

In Layman’s Terms: Semi-Open Relation Extraction from Scientific Texts

Ruben Kruiper*, Julian F.V. Vincent, Jessica Chen-Burger,
Marc P.Y. Desmulliez, Ioannis Konstas

Heriot-Watt University
Riccarton Campus, EH14 4AS
Edinburgh, United Kingdom

*Corresponding author: rk22@hw.ac.uk

Abstract

Information Extraction (IE) from scientific texts can be used to guide readers to the central information in scientific documents. But narrow IE systems extract only a fraction of the information captured, and Open IE systems do not perform well on the long and complex sentences encountered in scientific texts. In this work we combine the output of both types of systems to achieve Semi-Open Relation Extraction, a new task that we explore in the Biology domain. First, we present the Focused Open Biological Information Extraction (FOBIE) dataset and use FOBIE to train a state-of-the-art narrow scientific IE system to extract trade-off relations and arguments that are central to biology texts. We then run both the narrow IE system and a state-of-the-art Open IE system on a corpus of 10k open-access scientific biological texts. We show that a significant amount (65%) of erroneous and uninformative Open IE extractions can be filtered using narrow IE extractions. Furthermore, we show that the retained extractions are significantly more often informative to a reader.¹

1 Introduction

Identifying the central theme and concepts in scientific texts is a time-consuming task for experts and a hard task for laymen (Alper et al., 2004; El-Arini and Guestrin, 2011; Pain, 2016). This problem is even more pronounced in inter-disciplinary fields of study, where experts in a target domain often lack the deeper knowledge of a source domain (Carr et al., 2018). A specific example is biomimetics, an engineering problem-solving process in which one draws on analogous biological solutions (Kruiper et al., 2016). A major issue is that engineers (target domain) know little biology (source domain) or characteristics of plants or

¹We release FOBIE and code at <https://github.com/rubenkruiper/FOBIE>.

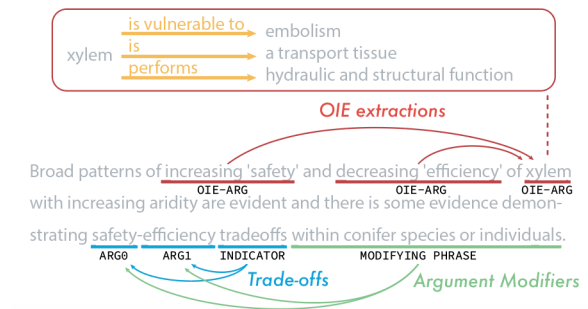


Figure 1: Example of Semi-Open Relation Extraction in the Biology domain (Burgess et al., 2006). We first extract the TRADE-OFF expressed between the central concepts ‘safety’ and ‘efficiency’ (in blue), that takes place specifically in ‘conifer species’ (in green). We further explore the content of a paper by investigating the results of an Open Information Extraction (OIE) system (in red). The central concepts captured by a TRADE-OFF mechanism enable the filtering of many irrelevant OIE extractions, which are found to be error-prone in scientific texts. Further OIE extractions found in the document can shed light on the semantic meaning of relevant concepts, e.g., ‘xylem’, as depicted at the top of the Figure (in the red box).

animals (Vattam and Goel, 2013). This domain-mismatch complicates searching for and reasoning over relevant scientific information, rendering biomimetics adventitious and solutions serendipitous (Kruiper et al., 2018).

Recently, TRADE-OFF relations have become of interest to biomimetics (Adriaens, 2019) because a trade-off defined in technology can be directly used to search for relevant texts in biology (Vincent, 2016). TRADE-OFF relations express a problem space in terms of mutual exclusivity constraints between competing demands. Therefore, trade-offs play a prominent role in evolutionary thinking (Agrawal et al., 2010) and are the principal relation under investigation in a significant portion of biology research papers (Garland, 2014). The

functional demands that are traded off are usually abstract and domain-independent terms, such as ‘*safety*’ and ‘*efficiency*’ in Figure 1. A gap remains in quickly comprehending the central information in a text, e.g., the biological mechanisms that are used to manipulate a trade-off.

