VERB': output += `<button class="btn btn-verb ml-1 mb-1" id="posLabel-${index}">`; break;
    case 'ADJ': output += `<button class="btn btn-adjective ml-1 mb-1" id="posLabel-${index}">`; break;
    case 'ORG':
    case 'GPE':
    case 'LOC':
    case 'PERSON':
      output += `<button class="btn btn-namedEntity ml-1 mb-1" id="posLabel-${index}">`; break;
    // Stop removing
    // Insert following lines and exchange variable YOURLABEL
    case 'YOURLABEL':
      output += `<button class="btn btn-YOURLABEL ml-1 mb-1" id="posLabel-${index}">`;
      break;
    // End Insertion
    default: output += `<button class="btn btn-secondary ml-1 mb-1" id="posLabel-${index}">`; break;
  }
  if (showTag) {
    output += `<text>${labelText}</text> <span class="badge badge-secondary">${index}</span><br><pos>${labelPos}</pos></button>`;
  } else {
    output += `<text>${labelText}</text> <span class="badge badge-secondary">${index}</span><pos hidden>${labelPos}</pos></button>`;
  }
  return output;
}

```

Figure 6: Example of a new coloring function

```

function createTaggedContent(words) {
  var output = "";
  for (var i = 0; i < words.length; i++) {
    //console.log(i)
    let labelText = words[i].text;
    let labelPos = words[i].posLabel;
    let type = words[i].type;
    let index = words[i].index;
    if (type == '') {
      switch (coloring) {
        case 'full': output += fullColoring(labelText, labelPos, index); break;
        case 'verbs': output += verbColoring(labelText, labelPos, index); break;
        case 'named-entities': output += namedEntitiesColoring(labelText, labelPos, index); break;
        case 'none': output += noneColoring(labelText, labelPos, index); break;
        // "Register" new created coloring function, replace "yourLabelColoring" with function name
        case 'YOURLABEL':
          output += yourLabelColoring(labelText, labelPos, index); break;
        // End
        default:
          output += verbColoring(labelText, labelPos, index);
          break;
      }
    }
    // ...
  }
}

```

Figure 7: "Register" new coloring function

```

function downgrade(targetElement) {
  // ...
  else {
    //Insert if-statement for YOURLABELCOLORING
    if (coloring == 'YOURLABELCOLORING') {
      switch (posLabel) {
        case 'NOUN': targetElement.className = "btn btn-noun ml-1 mb-1"; break;
        case 'VERB': targetElement.className = "btn btn-verb ml-1 mb-1"; break;
        case 'ADJ': targetElement.className = "btn btn-adjective ml-1 mb-1"; break;
        case 'ORG':
        case 'LOC':
        case 'PERSON':
        case 'GPE': targetElement.className = "btn btn-namedEntity ml-1 mb-1"; break;
        // Add case for YOURLABEL
        case 'YOURLABEL':
          targetElement.className = "btn btn-YOURLABEL ml-1 mb-1";
          break;
        //END
        default:
          targetElement.className = "btn btn-secondary ml-1 mb-1";
          break;
      }
    }
    // ...
  }
}

```

Figure 8: Necessary adjustments to downgrade ()




Figure 9: Data model structure used in the frontend

```
def read_config_file(self):
    """
    """
    if configs["Language"] == "English":
        try:
            self.nlp = spacy.load("en_core_web_sm")
        except:
            os.system('python -m spacy download en_core_web_sm')
            self.nlp = spacy.load("en_core_web_sm")
            print("Successfully loaded language: ENGLISH")
    """INSERT """
    if configs["Language"] == "YOURLANGUAGE":
        try:
            self.nlp = spacy.load("SPACY_LANGUAGE_MODULE")
        except:
            os.system('python -m spacy download SPACY_LANGUAGE_MODULE')
            self.nlp = spacy.load("SPACY_LANGUAGE_MODULE")
            print("Successfully loaded language: SPACY_LANGUAGE_MODULE")
    """END"""
```

Figure 10: Code snippet showing how to add further language support

because "His son" and "John Crozie" refer to the same entity.

Token placements. The annotator should consider placing certain tokens in different slots, without damaging the meaning of each the fact. Consider the input sentence "Michael Jordan was born in Brooklyn.". There is one fact synset and its corresponding triples:

f_1 ("M. J."; "was born in"; "Brooklyn")
 ("M. J."; "was born"; "in Brooklyn")

In the first triple, the preposition "in" is in the relation, while in the second it is in the object. The annotator should allow for such variations, because OIE systems should not be penalized for placing such words in different slots.

