

BenchIE^{FL}: A Manually Re-Annotated Fact-Based Open Information Extraction Benchmark

Fabrice Lamarche and Philippe Langlais
RALI, DIRO, Université de Montréal, Canada
fabrice.lamarche@umontreal.ca
felipe@iro.umontreal.ca

Abstract

Open Information Extraction (OIE) is a field of natural language processing that aims to present textual information in a format that allows it to be organized, analyzed and reflected upon. Numerous OIE systems are developed, claiming ever-increasing performance, marking the need for objective benchmarks. *BenchIE* is the latest reference we know of. Despite being very well thought out, we noticed a number of issues we believe are limiting. Therefore, we propose *BenchIE^{FL}*, a new OIE benchmark which fully enforces the principles of *BenchIE* while containing fewer errors, omissions and shortcomings when candidate facts are matched towards reference ones. *BenchIE^{FL}* allows insightful conclusions to be drawn on the actual performance of OIE extractors.

1 Introduction

Open Information Extraction (OIE), the task of extracting organized tuples containing information expressed in a sentence (Yates et al., 2007) has numerous downstream applications ranging from Question Answering (Fader et al., 2014) to Text Comprehension (Stanovsky et al., 2015).

Earlier works evaluated OIE extractors mainly by examining their output and manually determining whether the extracted tuples were expressed in a given sentence. This method lacks the capacity to measure the recall of systems, which led to the creation of OIE benchmarks consisting of annotations of all possible tuples from a corpus and a matching function establishing the concordance between extracted and annotated facts. Our analysis of the most recent benchmark, *BenchIE* (Gashteovski et al., 2022), shows that although very well thought out, its results are noisy and prone to biases, making its conclusion less trustworthy.

Contributions Our main contribution in this work is the release of a new OIE benchmark,

BenchIE^{FL}, that we created by re-annotating *BenchIE*, correcting frequent errors, inconsistencies and methodology limitations, resulting in more concise, precise and pertinent annotations. *BenchIE^{FL}* also benefits from a new matching function that is more flexible, and which captures more valid extractions, thus producing — as we shall see — a fairer ranking of evaluated systems.

In doing so, we produce a number of useful resources¹, including new OIE guidelines, both for sentence annotation and tuple matching. We compare seven OIE extractors — neural and non neural — and show that believed state-of-the-art systems are not necessarily the best. We further demonstrate on three downstream tasks that the scores of extractors on *BenchIE^{FL}* exhibits a much stronger correlation with their performance on each task.

This paper is organized as follows. In section 2, we review some existing OIE systems and benchmarks. We introduce issues with *BenchIE* in section 3, leading us to define guidelines discussed in section 4. In section 5 we characterize *BenchIE^{FL}*, a new version of *BenchIE* we annotated by following our guidelines. Experiments are reported in section 6 and we conclude in section 7.

2 Related Work

2.1 OIE Systems

Open information extraction systems vary in their approaches to the task. Earlier systems are mostly rule-based. These systems, like *ReVerb* (Fader et al., 2011), *ClausIE* (Del Corro and Gemulla, 2013) or *MinIE* (Gashteovski et al., 2017) make use of parts-of-speech tags, syntactical analysis and other grammatical characteristics to derive simple extraction rules. Newer systems are almost exclusively using neural approaches, and be-

¹Resources, guidelines as well as evaluation scripts used for our experiments are available at https://github.com/rali-udem/benchie_fl.git.

longing to one of two distinct categories, either sequence-to-sequence or sequence tagging techniques. Sequence-to-sequence models, like *ImojIE* (Kolluru et al., 2020b) and *M2OIE* (Ro et al., 2020) output the original sentence one word at a time, as well as markers that delimit the arguments and relations, making use of copying and attention mechanisms to output the original sentence words. Neural tagging models like *OpenIE6* (Kolluru et al., 2020a) and *CompactIE* (Fatahi Bayat et al., 2022) use different approaches but mainly focus on identifying relations initially, and then tagging other words as arguments related to identified relations. Models from both categories are typically trained on bootstrapped data from previously released extractors selected with confidence scores and heuristics.

Further information on those extractors and implementation details are provided in Appendix A.

2.2 OIE Benchmarks

The first complete OIE benchmark, *OIE2016* (Stanovsky and Dagan, 2016) was created by automatically transforming two Question Answer Driven Semantic Role Labeling (QA-SRL) datasets (He et al., 2015) into a benchmark comprising 3200 sentences (from the Wall Street Journal and Wikipedia) and 10359 extractions. The matching function used by *OIE2016* matches an extraction and a reference tuple if the grammatical head of both their arguments and their relation are the same.

Lechelle et al. (2019) identified limitations both in the conversion involved in *OIE2016* and its matching function (they demonstrate that it is straightforward to attack this benchmark with a dummy extractor). They released *WiRE57*, a benchmark made up of 57 expertly annotated sentences from Wikipedia and Reuters into 343 tuples, and use **token-level** scoring, meaning that for each annotated tuple, the precision is measured as the proportion of extracted words present in the annotations, and the recall, the proportion of annotated words present in the extractions. *WiRE57* is purposely criticized for its small size.

The authors of *CaRB* (Bhardwaj et al., 2019) proposed crowd-sourcing as a solution to the high cost of manually annotating sentences. They release their reference of 2714 tuples from a subset of 1200 sentences from *OIE2016* annotated thanks to Amazon Mechanical Turk. They also introduce slight modifications to *WiRE57*’s scoring function.

Gashteovski et al. (2022) criticize *CaRB* mainly

in the way it scores system extractions; showing that token-level scoring allows for incorrect extractions to be scored highly both in precision and recall. To counter this, they propose to use a conservative **exact matching** function, meaning that only extractions that are identical to an annotated tuple will count. This notion of exact match works because of the *fact synset* principle they introduce : instead of annotating only one formulation of a given fact, they aim to list all *possible* formulations of the fact in a single *synset* or *cluster*. Thus, if an extraction matches any of the formulations of a *synset*, is it said to match that *cluster*. They manually annotate 300 sentences of the original *OIE2016* dataset, resulting in 1354 *clusters*.

These benchmarks yield different conclusions regarding the best performing extractors. This is because their annotation principles, text corpora and matching functions are all different. Recently, Pei et al. (2023) made recommendations for assisting in deciding the best (neural) extractor for a given downstream task, and which benchmarks’ characteristics better correlate with it. While we agree that ultimately, we should test extractors on specific tasks (which we also revisit in Section 6.1), there is a need for sound references that will help appreciate limitations of current extraction technology, hopefully leading to better extractors. In their study, Pei et al. (2023) rejected *BenchIE* in comparing benchmarks,² while our inspections make us believe it is the most well-thought benchmark.³

3 Issues with *BenchIE*

By inspecting a (random) sample of 50 sentences of *BenchIE* - hereafter referred to as *B50* - we noticed a number of issues we think worth being taken care of for more meaningful comparisons of extractors.

3.1 Annotation problems

We identified five error types that are illustrated in Table 1.

Missing fact A fact is missing from the reference. This may be due to the lack of inference in *BenchIE* (see section 4.1), or from oversight or omission. The example shows the main information piece

²On the grounding that it is a fact-centered benchmark, which we see as a positive characteristic.

³Our review of existing benchmarks is not exhaustive. In particular, *ReOIE2016* (Zhan and Zhao, 2019) and *LSOIE* (Solawetz and Larson, 2021) have been proposed as updated versions of *OIE2016*, but they both use the same matching function, which has been largely criticized.

Error type	Annotation
<i>For example , when two such hydrophobic particles come very close , the clathrate-like baskets surrounding them merge</i>	
Missing Fact	when two such hydrophobic particles – come very close – [the] clathrate-like baskets surrounding them merge
False Fact	two [such] hydrophobic particles – come – [very] close
<i>He and his friends were said to have made bombs for fun on the outskirts of Murray , Utah .</i>	
Irrelevant fact	his friends – were said to have made bombs on – [the] outskirts [of Murray , Utah]
<i>They held the first Triangle workshop for thirty painters from the US , the UK and Canada at Pine Plains , New York .</i>	
Double annotation	They – held [the] [first] Triangle workshop at – [Pine Plains] New York
	They – held [the] [first] Triangle workshop at – New York
Double meaning	They – held [the] [first] Triangle workshop – for [thirty] painters from [the] US
	They – held [the] [first] Triangle workshop – for [thirty] painters from Canada

Table 1: Illustration of error types found in a sample of *BenchIE*. Square brackets indicate words that are optional when matching facts.

from the sentences and no formulation of that fact is present in the annotated clusters.

False fact A fact in the reference is false, or not necessarily implied by the sentence. The example shows a fact that is false because it lacks the mention that *hydrophobic particles can come very close* (when *baskets surrounding them merge*), instead conveying the fact that they *do*.