Information Extraction (IE), and specifically Relation Extraction (RE), can improve the access to central information for downstream tasks (Santos et al., 2015; Zeng et al., 2014; Jiang et al., 2016; Miwa and Bansal, 2016; Luan et al., 2018a). However, the focus of current RE systems and datasets is either too *narrow*, i.e., a handful of semantic relations, such as ‘USED-FOR’ and ‘SYNONYMY’, or too *broad*, i.e., an *unbounded* number of generic relations extracted from large, heterogeneous corpora (Niklaus et al., 2018), referred to as Open IE (OIE) (Etzioni et al., 2005; Banko et al., 2007). Narrow approaches to IE from scientific text (Augenstein et al., 2017; Gábor et al., 2018; Luan et al., 2018a) cover only a fraction of the information captured in a paper – usually what is within an abstract. It has been shown that scientific texts contain many unique relation types and, therefore, it is not feasible to create separate narrow IE classifiers for these (Groth et al., 2018). On the other hand, OIE systems are primarily developed for the Web and news-wire domain and have been shown to perform poorly on scientific texts. What laymen really need is a bit of both: the accuracy of narrow RE systems to extract central relations from scientific texts *and* the flexibility of an OIE system to capture a much larger fraction of the possible relations expressed in scientific texts.

This work aims to enable rapid comprehension of a large scientific document by identifying a) the central concepts in a text and b) the most significant relations that govern these central concepts. To this end, we introduce the task of Semi-Open Relation Extraction (SORE); Figure 1 illustrates the SORE process. First, we find the central concepts ‘*safety*’ and ‘*efficiency*’ involved in a TRADE-OFF relation. Then, by using the argument concepts of the relation as anchor points, we can explore further concepts and relations, e.g., ‘*xylem*’ in Figure 1. Uncovering these relations can elucidate the meaning of unfamiliar concepts to a layperson (Mausam, 2016). The SORE approach is hypothesized to reduce the number of uninformative extractions without limiting RE to a finite set of relations, which could generally benefit IE from sci-

entific articles, e.g., materials discovery (Kononova et al., 2019) and drug-gene-mutation interactions (Jia et al., 2019).

To address SORE we create the Focused Open Biological Information Extraction (FOBIE) dataset. FOBIE includes manually-annotated sentences that express explicit trade-offs, or syntactically similar relations, that capture the central concepts in full-text biology papers. We train a span-based RE model used in a strong scientific IE system (Luan et al., 2018a) to jointly extract these relation structures. We explore SORE and use the output of our model to filter the output of an OIE system (Saha and Mausam; Saha et al., 2017; Pal and Mausam, 2016; Christensen et al., 2011) on a corpus of biology papers. Qualitative analyses show that the output of a narrow RE model can speed up expert analysis of trade-offs in biological texts, and be used to filter out both erroneous and uninformative OIE extractions.

2 Related work

2.1 Open Information Extraction

OIE systems use a set of handcrafted or learned extraction rules and rely on dependency features to extract open-domain relational tuples from text (Yu et al., 2017; Niklaus et al., 2018). As OIE systems rely on syntactic features they require little fine-tuning when applied to different domains and the extraction rules work for a variety of relation types (Mausam, 2016). These properties can be especially useful on scientific texts where additional knowledge on unknown concepts can ease the textual comprehension for non-experts. Consider the example OIE extractions for ‘*xylem*’ in the top part of Figure 1.