A.4.3 Optional Tokens

If possible, the annotator should label as *optional* all tokens that can be omitted in an extraction without damaging its semantics. Such tokens include determiners (e.g., *a, the, an*), honorifics (e.g., *[Prof.] Michael Jordan*) or certain quantities (e.g., *[some] major projects*). The optional tokens are marked with square brackets [].

In what follows, we show examples of considered optional token(s).

Determiners. Unless a determiner is a part of a named entity (e.g., "The Times"), it is considered as optional. For instance, the following triples are considered to be semantically equivalent:

- ("Michael Jordan"; "took"; "the ball")
- ("Michael Jordan"; "took"; "ball")

The annotator, therefore, should annotate ("Michael Jordan"; "took"; "[the] ball"), where the optional token is in square brackets.

Titles. Titles of people are considered optional; e.g., ("[Prof.] Michael Jordan"; "lives in"; "USA").

Adjectives. The annotator should label adjectives as optional if possible. For example, in the following triple, the adjective *outstanding* can be considered optional: ("Albert Einstein"; "is"; "[an] [outstanding] scientist"). Note that the annotator should be careful not to label adjectives as optional if they are essential to the meaning of the triple. For instance, the adjective *cold* should not be labeled as optional in the triple ("Berlin Wall"; "is infamous symbol of"; "[the] cold war").

Quantities. Certain quantities that modify a noun phrase can be considered as optional; e.g.,

("Mitsubishi"; "has control of"; "[some] major projects").

Words indicating some tenses. The annotator can treat certain verbs that indicate tense as optional. For instance, the word *have* in ("FDA"; "[has] approved"; "Proleukin") can be considered as optional, since both VPs "have approved" and "approved" contain the same core meaning.

Verb phrases. It is allowed for the annotator to mark verb phrases as optional if possible; e.g. ("John"; "[continues to] reside in"; "Berlin").

Passive voice. When possible, if an extraction is in passive voice, the annotator should place its active voice equivalent into the appropriate fact synset. For instance, suppose we have the sentence "The ball was kicked by John.". Then, the fact synset should contain the following triples:

- ("[The] ball"; "was kicked by"; "John")
- ("John"; "kicked"; "[The] ball")

Note that the opposite direction is not allowed. If the sentence was "John kicked the ball.", then the annotator is not allowed to manually extract the triple ("[The] ball"; "was kicked by"; "John") because such extraction contains words that are not originally found in the input sentence ("was" and "by"). These are so-called implicit extractions and we do not consider them (details in Sec. A.4.7)).

Attribution clauses. Extractions that indicate attribution of information should be placed in the same fact synset as the original information statement. For example, the core information of the sentence "Conspiracy theorists say that Barack Obama was born in Kenya." is that Obama was born in Kenya. As indicated by [Mausam et al. \(2012\)](#), it is important not to penalize OIE systems that would also extract the context about the attribution of such information. Therefore, the annotator should include the following triples into the same fact synset: ("Barack Obama"; "was born in"; "Kenya") and ("Conspiracy theorists"; "say that"; "Barack Obama was born in Kenya").

A.4.4 Incomplete Clauses

The annotator should not manually extract incomplete clauses, i.e., triples such that they lack crucial piece of information. Suppose there is the input sentence "He was honored by the river being named after him". The following triple should not be manually extracted: ("He"; "was honored by"; "[the] river"), but the following triples should be: ("He";

"was honored by [the] river being named after"; "him") and ("[the] river"; "being named after"; "him").

A.4.5 Overly Complex Extractions

The annotators should not manually extract overly specific triples, such that their arguments are complex clauses. For instance, for the input sentence "Vaccinations against other viral diseases followed, including the successful rabies vaccination by Louis Pasteur in 1886.", the following triple should not be extracted: ("Vaccinations against other viral diseases"; "followed"; "including the successful rabies vaccination by Louis Pasteur in 1886") because the object is a complex clause which does not describe a single concept precisely, but rather it is composed of several concepts.

A.4.6 Conjunctions

The annotator should not allow for conjunctive phrases to form an argument (i.e., subject or object). Such arguments should be placed into separate extractions (and in separate fact synsets). Consider the sentence "Michael Jordan and Scottie Pippen played for Chicago Bulls.". The annotator should manually extract the following triples:

- ("M. Jordan"; "played for"; "Chicago Bulls")
- ("S. Pippen"; "played for"; "Chicago Bulls")

The annotator should not, however, manually extract ("Michael Jordan and Scottie Pippen"; "played for"; "Chicago Bulls").