Irrelevant fact A fact in the reference lacks relevance, because of missing context or other information. In the example, without the optional *Murray, Utah*, the fact is not relevant because *bombs* being made on *outskirts* is not relevant information.

Double annotation Two clusters have at least one formulation that expresses the same information. The example shows two clusters being exactly the same without the optional words. The optional words in the first cluster should be *New York* and not *Pine Plains*.

Double meaning A single cluster has at least two formulations that express different information. The example shows two formulations of the same cluster conveying very distinct information.

Table 2 shows the number of sentences in *B50* for which a cluster or a formulation shows a given error type. We observe that more than half the sentences have one or more missing facts. In many cases, this is because of inference (see section 4.1), but some facts have simply been omitted. Both irrelevant or false facts are present in about a third of the sentences. This high frequency of issues in *BenchIE* shows the need for a more thorough annotation of the original sentences. These new annotations need to be motivated and conducted by solid guidelines that are typically lacking in the OIE task.

3.2 Matching problem

Alongside the errors found in *BenchIE*’s annotations, we annotated facts output by *ReVerb*, *IMojIE*, *OpenIE6*, and *CompactIE* and found many cases where a system made an extraction that in our eye was valid and should match an annotated fact that was not matched by the exact match used in the *BenchIE* benchmark. This is because while [Gash-teovski et al. \(2022\)](#) argue that they listed all possible valid formulations of a given fact in each cluster, we find that this is not the case (and we argue that this is in practice very hard to do).

In the 50 sentences of *B50*, we found 26 (52%) sentences with at least one fact from one of the extractors that was not matched when it should have. Because of this, we develop a new matching function aiming to capture more matches between extracted and annotated tuples.

Error type	Count	%
Missing Fact	26	52
False fact	15	30
Irrelevant fact	32	64
Double annotation	28	56
Double meaning	15	30

Table 2: Count of sentences with a given error type in the 50 sentences of *B50* in *BenchIE*.

4 Guidelines

4.1 Annotation Guidelines

Few principles are universally accepted by OIE benchmarks and systems authors. Here, we try to list crucial principles that make most sense for OIE output to be useful for downstream tasks, and aim for those to guide annotation of our and future

Principle	Sentence	Annotations
Informativity	<i>Alex lived in Paris and is now living in Cologne.</i>	✓ Alex – is now living in – Cologne ✗ Alex – is now – living
Minimality	<i>The group was created in 2020 by three people.</i>	✓ The group – was – created ✓ The group – was created in – 2020 ✓ The group – was created by – three people ✗ The group – was created – in 2020 by three people
Exhaustivity	<i>He has written several newspaper and magazine opinion pieces.</i>	✓ He – has written – [several] pieces ✓ He – has written – [several] opinion pieces ✓ He – has written – [several] newspaper pieces ✓ He – has written – [several] magazine pieces
Relation completeness	<i>Tokyo’s population is over 13 millions.</i>	✓ Tokyo’s population – is over – 13 millions ✗ Tokyo’s population – is – over 13 millions
Inference	<i>‘My classical way’ was released in 2010 on Marc’s own label, Frazzy Frog Music.</i>	✓ Marc’s [own] label – is – Frazzy Frog Music

Table 3: Illustration of annotation principles. Facts in **green** (preceded by a check mark) should be included in the annotation, while facts in **red** (preceded by an cross mark) should not.

references. Our full annotation guidelines, along with examples illustrating the principles can be found in Appendix J. Examples for the following principles are shown in Table 3.

Informativity Tuples should be informative and not contain generalities. This general principle, also present in *BenchIE*, means that as long as a tuple contains an information expressed in the sentence and that it respects all other principles, it should be included in the annotations. In the negative example, the second argument *living* is not informative, thus that fact should not be annotated.

Minimality Each tuple should be minimal, meaning that it can not be separated into multiple distinct tuples. The faulty annotation in line 2 of Table 3 combines three facts, and should therefore not be annotated. No mention of minimality is made in *BenchIE*, and their guidelines only suggest making non-necessary words optional, which in practice allows for overly long and imprecise annotations.

Exhaustivity The set of tuples for a given sentence should cover all pieces of information expressed in the sentence. At the cluster level, all of the possible formulations for which any of the arguments or the relation is different should be listed. In the example, all the types of *writing* that *He* has done should be annotated in separate clusters. *BenchIE*’s guidelines try to handle exhaustivity by listing all verb-mediated facts, which does not capture all information.

Relation Completeness Relations are responsible for the information; as such they should be complete, meaning that the information in the argu-

ments do not change the core meaning of the relation. In the example, the negative fact should not be annotated because its relation, *is*, is not complete, its meaning is modified by the word *over* in the second argument. In contrast, *BenchIE*’s guideline specifically encourages annotators to place words that are not verbs in or out of relations without regard for relation integrity.

Inference We define inferred tuples as facts that are implied by the sentence (true if the sentence is true), but for which the relation linking the arguments is not present in the text. Inference should be carried out because it is useful in downstream tasks such as QA or knowledge base (KB) population (Gashteovski et al., 2020). However, limits should be set in regards to the information inferred. These nuances are explained in detail in the complete guidelines but simply put, tuples that can be inferred without needing complex reasoning or external knowledge should be annotated. In the last example of Table 3, the tuple is included because it is clearly implied by the sentence. Since *BenchIE* only includes facts mediated by words present in the sentence, they annotate almost no inferred facts.

4.2 Matching Guidelines

To guide which extracted tuples should match with which annotated clusters, we need matching guidelines. Exactly identical extraction and annotation should obviously match and these cases are the only matches scored in *BenchIE*. However, we believe other nuances exist and the following principles aim to illustrate these nuances. Examples are shown in Table 4. Our full matching guidelines can be found in Appendix K.

Principle	Sentence	Annotations
Relation specificity	<i>The party was thrown out of the government.</i>	The party – was – thrown out of [the] government
		The party – was thrown out of – [the] government
		✗ The party – was thrown – out of [the] government
Word choice	<i>He is the older brother of Alex.</i>	He – is – [an] older brother
		✓ He – is – the older brother
Level of detail	<i>Alex broadcasts a web series Music on a website.</i>	Alex – broadcasts, [a] web series
		Alex – broadcasts, Music
		Alex – broadcasts Music on, a website
		✓ Alex – broadcasts – Music on a website
		✗ Alex – broadcasts – a web series Music on a website

Table 4: Examples of matching principles. Annotated facts are in black, facts in **green** (preceded by a check mark) should be matched to the annotations, while facts in **red** (preceded by a cross mark) should not.

Relation specificity The extracted and annotated relations should be as specific, meaning that prepositions and linking words should not be arbitrarily placed in the arguments or in the relation. In the example, the relation *was thrown* is not as precise and complete as the annotated relations and its meaning is changed by the words *out of* in the second argument, thus the extraction should not be matched.

Word choice Certain extractions may contain syntax errors, misplaced words or other word choices that are different to those of the annotations. These are not inherently bad but if they affect the sense of the extraction, then the extraction should not match. The example shows the case where even if the extraction and the annotation do not use the same determiner (*an* and *the*), they both convey the same meaning, i.e. that *He* is older than his brother.

Level of detail Many system extractions carry information from more than one annotated tuple. We want to match these extractions only if the extraction combines information from no more than two clusters, otherwise, we consider it too noisy for a downstream task. The example shows extractions that combine information from two and from three clusters respectively. The negative example is noisier since its third argument is too long and lacks preciseness.

	<i>BenchIE^{FL}</i>	<i>BenchIE</i>
Total clusters	1798	1354
Avg cluster/sentence	6	4.5
Avg formulation/cluster	3	6
Avg formulation length	10.6	12.5
Avg relation length	3.9	4.0

Table 5: Annotation statistics of *BenchIE* and our re-annotated version: *BenchIE^{FL}*.

5 *BenchIE^{FL}*

Following our new annotation guidelines, the first author of this paper (NLP scientist with 2 years of expertise) annotated *BenchIE*’s original sentences resulting in *BenchIE^{FL}*, a re-annotated corpus of 300 sentences. This annotation effort was conducted using the *AnnIE* annotation platform (Friedrich et al., 2022).

Different statistics regarding both annotation sets are reported in Table 5. First, we annotate more facts in total, that is, a higher average number of clusters per sentence. This is because we include inferred information and follow the minimality principle (meaning that we divide the information as much as possible). Second, we annotate (far) fewer formulations per cluster, both because relationship specificity is of great importance in our guidelines, and because we don’t rely on an exact match function (see Section 5.2), so there is no need for *BenchIE^{FL}* to list all possible formulations of the same fact. Third, we note that on average, our annotations are shorter; again due to the minimality principle. On the other hand, the mean lengths of the relations are more or less equivalent, due to our desire to preserve the specificity of the relations. Appendix I shows an example highlighting differences in both annotations of the same sentence.