Existing OIE systems have been shown to perform significantly worse on the longer and more complex sentences found in scientific texts than on Wikipedia texts (Groth et al., 2018). Common issues of OIE systems on Web, News, and Wikipedia texts include the correct identification of the boundaries of an argument, handling latent *n*-ary relations, difficulty handling negations, and generating uninformative extractions (Schneider et al., 2017). Groth *et al.* (2018) evaluate the output of two state-of-the-art OIE systems based on correctness, rather than, e.g., the number of missed extractions. They note that the crux of the IE challenge is that extractions reflect the consequence of the sentence. As an example of an uninformative extraction Fader *et al.*

(2011) note how ‘(*Faust, made, a deal*)’ captures the consequence, but not the critical information of whom Faust made a deal with in the sentence “*Faust made a deal with the devil.*”. In this work, we explore filtering both incorrect and uninformative OIE extractions from scientific texts using the central concepts that we extract through narrow IE (cf. Section 5.3).

2.2 Narrow Relation Extraction from scientific text

Narrow RE entails identifying two or more related entities in a text and classifying the relation that holds between them. Early works on the combined task of Named Entity Recognition and labeling of relations between extracted entities used pre-computed dependency features (Liu et al., 2013; Chen et al., 2015; Lin et al., 2016), word position embeddings (Zeng et al., 2014), or considered only the Shortest Dependency Path between two entities as input (Bunescu and Mooney, 2005; Santos et al., 2015; Zeng et al., 2015). Later work aimed to reduce errors propagated by pre-computed dependency features (Nguyen and Grishman, 2015), or by joint modeling of entities and relations (Miwa and Bansal, 2016). Poor performance of these RE systems on scientific texts has led to the development of domain-specific datasets².

The SCIENCEIE dataset focuses on the extraction of 3 types of key-phrases, rather than Named Entities, and hyponymy and synonymy relations between these (Augenstein et al., 2017). The SemEval 2018 task 7 dataset focuses on 6 narrow relations between 7 entity types (Gábor et al., 2018). And the SCIERC dataset focuses on 7 relation types, including co-reference, between 6 types of entities (Luan et al., 2018a). Top systems developed for both SemEval tasks adapt the LSTM-based approach of Miwa & Bansal (2016), combined with semi-supervised learning and ensemble (Ammar et al., 2018), as well as pre-trained concept embeddings (Luan et al., 2018b).

2.3 BioNLP and BioCreAtIvE

In the past, several BioNLP and BioCreAtIvE shared tasks were organized that aimed at identifying relations in the biology domain (Hirschman

²SCIENCEIE SemEval 2017: 500 paragraphs from full-text Computer Science, Material Science, and Physics journal articles, SemEval 2018: 500 abstracts within the domain of Computational Linguistics. SCIERC: 500 abstracts from Artificial Intelligence conference and workshop proceedings.

#	FOBIE	SCIENCEIE	SE '18	SCIERC
Arg's	5834	9946	7483	8089
Rel's	4788	672	1595	4716
R/doc	3.09*	1.3	3.2	9.4

Table 1: Number of arguments, relations and relations per instance for FOBIE, SCIENCEIE, the SemEval 2018 task 7 dataset and SCIERC. R/doc stands for relations per sentence* for FOBIE (and per abstract or paragraph for the other datasets).

et al., 2005; Kim et al., 2009; Nédellec et al., 2013; Zhou et al., 2014). Many datasets focus primarily on a predefined set of biomedical relations, such as interactions between known proteins, genes, diseases, drugs, and chemicals (Kim et al., 2003; Krallinger et al., 2017; Cohen et al., 2017; Islamaj Doğan et al., 2019). Examples of more biology-oriented corpora include the BB corpus (Deléger et al., 2016) and the SEEDDEV corpus (Chaix et al., 2016). The BB corpus includes 4 entity types and 2 relation types that revolve around microorganisms of food interest. Besides abstracts and titles, it contains paragraphs and sentences from 20 full-text documents (Bossy et al., 2019). Similarly, SEEDDEV consists of 86 paragraphs from 20 full-text articles about seed development in a specific plant, the *Arabidopsis thaliana*. Considering the small size of the dataset, a relatively large number of many entity and relation types are used; 16 types of Named Entities and 21 types of relations. This results in an imbalanced dataset with 7 relations making up less than 1% of all relations. Furthermore, there is some overlap in source documents for the train/dev/test split (Chaix et al., 2016).