A.4.7 Implicit Extractions

We focus on explicit extractions, which means that every word in the extracted triple must be present in the original input sentence. Therefore, implicit extractions—i.e., extractions that contain inferred information which is not found in the sentence explicitly—are not considered. One example implicit extraction is ("Michael Jordan"; "be"; "Prof.") from the input sentence "Prof. Michael Jordan lives in USA.", where the triple infers that Michael Jordan is professor without being explicitly indicated in the sentence (i.e., the word "be" is not present in the input sentence, it is inferred).

A.5 Annotation Guidelines (Chinese)

The annotator followed the same general principles as with the English annotation guidelines (Sec. A.4). Due to the difference of the languages, we slightly adapted the annotation guidelines for the Chinese

language. In what follows, we list those differences.

A.5.1 Articles

Chinese language does not contain articles (i.e., "a", "an", "the"). Therefore, in the manual translation of the sentences, there are no articles in the Chinese counterparts, which also results in labeling such words as optional (for English, see Sec. A.4.3).

A.5.2 Prepositional Phrases within a Noun Phrase

Certain noun phrases with nested prepositional phrase cannot be translated directly into Chinese the same way as in English. For example, suppose we have the phrase "*Prime Minister of Australia*". In Chinese, the literal translation of this phrase would be "*Australia's Prime Minister*". For instance, in the English annotations the sentence "*He was the Prime Minister of Australia*" would have two fact synsets:

f_1 ("*He*"; "*was [the] Pr. Min. of*"; "*Australia*")

f_2 ("*He*"; "*was*"; "*[the] Pr. Min. [of Australia]*")

This is because the fact synset f_1 relates the concepts "*he*" and "*Australia*" with the relation "*was [the] Prime Minister of*", while the second fact synset relates the concepts "*he*" and "*Prime Minister [of Australia]*" with the relation "*was*".

In Chinese language, however, the construction of f_1 would not be possible, because the phrase "*Prime Minister of Australia*" cannot be separated into "*Prime Minister*" and "*Australia*". Therefore, the golden annotation for this particular example in Chinese would be only one fact synset: ("*He*"; "*was*"; "*[Australia's] Prime Minister*"), which is equivalent with f_2 .

A.6 Manual Labeling of OIE Triples

See Figure 11 for a screenshot from ANNIE's GUI, which shows the manual labeling process of OIE triples given an input sentence.

Sentence # 1 Edmund Barton, who was born in Australia, was a judge

Edmund 1 PERSON Barton 2 PERSON , 3 PUNCT who 4 PRON was 5 AUX born 6 VERB in 7 ADP Australia 8 GPE

, 9 PUNCT was 10 AUX a 11 DET judge 12 NOUN

Select subject(s) Select predicate(s) Select object(s) Optional Clear

Selection

Edmund 1 PERSON Barton 2 PERSON was 10 AUX a 11 DET judge 12 NOUN

... Add To Add to new Cluster Add to active Cluster

Cluster #:

Cluster 1

Triple 1

Edmund 1 PERSON Barton 2 PERSON

was 5 AUX born 6 VERB in 7 ADP

Australia 8 GPE

Triple 2

Edmund 1 PERSON Barton 2 PERSON

was 5 AUX born 6 VERB in 7 ADP

Australia 8 GPE

Cluster 2

Triple 1

Edmund 1 PERSON Barton 2 PERSON was 10 AUX a 11 DET judge 12 NOUN

Figure 11: Manual labeling of OIE triples. The user selects tokens from the tokenized input sentence and places them into the correct slot: subject (green), predicate (yellow) or object (blue). Then, the user adds the extracted triple either to an active fact cluster (i.e., fact synset) or to a new one. The user can also select which tokens are optional by clicking the "Optional" button on an active token selection.

Input sentence (EN): He served as the first Prime Minister of Australia and became a founding justice of the High Court of Australia .		
f_1	("He"; "served as"; ("He"; "served";	"[the] [first] Prime Minister [of Australia]" "as [the] [first] Prime Min. [of Australia]"
f_2	("He"; "served as [the] [first] Prime Min. of"; ("He"; "served as [the] [first] Prime Min.;"	"Australia" "of Australia"
f_3	("He"; "became"; ("He"; "became";	"[a] [founding] justice" "[a] [found.] just. of [the] High Court [of Aus.]"
f_4	("He"; "became [a] [founding] justice of"; ("He"; "became [a] [founding] justice";	"[the] High Court [of Australia]" "of [the] High Court [of Australia]"
f_5	("He"; "bec. [a] [found.] just. of [the] H. C. of"; ("He"; "bec. [a] [found.] just. of [the] H. C.;"	"Australia" "of Australia"
Input sentence (DE): Er diente als erster Premierminister von Australien und wurde Gründungsrichter des obersten Gerichts von Australien.		