5.1 Manual evaluation

Since all of our annotations are carried out by a single annotator, we set out to validate them using two other annotators (NLP scientists with over 20 years of expertise, *a₂* not being a co-author). We present them with sentences of *B50* and their annotations from both sets of annotations (25 each)⁴,

⁴See Appendix C for details of slight modifications we made to our annotations for this experiment.

and then ask them to (blindly) annotate : **exhaustivity** by indicating if a set of annotations for a given sentence fail to covers all the facts expressed, **minimality** by indicating clusters that can be separated into several (smaller) ones, and **relation completeness** by marking relations that are modified by their arguments or that do not hold on their own.

	<i>BenchIE</i> ^{FL}		BenchIE		κ
	a_1	a_2	a_1	a_2	
Exhaustivity	5	9	7	14	0.62
Minimality	3	16	5	9	0.68
Incomplete relations	4	14	9	18	0.68

Table 6: Validation statistics. Non-minimal sentences are those with at least one non-minimal annotation, and incomplete sentences are sentences with at least one cluster whose relation is incomplete. a_1 and a_2 represent the first and second annotators respectively.

Results are presented in Table 6 where each cell represents the count of error-full sentences for each criteria by each annotator for both sets of annotations. We observe that for all criterion, the annotators found our annotations to follow the guidelines much more closely. Both annotators find more than double the errors in *BenchIE* compared to our annotation for almost all error types. Table 6 also shows Cohen’s kappa scores for the agreement between the two annotators, which indicate a moderate agreement. We found validating annotations requires annotators to fully embrace guidelines, a too demanding task. We believe that OIE benchmarks are better evaluated by how strongly their results resemble performance of systems on downstream tasks, as explored in section 6.1. In Appendix G, we also devise (objective) tests that highlight annotation issues and show that *BenchIE*’s annotations contain a lot more potential errors than ours.

5.2 Matching function

As a reminder, previous OIE benchmarks (*WiRE57*, *CaRB* and others) used **token-level** matching, where a candidate extraction would be scored in precision and recall by the number of words overlapping with the best-matching annotation. *BenchIE*’s authors introduced the **exact** matching, only possible thanks to *synsets*, where a candidate extraction only matches an annotation if they are identical. We set out to develop a new matching function that captures more matches than the exact matching used in *BenchIE*. In order to evaluate this matching function, we annotate the extractions

produced by the seven systems introduced in Section 2.1 for the 50 sentences of *B50*, and indicate for each extraction the index of the cluster it should match to according to our matching guidelines; an extraction that should not match any cluster is marked by 0. Examples of such annotations are provided in Table 7. The resulting resource, named *BenchIE*_{match}^{FL} contains 9400 extraction-annotation pairs, 96.8% of which have no associated cluster.

We then gathered three simple types of heuristics that we found capture more of those annotated matches that are not exact matches :

Alternative formulations (AF) In two specific (yet frequent) situations, we do credit an extraction which does not match a reference cluster because it regroups information from two clusters: when its relation is reducible⁵ to is, and when one of its argument contains a coordinate conjunction and. Implementation details are provided in Appendix B, but as an illustration, the tuple (He – is – Canadian and a musician) might yield to two alternatives (He – is – Canadian) and (He – is – a musician), that will more likely (exact) match the reference.

Level of detail Matching (LoD) We match an extraction which linearization⁶ is verbatim the one of a reference formulation and which one of its argument and its relation are also present in another cluster. A typical example is illustrated in Table 4 where the candidate tuple (Alex – broadcasts – Music on a website) is matched because it has the same linearization as the reference tuple (Alex – broadcasts Music on – a website) and *Alex*, and *broadcast* are found in the second cluster: (Alex – broadcast – Music) in the corresponding slots.

Punctuation (Punc) We carry out matching removing all punctuation characters and lowercasing strings. This is because we have not listed all possible combinations of capitalization and punctuation in our annotations, and consider we should not.

Both AF and LoD heuristics help identify matches that occur because of the *Level of Detail* principle from our guidelines, while the Punc heuristic applies to the *Word Choice* principle. The F1 score of matching function variants on *BenchIE*_{match}^{FL} is reported in Table 8. We observe that the exact match function is outperformed by

⁵We say r reduces to r' , if removing optional words from r leads to r' .

⁶By linearization of a tuple, we mean the string obtained by concatenating in that order arg1, relation and arg2.

System extraction	Annotation cluster	Match
He – served as first Prime Minister of – Australia	He – served as – [the][a] Prime Minister	0
He – became – founding justice of High Court of Australia	He – became – [a] founding justice	1

Table 7: Examples of matching annotations. Index of clusters are not relevant here.

each addition of the heuristics we described, and that using them all leads to the best performance. Therefore we selected (EM + AF + LoD + Punc) as the default matching function in the evaluation toolkit accompanying *BenchIE^{FL}*. We further verify in Appendix D, that the ranking of systems thanks to our scoring function correlates better than other scoring functions to the ranking obtained by evaluating systems based on the (manual) matching annotations we conducted.⁷

Matching function	F1
Exact Match (EM)	0.83
EM + AF	0.87
EM + LoD	0.87
EM + AF + LoD	0.96
EM + AF + LoD + Punc	0.97

Table 8: Matching F1 of matching function variants on *BenchIE^{FL}_{match}*.

6 Experiments

We use the default configuration for the 7 extractors mentioned in Section 2.1 without attempting to optimize them for specific benchmarks. Pei et al. (2023) observed that optimizing such systems do not lead to significant performance differences.

6.1 Downstream tasks

Ultimately, OIE extractions, and thus systems, are only useful in downstream applications. This means that a useful benchmark’s rankings should follow closely the rankings obtained by evaluating OIE extractions usefulness in these downstream tasks. To test whether this is the case for our reference, we study three tasks that directly use OIE triples: Assertion-Based Question Answering (Yan et al., 2018), Complex Questions Answering with quasi Knowledge Graphs (Lu et al., 2019) and Knowledge Base Population (Mesquita et al., 2019). See Figure 1 for a visual description of the

⁷We also attempted training a matching function on *BenchIE^{FL}_{match}*, but found disappointing result with low generalization on unseen matching pairs.

tasks, Table 9⁸ for systems scores and Annex F for examples and details on each task’s setup.

System	ABQA	C-QA	KBP
<i>ReVerb</i>	0.230	0.092	0.149
<i>ClausIE</i>	0.180	0.089	0.026
<i>MinIE</i>	0.270	0.095	0.396
<i>IMojIE</i>	0.170	-	-
<i>OpenIE6</i>	0.170	0.087	0.064
<i>M2OIE</i>	0.170	0.090	0.014
<i>CompactIE</i>	0.160	0.093	0.006

Table 9: Scores of systems on downstream tasks.

ABQA In *ABQA*, the input is a passage of a few sentences along with a question, and the output is the answer to that question identified in the passage. Yan et al. (2018) created an *ABQA* reference, *WebAssertions* by using *ClausIE* to extract all facts from the passages and asking annotators to identify extractions answering the given question. We run all tested systems on the first 100 passages, using our matching function to match extractions of systems to answers, and compute scores for the task by giving systems one point for a single sentence if one of their extractions match any answer cluster.

C-QA Lu et al. (2019) introduce a novel approach to answering complex questions called *QUEST*. This system uses OIE extractions from web documents to construct a quasi Knowledge Graph and uses that graph to answer questions. They evaluate this system on a few QA benchmarks, including *WikiAnswers* (CQ-W) (Abujabal et al., 2017). We use our tested systems to extract facts from the web documents on the first 50 questions of *WikiAnswers*. We then use *QUEST* to construct answers and report the scores measured by the Mean Reciprocal Rank.

KBP *KnowledgeNet* is a dataset of more than 7000 annotated sentences introduced by Mesquita et al. (2019) for evaluating the task of automatically populating a Knowledge Base. It’s annotations contain triples similar to OIE ones, with a fixed subset

⁸*IMojIE* is excluded from C-QA and KBP because of compute constraints.




Figure 1: Downstream tasks flowchart.

of 15 relations, akin to the traditional Information Extraction task. Here we measure OIE system’s ability by running them on the input sentence and counting how many triples from *KnowledgeNet* they were able to extract correctly.

We compute Pearson product-moment correlation coefficients between these rankings and the ones obtained using OIE benchmarks and observe (see Table 10) that $BenchIE^{FL}$ has the highest correlation on these varied tasks, leading us to hypothesize that it is a better indicator of real performances of extractors. We also observe strong trends in the rankings (the performance on $BenchIE^{FL}$ always being the best indicator, followed by those measured on $BenchIE$, $WiRE57$ and finally $CaRB$), further reinforcing our claim that these results should hold for most tasks.