In contrast to the previously described datasets, FOBIE does not classify arguments of relations into specific entity-types. FOBIE contains annotations of key-phrases found in full-text scientific papers, similar to SCIENCEIE. The key-phrases and relations are annotated in 1,548 relatively long and complex sentences, which were sourced from 1,215 full-text scientific biological texts using a Rule-Base System. Table 1 provides an overview of the size of FOBIE in comparison to SCIENCEIE, the SemEval 2018 task 7 dataset and SCIERC. Both the BB and SEEDDEV corpus contain approximately 3,500 relations within a small sub-domain of biology, while FOBIE focuses more generally on the domain of biology. Section 3 describes the collection of FOBIE and dataset statistics in detail.

3 Dataset description

3.1 Dataset collection

A variety of words are able to indicate a trade-off, e.g., *compromise*, *optimization*, *balance*, *interplay* and *conflict* (Kruiper et al., 2018). We adapt these terms as trigger words in a Rule-Based System (RBS) and run it on 10k open-access papers that were collected from the Journal of Experimental Biology (JEB) and BioMed Central (BMC) journals on ‘*Biology*’, ‘*Evolutionary Biology*’, and ‘*Systems Biology*’. The selection of journals was made only to the extent that the articles focus on the biological domain. We retained the abstract, introduction, results, discussion and conclusion sections. We used spaCy³ to split the texts into sentences and identify POS tags and dependency structure. The FOBIE dataset contains only sentences that the RBS identified as expressing a TRADE-OFF relation.

3.2 Annotation

The initial annotations extracted by the RBS were manually corrected and extended by a biology expert using the BRAT interface (Stenetorp et al., 2012). We define three relation types: TRADE-OFF, ARGUMENT-MODIFIER and NOT-A-TRADE-OFF. The latter denotes phrases that are related to a trigger word, but not by a TRADE-OFF relation. These syntactically similar relations provide useful training signal as negative samples. Negative samples are important because possible trigger words can be contiguous, e.g., the phrase ‘*negative correlation*’ denotes a TRADE-OFF relation, whereas ‘*correlation*’ by itself does not. As a result, the annotation of training examples is harder, and lexical and syntactic patterns that correctly signify the relation are sparse (Peng et al., 2017). For simplicity’s sake, with some abuse of terminology, we refer to all such relations collectively as *trade-offs*.

We found a substantial amount of arguments to be nested or in a non-projective relationship. In Figure 2 the prepositional phrase ‘*in jumping*’, conceptually refers to both central concept arguments of the relation, i.e., ‘*the need for energy storage*’ and ‘*the presence of resilin*’. We adopt the following annotation heuristic: prepositional phrases are treated as modifying phrases when they apply to multiple arguments (as is the case in Figure 2) or can be distinctly separated from the argument, e.g., by punctuation.

³<https://spacy.io/>

# Sentences	1548
Avg. sent. length	37.77
% of sents \geq 25 tokens	79.26%
Relations:	
- TRADE-OFF	765
- NOT-A-TRADE-OFF	2502
- ARG-MODIFIER	1521
Triggers	1600
Arguments	4234
Spans	6309
Unique spans	2701
Unique triggers	41
Max triggers/sent	2
Max spans/sent	7
Spans w/ multiple relations	2075
# single-word arguments	498
Avg. tokens per argument	3.44

Table 2: Aggregated statistics for FOBIE.