f_1	("Er"; "diente als"; ("Er"; "diente";	"[erster] Premierminister [von Australien]" "als [erster] Premierminister [von Australien]"
f_2	("Er"; "diente als [erster] Premiermin. von"; ("Er"; "diente als [erster] Premierminister";	"Australien" "von Australien"
f_3	("Er"; "wurde";	"Gründungs. [des obersten Gerichts] [von Aus.]"
f_4	("Er"; "wurde Gründungsrichter";	"[des] obersten Gerichts [von Australien]"
f_5	("Er"; "wurde Gründungs. [des] oberst. Ger.;" ("Er"; "wurde Gründungs. ... Gerichts von";	"von Australien" "Australien"
Input sentence (GL): Serviu como primeiro primeiro ministro de Australia e converteuse nun xuíz fundador do Tribunal Superior de Xustiza de Australia.		
f_1	(/ "Serviu"; (/ "Serviu como";	"como [primeiro] primeiro ministro [de Aus.]" "[primeiro] primeiro ministro [de Aus.]"
f_2	(/ "Serviu como prim. primeiro ministro"; (/ "Serviu como prim. primeiro ministro";	"de Australia" "Australia"
f_3	(/ "converteuse"; (/ "converteuse nun";	"nun xuíz [fund.] do T. Sup. de Xus. [de Aus.]" "xuíz [fund.] do Trib. Sup. de Xus. [de Aus.]"
f_4	(/ "converteuse nun xuíz [fundador] do";	"Tribunal Superior de Xustiza [de Australia]"
f_5	(/ "con. nun xuíz [fund.] do T. S. de X. de"; (/ "con. nun xuíz [fund.] do T. S. de X.;"	"Australia" "de Australia"

Figure 12: Example sentence with five *fact synsets* (f_1 – f_5) in several languages: English (EN), German (DE) and Galician (Galician). Square brackets ([]) are a shorthand notation for multiple extractions: triples both with and without the expression(s) in the brackets are considered correct. For continuation of this figure, see Figure 13, whereas the same input sentence and its corresponding OIE annotations are written in Chinese (ZH), Japanese (JA) and Arabic (AR).

Input sentence (ZH): 他曾担任澳大利亚第一任总理，并成为澳大利亚高等法院的创始法官。		
f_1	("他"; "[曾]担任";	"[澳大利亚][第一任]总理")
f_2	("他"; "成为";	"[澳大利亚高等法院的][创始]法官")
Input sentence (JA): 彼はオーストラリアの初代首相を務め、オーストラリア高等裁判所の創設裁判官になりました。		
f_1	("彼[は]"; "務め";	"[オーストラリア][の][初代]首相を")
	("彼[は]"; "を務め";	"[オーストラリア][の][初代]首相")
f_2	("彼[は]"; "の初代首相[を]務め";	"オーストラリア")
	("彼[は]"; "初代首相[を]務め";	"オーストラリアの")
f_3	("彼[は]"; "なり[まし][た]";	"[創設]裁判官に")
	("彼[は]"; "なり[まし][た]";	"[オーストラリア]高等裁判所の[創設]裁判官に")
f_4	("彼[は]"; "の[創設][裁判]官[に]なり ました";	"[オーストラリア]高等裁判所")
	("彼[は]"; "[創設][裁判]官[に]なりま した";	"[オーストラリア]高等裁判所の")
f_5	("彼は"; "高等裁判所の[創設]裁判 官[に]なり[まし][た]")	"オーストラリア"
Input sentence (AR): أستراليا في العليا للمحكمة مؤسسًا قاضيًا وأصبح لأستراليا وزراء رئيس أول منصب شغل هو .		
f_1	("هو"; "غل";	"[الأستراليا] وزراء رئيس [أول منصب]")
f_2	("هو"; "وزراء رئيس [أول منصب] غل";	"أستراليا")
f_3	("هو"; "أصبح"	"[أستراليا في] العليا للمحكمة مؤسسًا [اضيًا]"
	("هو"; "أصبح"	"[مؤسسًا] [اضيًا]"
f_4	("هو"; "مؤسسًا [اضيًا] وأصبح"	"[أستراليا في] العليا للمحكمة"
f_5	("هو"; "للمحكمة مؤسسًا قاضيًا أصبح العليا";	"أستراليا [في]"

Figure 13: Example sentence and its corresponding OIE annotations in Chinese (ZH), Japanese (JA) and Arabic (AR). This figure is continuation of Figure 12.