Benchmark	ABQA	C-QA	KBP
<i>WiRE57</i>	0.044	-0.184	0.204
<i>CaRB</i>	-0.649	-0.575	-0.382
<i>BenchIE</i>	0.616	0.305	0.631
$BenchIE^{FL}$	0.940	0.504	0.941

Table 10: Correlation between system scores on benchmarks and on downstream tasks.

6.2 Benchmark comparison

We can compare benchmarks by the different system rankings they lead to. Figure 2 shows scores of tested systems on the four aforementioned benchmarks, using their default evaluation toolkit.

We observe that all rankings differ largely. Certain similarities can be observed between the results obtained by *BenchIE* and our reference, both of which rank *MinIE* as the best-performing system. However, for virtually all other systems, there are (major) differences between their final rankings according to these two references. It is also remarkable that both recent state-of-the-art benchmarks, *BenchIE* and *CaRB* largely overestimate neural models compared to $BenchIE^{FL}$. In fact, we observe that neural networks by the virtue of the




Figure 2: System performance by benchmarks, scored using default scoring function of each benchmark.

datasets they have been trained on, have a tendency of copying large chunks of input texts,⁹ that often lead to non informative tuples. See Appendix E for an analysis of the shortcomings of neural systems.

7 Discussion

We propose new annotated resources: most notably a re-annotated OIE corpora, $BenchIE^{FL}$ and a set of matching annotations, $BenchIE_{match}^{FL}$. We also propose new guidelines for the OIE task, both for the annotation of tuples and the matching of candidate extractions. We also deliver noticeable improvements on the exact matching function, while not compromising the principles behind *BenchIE*.

Thanks to those resources, we conduct a meaningful comparison of off-the-shelf extractors, showing that older rule-based systems are still competitive. Finally we conduct a study that shows on three downstream tasks that our benchmark better reflects the performance of OIE systems on those tasks, making $BenchIE^{FL}$ the best reference to guide OIE system choice, and to influence system development.

There are a number of avenues worth pursuing along this work, especially in annotating more sen-

⁹*ImojIE* and *M2OIE* produce tuples with an average number of words of 13.4 and 12.2 respectively, while for instance, *ReVerb* produces tuples of 7.5 words on average.

tences, both to serve as training sets for new systems and in other languages.

Limitations

BenchIE^{FL} aims to correct some of the shortcomings of *BenchIE*. However, it still has some limitations.

- While it contains a fair number of sentences to draw conclusions about which OIE systems to use, it lacks the needed size to be useful in training models. Annotation of sentences requires a lot of effort, but adding to our 300 sentences would likely be useful.
- We believe our annotations to be more rigorous than previous ones. Still, it is likely that some mistakes were made and should be corrected. It would have been great (but out of the scope of this work) to enrol more annotators to better measure their agreement while annotating according to our guidelines.
- Even if the matching function we propose performs better than the exact match function of *BenchIE*, it still lacks some flexibility. We were not able to train a better function, but some different features or methods may outperform our custom function.
- Still, *BenchIE^{FL}*, like most of OIE benchmarks, is only available in English.

Ethics Statement

References

- Abdalghani Abujabal, Mohamed Yahya, Mirek Riedewald, and Gerhard Weikum. 2017. Automated template generation for question answering over knowledge graphs. In *Proceedings of the 26th international conference on world wide web*, pages 1191–1200.
- Sangnie Bhardwaj, Samarth Aggarwal, and Mausam Mausam. 2019. *CaRB: A crowdsourced benchmark for open IE*. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6262–6267, Hong Kong, China. Association for Computational Linguistics.
- Luciano Del Corro and Rainer Gemulla. 2013. *Clausie: Clause-based open information extraction*. In *Proceedings of the 22nd International Conference on World Wide Web, WWW '13*, page 355–366, New York, NY, USA. Association for Computing Machinery.
- Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011. *Identifying relations for open information extraction*. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 1535–1545, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Anthony Fader, Luke Zettlemoyer, and Oren Etzioni. 2014. *Open question answering over curated and extracted knowledge bases*. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14*, page 1156–1165, New York, NY, USA. Association for Computing Machinery.
- Farima Fatahi Bayat, Nikita Bhutani, and H. Jagadish. 2022. *CompactIE: Compact facts in open information extraction*. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 900–910, Seattle, United States. Association for Computational Linguistics.
- Niklas Friedrich, Kiril Gashteovski, Mingying Yu, Bhushan Kotnis, Carolin Lawrence, Mathias Niepert, and Goran Glavaš. 2022. *AnnIE: An annotation platform for constructing complete open information extraction benchmark*. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 44–60, Dublin, Ireland. Association for Computational Linguistics.
- Kiril Gashteovski, Rainer Gemulla, and Luciano del Corro. 2017. *MinIE: Minimizing facts in open information extraction*. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2630–2640, Copenhagen, Denmark. Association for Computational Linguistics.
- 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, Online. Association for Computational Linguistics.
- 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 *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4472–4490, Dublin, Ireland. Association for Computational Linguistics.
- Luheng He, Mike Lewis, and Luke Zettlemoyer. 2015. *Question-answer driven semantic role labeling: Using natural language to annotate natural language*. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 643–653, Lisbon, Portugal. Association for Computational Linguistics.

- Keshav Kolluru, Vaibhav Adlakha, Samarth Aggarwal, Mausam, and Soumen Chakrabarti. 2020a. [OpenIE6: Iterative Grid Labeling and Coordination Analysis for Open Information Extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3748–3761, Online. Association for Computational Linguistics.
- Keshav Kolluru, Samarth Aggarwal, Vipul Rathore, Mausam, and Soumen Chakrabarti. 2020b. [IMO-JIE: Iterative memory-based joint open information extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5871–5886, Online. Association for Computational Linguistics.
- William Lechelle, Fabrizio Gotti, and Phillippe Langlais. 2019. [WiRe57 : A fine-grained benchmark for open information extraction](#). In *Proceedings of the 13th Linguistic Annotation Workshop*, pages 6–15, Florence, Italy. Association for Computational Linguistics.
- Xiaolu Lu, Soumajit Pramanik, Rishiraj Saha Roy, Abdalghani Abujabal, Yafang Wang, and Gerhard Weikum. 2019. [Answering complex questions by joining multi-document evidence with quasi knowledge graphs](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '19*. ACM.
- Filipe Mesquita, Matteo Cannavicchio, Jordan Schmidek, Paramita Mirza, and Denilson Barbosa. 2019. [KnowledgeNet: A benchmark dataset for knowledge base population](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 749–758, Hong Kong, China. Association for Computational Linguistics.
- Kevin Pei, Ishan Jindal, Kevin Chen-Chuan Chang, ChengXiang Zhai, and Yunyao Li. 2023. [When to use what: An in-depth comparative empirical analysis of OpenIE systems for downstream applications](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 929–949, Toronto, Canada. Association for Computational Linguistics.
- Youngbin Ro, Yukyung Lee, and Pilsung Kang. 2020. [Multi~2OIE: 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.
- Jacob Solawetz and Stefan Larson. 2021. [LSOIE: A large-scale dataset for supervised open information extraction](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2595–2600, Online. Association for Computational Linguistics.
- Gabriel Stanovsky and Ido Dagan. 2016. [Creating a large benchmark for open information extraction](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2300–2305, Austin, Texas. Association for Computational Linguistics.
- Gabriel Stanovsky, Ido Dagan, and Mausam. 2015. [Open IE as an intermediate structure for semantic tasks](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 303–308, Beijing, China. Association for Computational Linguistics.
- Zhao Yan, Duyu Tang, Nan Duan, Shujie Liu, Wendi Wang, Daxin Jiang, Ming Zhou, and Zhoujun Li. 2018. [Assertion-based qa with question-aware open information extraction](#).
- Alexander Yates, Michele Banko, Matthew Broadhead, Michael Cafarella, Oren Etzioni, and Stephen Soderland. 2007. [TextRunner: Open information extraction on the web](#). In *Proceedings of Human Language Technologies: The Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, pages 25–26, Rochester, New York, USA. Association for Computational Linguistics.
- Junlang Zhan and Hai Zhao. 2019. [Span model for open information extraction on accurate corpus](#).