We randomly selected 250 sentences (16.1%) for re-annotation and quality control by a second domain expert. The inter-annotator agreement Cohen k is found to be 0.93. Table 2 summarizes statistics on FOBIE. The final dataset consists of 1,548 single sentences from 1,292 unique documents, split into 1,248/150/150 train/dev/test. The split is controlled for source document overlap to avoid having identical arguments of relations appearing both during training and testing. FOBIE contains relatively long key-phrases with an average of 3.44 tokens and only 12% of them consist of a single token. In comparison SCIENCEIE and SCIERC both contain 31% singleton key-phrases, and the average entity length in SCIERC is 2.36. Furthermore, sentences taken from full-text documents are longer than those found in abstracts. The average sentence length in SCIERC is 24.31 tokens, while 79.26% of the sentences in FOBIE are longer than 25 tokens.

4 Narrow IE baseline

4.1 Task definition

Following Peng et al. (2017) we extract n -ary relations by (1) identifying the trigger and (2) extracting the binary relations between this trigger and the arguments – inspired by Davidsonian semantics. We define key-phrases as spans of consecutive words $s \in S$, with S all possible spans in a sentence, and relation-types as $r \in R^d$, with d the total number of unique relations. Then a binary relation is a triple $\langle \text{governor}, \text{relation}, \text{dependent} \rangle$ with *governor* and *dependent* elements of S . The union of the following binary relations found in a sentence may constitute a non-projective graph:



These observations indicate a clear correlation between the presence of resilin and the need for energy storage in jumping.

Figure 2: Example of an annotation in BRAT showing a trigger word, ‘correlation’, that is related to two arguments, which in turn are related to a single modifier. The trigger word does not indicate a TRADE-OFF relation, but a positive correlation.

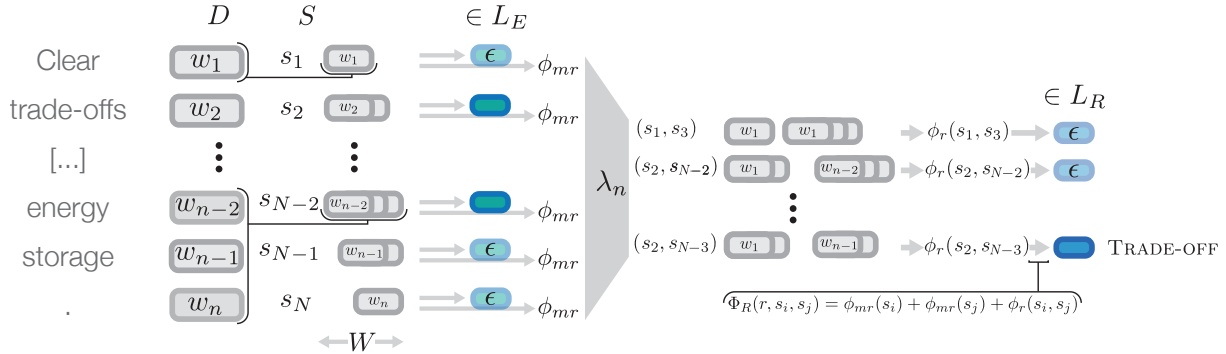


Figure 3: We provide the SCIIE system with single sentences D as input. For all possible spans up to width W a span label $\in L_E$ is computed and a mention score ϕ_{mr} . Spans with the lowest mention scores are pruned, with variable beam size λ_n . For combinations of remaining spans a relation label $\in L_R$ is predicted. The set of span labels L_E and the set of relation labels L_R both contain a dummy class ϵ .

Def. 1. An explicit trade-off is an instance of a directed relation $t \in T^o$, indicated by trigger word $p \in P^u$ with u the set of unique trigger words and $P \subset S$. A trade-off is a binary relation, $t \models o$, with *governor* $\in P$ and *dependent* $\in S$. A single trigger word p can be in n multiple relations.

Def. 2. An argument-modifier is a directed binary relation $a \in A^m$, where we omit the classification of a into a set of possible modification types $\in m$. An instance of a is then a tuple $\langle \text{governor}, \text{relation}, \text{dependent} \rangle$ where one of the arguments is related to a trigger word p , and both arguments $\in S$.

4.2 Baseline system

We adapt a span-based approach that has been used previously for the tasks of co-reference resolution (Lee et al., 2017), Semantic Role Labeling (He et al., 2018), and scientific IE (Luan et al., 2018a). The use of span representations as classifier features enables end-to-end learning by propagating information between multiple tasks without increasing the complexity of inference. We train the SCIIE system (Luan et al., 2018a) on FOBIE to extract spans that constitute trigger words and key-phrases, as well as the binary relations between these spans.

Figure 3 illustrates the input that we provide to SCIIE. All tokens are embedded using GloVe (Pennington et al., 2014) and ELMo embed-

dings (original) (Peters et al., 2018). For a single sentence $D = \{w_1, \dots, w_n\}$ all possible spans $S = \{s_1, \dots, s_N\}$ are computed, which are within-sentence word sequences.

The model deals with $\mathcal{O}(n^4)$ possible combinations of spans, where n is the number of words in a sentence. Therefore, pruning is required to make the classification of span-pairs into relation labels tractable at both training and test time (Lee et al., 2017; He et al., 2018). First, a score ϕ_{mr} of how likely a span is mentioned in a relation is computed. These mention scores enable beam pruning the number of spans considered for relation classification with a variable beam of size λ_n , where n is the number of tokens in the input sentence (Luan et al., 2018a). Second, the maximum width W of spans is limited to reduce the total number of spans. We set λ to .8 and W to 14 tokens, the maximum span length in FOBIE.

After pruning, a label $e_i \in L_E$ is predicted for the remaining spans s_i . Here L_E is the set of possible span labels, including a non-span class ϵ . For pairs of spans (s_i, s_j) the model predicts which relation $r_{ij} \in L_R$ holds between them. The set of possible relation types is L_R , which includes a non-relation class ϵ . The output consists of labeled spans and relation labels for pairs of spans. For a detailed description of the SCIIE system we refer to Luan et al. (2018a).