A Implementation details

We used implementations of various systems and benchmarks in order to test and compare them to each other. Here is a brief description of the systems we used:

ReVerb This system uses regular expressions to extract candidate facts and a simple feature-based classifier to filter duplicates and non-useful tuples.¹⁰

ClausIE This system uses grammatical analysis to create *clauses*, or minimal sentences and extracts tuples based on those clauses and their *type*. We could not find a working implementation of *ClausIE* so we used a simple script to generate extraction from a web page providing an API to demonstrate the system’s capabilities.¹¹

MinIE This system is built on top of *ClausIE*, adding more patterns to identify clauses and modifying final extractions to be minimal and have more solid relations. We used a Python wrapper of the official implementation.¹²

IMojIE This system is a sequence to sequence model using a BERT encoder, trained on extractions from *OpenIE4*, *RnnOIE* and *ClausIE*.¹³

OpenIE6 This system is a successor to *OpenIE*, and it is a grid-labelling model built on BERT embeddings with syntactic and grammatical constraints used at inference time trained on boot-strapped *IMojIE* extractions.¹⁴

M2OIE This system is one of the only multilingual system, only needing corresponding language BERT models, that works by first extracting relations and then extracting related arguments in a sequence-to-sequence architecture.¹⁵

CompactIE This system works by extracting constituents (arguments and relations) and linking them using a neural classifier, and is trained to output *compact*, or minimal extractions by adding constraints on the extraction components.¹⁶

All systems were ran on an Apple M2 processor. Most systems had dependencies issues out of the box and significant effort was needed to make them work as expected. Since our datasets were quite small, *CaRB* being the largest with 1200 sentences, computation time was not an issue, although all neural systems needed far more time to extract facts than rule-based systems. *IMojIE* was the slowest system, needing more than four hours to run on *CaRB*’s test set.

Regarding the benchmarks we have been using:

WiRe57 : This benchmark is comprised of a small, expertly annotated corpus of 57 sentences and uses a token-level matching function.¹⁷

CaRB : This benchmark which is a crowd-sourced annotation of 1200 sentences also uses a token-level scoring function. We used the test set which is made up of a total of 640 sentences.¹⁸

BenchIE : This benchmark uses *clusters* or *synsets* to group all formulations of a single fact, allowing usage of an exact matching function. It is made up of 300 manually annotated sentences. We used the default facet.¹⁹

B Matching function

B.1 Alternative formulations

Here we explain the details of implementation of the alternative formulations we introduce in section 5.2. The aim of these alternative formulations is to match extractions to annotations even when an exact match does not show correspondence. We introduce the notion of rewriting pairs (A,B) where we authorize an argument of an extraction that contains both A and B to be rewritten by removing either A or B. We called these modified extractions *alternative formulations*

We identify two cases where alternative formulations can be generated safely: in extractions with the relation *is*, and in those where an argument contains the token *and*. In the first case, we collect from the reference all pairs (A,B) from formulations (A,r,B) where *r* reduces to *is*,²⁰ while in the second case, we collect pairs (A,B) whenever we have two tuples (E,rel,A) and (E,rel,B) in the reference. Table 11 shows an example of how these

¹⁰<http://reverb.cs.washington.edu/>

¹¹<https://gate.d5.mpi-inf.mpg.de/ClausIEGate/ClausIEGate/>

¹²<https://github.com/mmxgn/miniepy>

¹³<https://github.com/dair-iitd/imojie>

¹⁴<https://github.com/dair-iitd/openie6>

¹⁵<https://github.com/youngbin-ro/Multi2OIE>

¹⁶<https://github.com/FarimaFatahi/CompactIE>

¹⁷<https://github.com/rali-udem/WiRe57>

¹⁸<https://github.com/dair-iitd/CaRB>

¹⁹<https://github.com/gkiril/benchie>

²⁰We say *r* reduces to *r'*, if removing optional words from *r* leads to *r'*.

pairs are generated from a sample annotation and extraction and how they can match annotations.

B.2 Scoring extractions

If a candidate extraction does not exact match a reference cluster, we consider alternative extractions obtained by rewriting pairs and give credit to the original extraction if one of those alternatives exact match the reference. We only consider the first matching alternative and its associated cluster for the computation of precision and recall as to not overly reward extractions for combining information and thus not respecting the minimality principle.

<i>Chilly Gonzales is a Canadian musician who lived in Paris, France and in Cologne, Germany.</i>
Annotations: (Chilly Gonzales – lived in – Paris) (Chilly Gonzales – lived in – Cologne) (Chilly Gonzales – is – Canadian) (Chilly Gonzales – is [a] – musician)
Rewriting pairs: is : (Canadian – Musician) and : (Paris – Cologne)
Extraction: (Chilly Gonzales – is a – Canadian musician)
Alternative formulations: (Chilly Gonzales – is – Canadian) (Chilly Gonzales – is – musician)
Extraction: (Chilly Gonzales – lived in – Paris and Cologne)
Alternative formulations: (Chilly Gonzales – lived in – Paris) (Chilly Gonzales – lived in – Cologne)

Table 11: Alternative formulation generation example

C About inferred clusters

Since our annotation and *BenchIE*’s differ largely in handling of inference, we present a modified version of our annotation set for which the original annotator transformed the inferred clusters to use only words present in the original sentence (light inference) and removed the inferred clusters that could not be modified as such (heavy inference).

This modified annotation is what is presented to the annotator when comparing ours and *BenchIE*’s annotations in section 5.1, in order to give a fair comparison in regards to exhaustiveness given that *BenchIE*’s annotation does not contain inference, other than in some very rare cases. Statistics regarding this modification and the use of inference in both references are presented in Table 12. We observe that in our original annotation, we had 22 annotated clusters that were instances of heavy inference which we were not able to transform into light inference and had to be removed, while 22 of them could be transformed. We also observe that *BenchIE* has very few instances of inference. This is because they decide to only annotate facts for which the relation is verbatim in the text, although we did find a few instances of light inference.

	<i>BenchIE</i> ^{FL}	Modified	<i>BenchIE</i>
Total clusters	139	117	119
Inferred clusters	47 (34%)	25 (21%)	3 (2.5%)
Heavy inference	44 (32%)	0	0
Light inference	3 (2%)	25 (21%)	3 (2.5%)

Table 12: Inference statistics

D Comparison of scoring functions

Using our manual matching annotations, we compute theoretical scores and rankings for tested systems. We then compute scores and rankings using the different matching scores and functions of previous benchmarks. We compare these matching methods by computing their Pearson product-moment correlation coefficients with the manual rankings on *BenchIE*_{match}^{FL} (See Table 13). We observe that token-level matching has very low correlation and that while both *BenchIE*’s exact match and our *Custom Match* have similar scores, ours correlates to a greater degree.

Matching function	Correlation with manual rankings
<i>WiRE57</i>	0.219
<i>BenchIE</i>	0.961
<i>Custom Match</i>	0.997

Table 13: Pearson product-moment correlation between rankings obtained by matching functions and manual rankings on *BenchIE*_{match}^{FL}

passage:	the internal frame backpack is a recent innovation , invented in 1967 by greg lowe, who went on to found lowe alpine and lowepro, companies specializing in backpacks and other forms of carrying bags for various equipment
question:	when was the backpack invented ?
<i>WebAssertions</i>	(a recent innovation – invented – in 1967 – by greg lowe)
<i>WebAssertions</i> ^{FL}	([the] [internal frame] backpack – [was]/[be]/[is] invented in – 1967)

Table 14: Question-Passage-Answer triples from *WebAssertions* and *WebAssertions*^{FL}.

E Neural system scores

Here we try to demonstrate why neural approaches get lower scores on our benchmark than on others. First we compare lengths of extractions by systems in Table 15, where we see that neural approaches all have longer mean extraction length compared to rule-based systems, except for *CompactIE* which has been specifically trained to output compact extractions. We also see that our two best performing systems, *MinIE* and *ReVerb* have the shortest extractions. We hypothesize that they obtain the highest scores in part because they are most capable at precisely separating facts into minimal extractions, which is something that neural systems fail to do. This is illustrated in Table 20 where we show all extractions for all systems on a given sentence. We see that both *IMojIE* and *M2OIE*, neural systems, make extremely long and useless extractions, recopying almost the whole sentence. Indeed, we observed that neural systems have a tendency to recopy large parts of the input sentences without being able to accurately separate facts. We hypothesize that this is because they have been trained on non-minimal and relation-complete extractions.

System	Mean extraction length (# words)
ReVerb	7.5
ClausIE	10.9
MinIE	7.9
IMojIE	13.4
OpenIE6	12.2
M2OIE	12.2
CompactIE	9.3

Table 15: Extraction length by system

F Downstream Tasks Details

Here we provide more details on how we ran the experiment on the downstream tasks and on some modifications that were made to either their data or their scoring procedures.

ABQA Table 14 shows an example of a triple from *WebAssertions*. We found many tuples in it that are too long or even fail to answer the question. Thus, we re-annotated the 100 first passage-question pairs of *WebAssertions*, following the *BenchIE* format and listing all clusters that can answer the question. We distribute the resulting resource named *WebAssertions*^{FL}, an example of which being in Table 14. We believe that this did not impact the results of system on the task as we simply corrected the mistakes and formatted the answers in *BenchIE*’s cluster format. We also only used the first 100 sentences of *WebAssertions* because of the high annotation effort required.

C-QA For this task, our results are far below what the original paper’s scores were. This is because we slightly modified the scoring used : *QUEST*’s original code allowed for various formulations of an entity to be regrouped in a single answer. However, in practice, multiple different entities were regrouped in a single answer, giving full points to answers containing multiple different elements. Using the original scoring, we found system’s score to be closer to the original papers (0.25 to 0.29). We thus separated all entities into distinct answers and computed MRR on those answers, which we believe to be a more accurate evaluation method. Table 16 shows an example of a single question from *WikiAnswers* along with the answer modification procedure.

IMojIE is not included in this experiment because of compute limitations. Indeed, *QUEST* uses 10 documents per questions, each document containing a few hundred sentences, making the total sentence count for 50 questions more than 50,000. *IMojIE* being the slowest system, running it on this whole corpus would have taken more than two months of compute time on our setup.

KBP Table 17 shows an example of annotation from *KnowledgeNet*. This reference’s annotated relations come from a fixed subset of 15 relations,

Question	which films were starred by julia roberts and richard gere?
Gold Answer	runaway bride, pretty woman
Documents	Document 1 : In his career Richard Gere has starred opposite some seriously attractive women and ... With Kim Basinger he starred in 2 movies "No Mercy" and "Final Analysis", with Diane Lane... ... Document 2 : Julia Fiona Roberts is an American actress and producer who made her debut in ... Roberts made her breakthrough the following year by starring in
Original QUEST Answer	Answer 1 : ('a year', 'golden globe', 'pretty woman', 'jason alexander') Answer 2 : ('phillip', '41', 'richard') ... MRR : 1
Reformatted QUEST Answer	Answer 1 : 'a year' Answer 2 : 'golden globe' Answer 3 : 'pretty woman' ... MRR : 0.333

Table 16: *WikiAnswers* Example with output modification.

like *PlaceOfBirth* or *FoundedBy*. Consequently, they do not annotate the relation words from the sentence, making it harder to compare to our system’s annotation. In order to find matches between extractions and annotations, we only compared arguments, and assumed that if both arguments were absolutely equal to annotations, the information should be very similar. This might not be perfectly accurate but does not penalise or reward systems differently and accurately measures systems ability to properly extract arguments. Furthermore, we only counted one annotation for each different relation per sentence, since the annotations in *KnowledgeNet* often contained more than one annotation with the same relation for the same fact, often listing pronouns and coreference to a single entity in different annotations.

Here, *IMojIE* is once again excluded from the experiment, because of the same compute limitations. *KnowledgeNet*’s corpus contains more than 7000 sentences, and running this system would have necessitated more than 10 days of compute time on our setup.

Sentence : <i>After moving to New York, where she established Euro Capital Properties along with her husband Jacques, her discerning eye for design was put to great use.</i>
Annotations : she - <i>PlaceOfResidence</i> - New York her - <i>PlaceOfResidence</i> - New York she - <i>Spouse</i> - Jacques her - <i>Spouse</i> - Jacques

Table 17: *KnowledgeNet* Annotation Example.

G Proxys

Here, we validate our re-annotation by devising objective proxies that highlight differences between *BenchIE*’s annotations and ours. These proxies try to identify annotations from both sets that might fall into one of the error categories from our experiment in Section 3.1.

Double Annotation We find distinct clusters that have identical formulations.

Double Meaning We find clusters that have two different meaning by selecting those that have identical first argument and relation but with different second argument.

False/Irrelevant/Missing Facts We find clusters from one set of annotation that don’t appear in the other set.

Table 18 shows the number of sentence that have at least one cluster identified by each proxy. We first observe that all proxies return a lot more occurrences from *BenchIE*’s annotation than from ours. This leads us to believe that this annotation contains more errors. However, not all annotations matched by the proxies are necessarily an error. Examples of such annotations from both sets are presented in Table 19.

For the *Double Annotation* proxy, we find that annotations from *BenchIE* are actually errors, the same fact with a differently formulated relation present in two distinct clusters. In contrast, the annotations from *BenchIE*^{FL} are cases where we find both the annotation regrouping the information

in a single cluster or separating it in two would be adequate. We verify that this is the case for the 5 *Double annotation* clusters from *BenchIE^{FL}*, and that most observed examples from the 135 cases from *BenchIE* are all double annotation errors.

The same is true for the *Double Meaning* proxy, where most cases from *BenchIE* are errors, and almost none from *BenchIE^{FL}* are. *BenchIE*'s example in Table 19 is especially strong, where five different meanings are present in a single cluster. However, we estimate that around 50% of annotations matched by that proxy from *BenchIE* are errors, still a very high count, and that most from *BenchIE^{FL}* are simply reformulation of the same entity or fact, like *Doctor* and *Dr.* from the example.

For the *False/Irrelevant/Missing Facts* proxy, it is important to note that counts are inverted : the proxy returns 17 annotations from *BenchIE* that are not present in *BenchIE^{FL}*, and 68 from the inverse, leading us to believe that our annotation is more exhaustive. Looking at the example, we once again find that the fact present in *BenchIE* and not in *BenchIE^{FL}* is actually an irrelevant fact, and that the inverse is a simple case of inference, not present in *BenchIE*. We estimate that 90% of facts present in our annotation and not in *BenchIE*'s are actually *Missing Facts*, and that more than 50% of facts present in *BenchIE* and not in our annotation are either *Irrelevant* or *False*.

	BenchIE	BenchIE^{FL}
Double annotation	135	5
Double meaning	55	26
False/Irrelevant/Missing fact	17	68

Table 18: Proxies sentence count

in *BenchIE*. We also see that *BenchIE*'s annotation contain a lot of different formulations of the same fact, most of which we do not consider to be valid because of relation integrity.

H Examples of Extractions

Table 20 shows example of extractions from all 7 tested systems on a single sentence.

I Annotation Examples

Table 21 shows differences between our annotation and *BenchIE*'s original annotations for the same sentences. We see that our annotation has a lot more clusters, partly because of our inclusion of inference but also because of the minimality principle that guides us to annotate the fact (My Classical Way – was – released), which is not annotated

Double Annotation
<p><i>BenchIE</i> : He served as the first Prime Minister of Australia and became a founding justice of the High Court of Australia .</p> <p>He - became - [a] [founding] justice of [the] High Court [of Australia]</p> <p>He - became [a] [founding] justice of - [the] High Court [of Australia]</p> <p>He - became [a] [founding] justice - of [the] High Court [of Australia]</p>
<p><i>BenchIE^{FL}</i> : For patients who do not recover quickly , the protocol also includes support groups and/or psychotherapy .</p> <p>[the] protocol - includes - support groups</p> <p>[the] protocol - includes - support groups and/or psychotherapy</p> <p>[the] protocol - includes - psychotherapy</p> <p>[the] protocol - includes - support groups and/or psychotherapy</p>
Double Meaning
<p><i>BenchIE</i> : It deals with cases of fraud in relation to direct taxes and indirect taxes , tax credits , drug smuggling , and money laundering , cases involving United Nations trade sanctions , conflict diamonds and CITES .</p> <p>It - deals - with cases of fraud in [relation to] money laundering</p> <p>It - deals - with cases of fraud in [relation to] drug smuggling</p> <p>It - deals - with cases of fraud in [relation to] indirect taxes</p> <p>It - deals - with cases of fraud in [relation to] direct taxes</p> <p>It - deals - with cases of fraud in [relation to] tax credits</p>
<p><i>BenchIE^{FL}</i> : Dr. Jagan himself was personally involved in the organization of the strike , and helped to raise funds across the country to it .</p> <p>Jagan - is [a] - Dr.</p> <p>Jagan - is [a] - Doctor</p>
False/Irrelevant/Missing Facts
<p><i>BenchIE</i> : Graner handcuffed him to the bars of a cell window and left him there , feet dangling off the floor , for nearly five hours .</p> <p>feet - dangling off [the] floor for - nearly five hours</p>
<p><i>BenchIE^{FL}</i> : He served as the first Prime Minister of Australia and became a founding justice of the High Court of Australia .</p> <p>Australia - has [had] - [a] High Court</p>

Table 19: Examples of proxy outputs

His son , John Crozier , Jr. , was an early aviation pioneer who began building a human-powered flying machine in the 1890s , but was killed in a feud in Grainger County before he could complete it .	
System	Extractions
ReVerb	(John Crozier – was – an early aviation pioneer) (an early aviation pioneer – began building – a human-powered flying machine) (the 1890s – was killed in – a feud) (he – could complete – it)
ClausIE	(His – has – son) (His son – is – John Crozier) (John Crozier – is – Jr.) (His son – was – an early aviation pioneer) (an early aviation pioneer – began – building a human-powered flying machine in the 1890s) (an early aviation pioneer – began – building a human-powered flying machine) (an early aviation pioneer – was killed – in a feud in Grainger County before he could complete it) (an early aviation pioneer – was killed – in a feud in Grainger County) (he – could complete – it)
MinIE	(His – has – son) (son – is – John Crozier) (John Crozier – is – son) (John Crozier – is – Jr.) (son – was – early aviation pioneer) (early aviation pioneer – began building human-powered flying machine in – the 1890s) (early aviation pioneer – began – building human-powered flying machine) (early aviation pioneer – was killed in feud in – Grainger County) (he – complete – it)
IMojIE	(His son – was – an early aviation pioneer who began building a human-powered flying machine in the 1890s , but was killed in a feud in Grainger County) (he – could complete – it)
OpenIE6	(His son – was – an early aviation pioneer who began building a human-powered flying machine in the 1890s) (an early aviation pioneer – began – building a human-powered flying machine in the 1890s) (an early aviation pioneer – was killed – in a feud in Grainger County before he could complete it) (he – could complete – it) (an early aviation pioneer – began building – a human-powered flying machine in the 1890s)
M2OIE	(His son , John Crozier , Jr. – was – an early aviation pioneer who began building a human-powered flying machine in the 1890s , but was killed in a feud in Grainger County) (an early aviation pioneer – began building – a human-powered flying machine in the 1890s) (an early aviation pioneer – was killed – in a feud in Grainger County) (he – could complete – it)
CompactIE	(John Crozier – was – an early aviation pioneer) (an early aviation – was killed – in a feud in Grainger County) (an early aviation – could complete – it)

Table 20: Extraction examples

“ My Classical Way ” was released on 21 September 2010 on Marc ’s own label , Frazzy Frog Music .	
BenchIE’s annotations	BenchIE ^{FL} ’s annotations
<p>Cluster 1:</p> <p>[‘ ’] My Classical Way [‘ ’] – was released on – 21 September 2010</p> <p>[‘ ’] My Classical Way [‘ ’] – was – released on 21 September 2010</p> <p>[‘ ’] My Classical Way [‘ ’] – was released – on 21 September 2010</p> <p>[“] My Classical Way [”] – was released on – 21 September 2010</p> <p>[“] My Classical Way [”] – was – released on 21 September 2010</p> <p>[“] My Classical Way [”] – was released – on 21 September 2010</p> <p>“My Classical Way” – was released on – 21 September 2010</p> <p>“My Classical Way” – was – released on 21 September 2010</p> <p>“My Classical Way” – was released – on 21 September 2010</p> <p>Cluster 2:</p> <p>[‘ ’] My Classical Way [‘ ’] – was released on – Marc ’s [own] label</p> <p>[‘ ’] My Classical Way [‘ ’] – was – released on Marc ’s [own] label</p> <p>[‘ ’] My Classical Way [‘ ’] – was released – on Marc ’s [own] label</p> <p>[‘ ’] My Classical Way [‘ ’] – was released on – Frazzy Frog Music</p> <p>[‘ ’] My Classical Way [‘ ’] – was – released on Frazzy Frog Music</p> <p>[‘ ’] My Classical Way [‘ ’] – was released – on Frazzy Frog Music</p> <p>[“] My Classical Way [”] – was released on – Marc ’s [own] label</p> <p>[“] My Classical Way [”] – was – released on Marc ’s [own] label</p> <p>[“] My Classical Way [”] – was released – on Marc ’s [own] label</p> <p>[“] My Classical Way [”] – was released on – Frazzy Frog Music</p> <p>[“] My Classical Way [”] – was – released on Frazzy Frog Music</p> <p>[“] My Classical Way [”] – was released – on Frazzy Frog Music</p> <p>“My Classical Way” – was released on – Marc ’s [own] label</p> <p>“My Classical Way” – was – released on Marc ’s [own] label</p> <p>“My Classical Way” – was released – on Marc ’s [own] label</p> <p>“My Classical Way” – was released on – Frazzy Frog Music</p> <p>“My Classical Way” – was – released on Frazzy Frog Music</p> <p>“My Classical Way” – was released – on Frazzy Frog Music</p>	<p>Cluster 1:</p> <p>[‘] [‘] My Classical Way [”] – was – released</p> <p>Cluster 2:</p> <p>[‘] [‘] My Classical Way [”] – was released on – 21 September 2010</p> <p>Cluster 3:</p> <p>[‘] [‘] My Classical Way [”] – was released on – Frazzy Frog Music</p> <p>Cluster 4:</p> <p>Frazzy Frog Music – is – Marc [’s] own label</p> <p>Frazzy Frog Music – is own label of – Marc</p> <p>Frazzy Frog Music – is owned by – Marc</p> <p>Marc [’s] own label – is – Frazzy Frog Music</p> <p>Cluster 5:</p> <p>Frazzy Frog Music – is [a] – label</p> <p>Frazzy Frog Music – is – [a] label</p> <p>Cluster 6:</p> <p>Marc – has [a] – label</p> <p>Marc – has – [a] label</p> <p>Marc – has – [own] label</p> <p>Marc – owns [a] – label</p> <p>Marc – owns – [a] label</p>

Table 21: Annotations examples from *BenchIE* and *BenchIE^{FL}*

J Annotation Guidelines

This Appendix contains annotation guidelines for the open information extraction task and has been used in the annotation process for the *BenchIE^{FL}* reference. The various principles dictate which facts should and should not be annotated. The information is presented in the following format: Sentences are in the top cell and in the cell below, examples that should be annotated are in **green** and preceded by a check mark, ones that should not be included are in **red** and preceded by a cross mark.

J.1 Number of arguments

All tuples must contain between 1 and 2 arguments. Extractions with more than two arguments can be split into more compact extractions. The first example shows this principle, while the second example shows how some tuples only have a single argument (we write XXX in the second argument for convenience).

J.2 Informativeness

Annotated tuples must contain relevant information that is expressed in the sentence. Tuples must be informative and relevant. They must not contain

<i>Kyle left for school on Monday.</i>
✗ (Kyle – left – for school – on Monday)
✓ (Kyle – left on – Monday)
✓ (Kyle – left for – school)

Number of arguments : First example

<i>Gideon Rodan taught at the University of Connecticut School of Dental Medicine.</i>
✓ (Gideon Rodan – taught – XXX)
✓ (Gideon Rodan – was – a teacher)

Number of arguments : Second example

generality or empty words that convey no information. In the example, the fact that *he has written* is not relevant since it is a generality, most people *have written* and it is not the information presented in the sentence.

<i>He has written several newspaper and magazine opinion pieces in The Guardian, Vice, Billboard, and others.</i>
✓ (He – has written – opinion pieces)
✗ (He – has – written)

Informativeness

J.3 Minimality

Annotated tuples must contain minimal information, which cannot be subdivided into smaller pieces of information. No argument should contain information about two different entities if this is true for both, and no tuple should contain more than one piece of information about an entity if these can be divided. In the first example, all the different minimal pieces of information (creators, time of creation) must be separated in minimal clusters and not grouped like in the example that should not be included. It is sometimes necessary to separate information, if and only if it is also true when separated. In the second example, the dog is neither *black* nor *brown*, but *black and brown*, whereas in the first example, *He has Cornish ancestors* **and** *He has Welsh ancestors*.

J.4 Exhaustivity

All the minimal information present in the sentence must be included in the annotations. Some argu-

<i>The group was created in 2020 by three people</i>
✓ (The group – was – created)
✓ (The group – was created in – 2020)
✓ (The group – was created by – three people)
✗ (The group – was – created in 2020 by three people)

Minimality : First example

<i>He has Cornish as well as Welsh ancestry.</i>
✓ (He – has – Cornish ancestry)
✓ (He – has – Welsh ancestry)

Minimality : Second example

ments or relations may be affected by modifications but remain true without them, so it's necessary to list all possible formulations that respect the other principles. In the example, it is true that *he wrote opinion pieces, newspaper opinion pieces* and *magazine opinion pieces*, so all these facts must be listed in three separate clusters.

J.5 Relation completeness

Relations are the vehicles of information; arguments must not contain information that changes their meaning. Relations can be complicated but necessary, while they can sometimes be simplified. They must be simplified as much as possible to respect the principle of minimality, without losing their original meaning, expressed in the sentence. In the example, the second argument of the erroneous annotation contains the word *over*, which modifies the meaning of the relation *is*, whereas in the positive example, the second argument, *13 Millions*, is only the object of the relation.

Sometimes, relationships can be complicated but necessary, while sometimes they can be simplified, keeping the additional part optional only if it's made necessary by the lack of other tuples explaining that additional part, as in the second example where the part *from Hungary* in the second cluster is optional because the place of origin of their

<i>The dog is black and brown.</i>
✓ (The dog – is – black and brown)

Minimality : Third example

<i>He has written several newspaper and magazine opinion pieces in The Guardian, Vice, Billboard, and others.</i>
✓ (He – has written – several newspaper opinion pieces)
✓ (He – has written – several magazine opinion pieces)
✓ (He – has written – several opinion pieces)

Exhaustivity

<i>Tokyo's population is over 13 Millions</i>
✓ (Tokyo's population – is over – 13 Millions)
✗ (Tokyo's population – is – over 13 Millions)

Relation completeness : First example

escape is present in the first cluster. In the third example, *in Paris* is not optional because without this information, the relation no longer holds.

<i>His parents are Ashkenazi Jews who had to flee from Hungary during World War II.</i>
✓ (His parents – had to flee from – Hungary)
✓ (His parents – had to flee [from Hungary] during – World War II)

Relation completeness : Second example

J.6 Coreference resolution

No coreference resolution is performed outside sentences. Even if a given sentence comes from a document that allows us to resolve a coreference, as OIE is intended to be a task performed on isolated sentences, we only resolve the coreferences of entities included in sentences taken in isolation. Tuples using pronouns for which we can't identify the substitution element may seem meaningless, but coreference resolution must take place outside OIE, being a task in itself. In the example, we don't do coreference resolution for the pronoun *He*, as no information about it is available in the sentence. However, we include a formulation replacing *them* with *tax reductions* in the annotation.

Chilly Gonzales is a Grammy-winning Canadian musician who resided in Paris, France for several years, and now lives in Cologne, Germany.

✓ (Chilly Gonzales – resided in Paris for – several years)

Relation completeness : Third example

He did not go as far as he could have in tax reductions ; indeed he combined them with increases in indirect taxes .

✓ (He – combined them with – increases in indirect taxes)
(He – combined tax reductions with – increases in indirect taxes)

Coreference resolution

J.7 Inference

Inference is necessary: facts directly implied by the sentence, even if not expressed verbatim, are relevant pieces of information. A nuance is necessary here with regard to potential implicit facts. These are not necessarily implied by the sentence and should therefore be omitted. In the first example, it is necessarily implied by the sentence that *Paul Johanson* is *Monsanto's Director of Science*. On the other hand, the fact that *Monsanto's spray* is *gentle on the female organ* is not necessarily true, what is true that this information is said by *Paul Johanson*.

However , Paul Johanson , Monsanto 's director of plant sciences , said the company 's chemical spray overcomes these problems and is gentle onthe female organ .

✓ (Paul Johanson – is – Monsanto's director of science)

✓ (Paul Johanson – says – the company's chemical spray is gentle on the female organ)

✗ (the company's chemical spray – is gentle on – the female organ)

Inference : First example

It is then necessary to distinguish between light and heavy inference. We define light inference as a form of inference that does not require logical

reflection with respect to the sentence to deduce the fact, which is simply true as long as the sentence is also true. In the second example, the relation is implicit, but the annotated fact is obviously true. Heavy inference, on the other hand, requires some reflection or combination of logical operations to imply the fact. A case of heavy inference is that which requires external knowledge, as in the third example, where knowledge of human culture and the principle of heredity is necessary to make the inference. Another example of heavy inference is generalization, as in the fourth example, where a stronger fact is implied, a generalization of what is expressed in the sentence using a single example. A final example of heavy inference is that of lower or upper limits. As in the fifth example, we don't want to generalize lower or upper bounds to entities that are not directly expressed in the sentence. We therefore include in the reference facts that can be inferred using light inference, but not those resulting from heavy inference.

Jason Charles Beck, a Jewish Canadian musician, was born in 1972.

✓ (Jason Charles Beck – is – Jewish)

Inference : Second example

Gonzales is the son of Ashkenazi Jews who were forced to flee from Hungary during World War II.

✗ (Gonzales – is – Jewish)

Inference : Third example

Gonzales is a McGill-trained virtuoso pianist.

✗ (McGill – trains – pianists)

Inference : Fourth example

J.8 Reformulation

If a relation or argument is expressed in a complex way in the text, a simpler re-formulation of the same fact is added in the same cluster, even if the relation in the two formulations is not the same and the level of detail may be different. This is a compromise between the goal of OIE of collecting all the factual information expressed in the text and the importance of formulating these facts in simple language, which is relevant but not necessarily OIE's primary goal. The example shows a

The prefecture is part of the world's most populous metropolitan area with upwards of 37.8 million people and the world's largest urban agglomeration economy.

✗ (the world – has upwards of – 37.8 million people)

Inference : Fifth example

case where the reformulated relation is different but conveys the same meaning.

Sam managed to convince John

✓ (Sam – managed to convince – John)
(Sam – convinced – John)

Reformulation

J.9 Active and Passive Voice

Clusters are used to group together the active and passive formulations of an extraction in a single fact. If the active formulation is not present in the text, it should still be added in the same cluster if it is simpler than the original formulation present in the sentence. The example shows an originally passive tuple and it's active formulation added in the same cluster.

The apple was eaten by Kyle

✓ (The apple – was eaten by – Kyle)
(Kyle – ate – the apple)

Active and Passive Voice

J.10 Attribution and Speculation

Some information in the text is speculative or attributed to an entity, so this characteristic must be included in the relationship, in the way it is formulated in the sentence. This makes it possible to preserve this information without having to introduce a particular structure. This information must be included in the relation, as it is in no way related to the arguments. The example shows a case where the attribution is added in the relation of the tuple.

J.11 Correction

Occasionally, some tuples may consist of words from the original sentence but contain grammatical errors. In this case, the tuple formed from the original words and the corrected tuple should be included in the same cluster. This ensures that

The earth is flat, according to an Apple Valley man.

✓ (The earth – is according to an Apple Valley man – flat)

✗ (The earth – is – flat)

Attribution and Speculation

neither the systems making the correction nor those using the original text are penalized. The example shows that *newspaper* without an *s* is a grammatical error, so both the original and the correction should be included.

He has written several newspaper and magazine opinion pieces in The Guardian, Vice, Billboard, and others.

✓ (He – has written in – newspaper)
(He – has written in – newspapers)

Correction

K Matching Guidelines

This Appendix contains the matching guidelines for the open information extraction task and has been used in the development of the *BenchIE^{FL}* reference matching function. The various principles dictate which pairs of extractions made by systems and annotations should and should not match. The information is presented in the following format: Sentences are in the top cell and in the cell below, the different formulations of the same cluster (of the same fact) are in a paragraph and a line break separates them. Clusters in black represent annotations. Examples in green, preceded by a check mark are examples that match an annotation in the reference, while examples in red, preceded by a cross mark do not.

K.1 Exact match

Two absolutely identical extractions should match.

K.2 Relation specificity

Extractions are allowed very little flexibility in the specificity of the relation: the relation is the vehicle of information, so it's important that it's almost as specific as the reference. That said, a different formulation that is just as specific should be accepted. In the example, *was thrown* is not a relevant relation in the context of this extraction, as *was* or *was thrown out of* would have been (the word *out* in argument 2 changes the meaning of the relation).

<i>The Finns party was thrown out of the government and the new “Blue Reform” group kept its cabinet seat.</i>
(The Finns party – was – thrown out of the government) (The Finns party – was thrown out of – the government)
✗ (The Finns party – was thrown – out of the government)

Relation specificity

K.3 Errors

Some extractions made by systems may present syntax or grammatical errors, when a word is misplaced or unnecessary. If this error changes the meaning of the relation or one of the arguments, the extraction should not be matched. If not, it should match the corresponding annotation. In the example, the word *also* refers to the relation *is*, and does not change the meaning of the relation, whereas the word *and* changes the meaning of the extraction, making it nonsensical.

<i>Known for his albums of classical piano compositions, he is also a producer and songwriter.</i>
(He – is [also] – a songwriter)
✗ (He – is – a songwriter and)
✓ (He – is – a songwriter also)

Errors

K.4 Word Choice

Some words may be equivalent to those present in the annotations in certain contexts, even if we have chosen not to include them in the reference. If these words are used in the system extractions instead of those used in the reference, we still accept the system extraction. Some word choices may be wrong, but we still accept the extraction if the meaning remains. In the example, the determiner *the* is used instead of *a* in the extraction because it's the word found in the original sentence, but both are equally appropriate, so we accept the extraction.

<i>He is the younger brother of the prolific film composer Christophe Beck.</i>
(He – is – a younger brother)
✓ (He – is – the younger brother)

Word Choice

K.5 Level of Detail

We want to match extractions which have a level of detail higher than the annotation but that convey the same information. By level of detail we mean that they combine information from two annotated clusters. On the other hand, if an extraction combines information from three or more annotated clusters, we consider it to be too noisy and not precise enough to be useful. The positive example is matched because it conveys the same information as the second annotated cluster, and only adds a single level of detail from the third cluster. The negative example is not matched because it combines information from all three annotated tuples into a long and imprecise second argument.

<i>Alex broadcasts a web series Music on a website.</i>
(Alex – broadcasts – a web series)
(Alex – broadcasts – Music)
(Alex – broadcasts Music on – a website)
✓ (Alex – broadcasts – Music on a website)
✗ (Alex – broadcasts – a web series Music on a website)

Level of detail