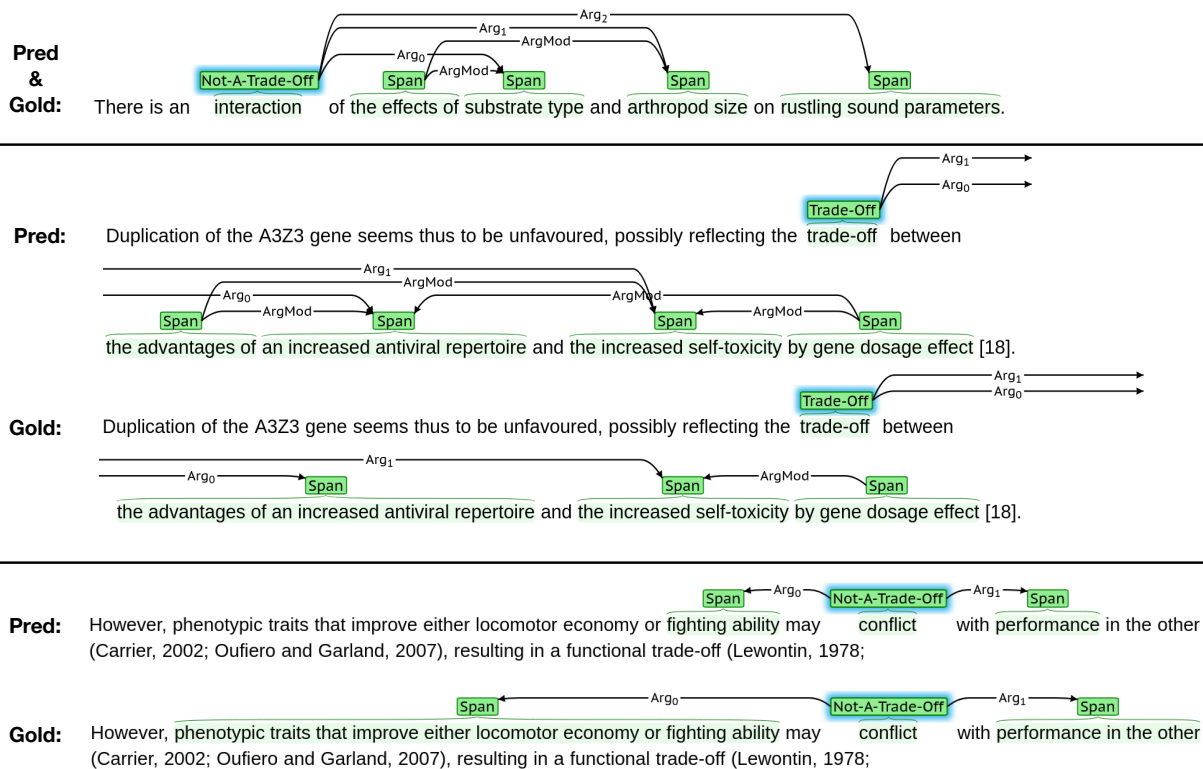


Figure 4: Examples of output from the SCIIE system trained on FOBIE, in comparison to gold annotation.

4.3 Narrow IE results

We evaluate SCIIE on two sub-tasks: (1) Argument Recognition, and (2) Relation Extraction. Table 3 summarizes the results on the sub-tasks of Argument Recognition and RE. With regards to the first sub-task, we train two SCIIE models. One model only predicts whether a span is a valid span or not, while a second model predicts whether the span is a trigger word or a key-phrase. For the first sub-task we also report the results of the RBS described in Section 3.1. The RBS performs significantly worse; it identifies trigger words exceptionally well (F1=95.89 on test set) but does not correctly recognize many of the remaining key-phrases (F1=22.36 on test set), resulting in a low overall performance.

Figure 4 shows example outputs of the narrow RE model. The predicted relation (NOT-A-TRADE-OFF) and its accompanying structure for the first example are completely correct. Note how the argument modifiers result in a non-projective structure. The second example is more challenging, with a longer range dependency between the trade-off span and the second dependent argument. Our model predicts the correct relation, TRADE-OFF, but only extracts partial argument spans and essentially fragments them into several modifying

argument relations. The third example exhibits a relatively long argument – which is common in scientific literature – where only a small part of the span is predicted.

4.4 Supporting trade-off annotation

A qualitative analysis confirms the ability of the trained narrow IE system to support a domain expert during trade-off annotation. We predict trade-offs for 523 unlabeled, scientific papers that have been annotated with a trade-off in an ontology of biomimetics (Vincent, 2014, 2016). A domain expert compares the trade-offs found in the ontology of biomimetics against the output of the SCIIE system, see Table 4. Narrow IE is found to locate the central TRADE-OFF relations and arguments for 41.68% of the total 523 papers. Explicit trade-offs were found in 243 documents. At least one of the extracted TRADE-OFF relations for each document is identical to the expert annotation in 77.37% of these documents. For 89.71% of the 243 documents a trade-off was found to be correct after some interpretation by the expert. Two main types of uninformative trade-offs were found: trade-offs from a cited source and trade-offs between generic terms, e.g., a trade-off between cost and benefit without defining what the cost and benefit are.

<i>Argument Recognition</i>	P	R	F1
RBS	44.31	35.32	39.31
SciIE unlabeled spans	86.76	79.39	82.91
SciIE labeled spans	83.86	80.83	82.32
<i>Relation Extraction</i>	P	R	F1
RBS	—	—	—
SciIE unlabeled spans	68.53	65.48	66.97
SciIE labeled spans	67.71	67.71	67.71

Table 3: Results on test set for our Rule-Based System (RBS) and SciIE (Luan et al., 2018a) w.r.t argument recognition and the combined task of extracting and classifying relations (RE). Providing the model with labels to distinguish trigger words from key-phrases slightly improves performance.

Documents with identified trade-offs	243
Exact match	77.37%
Match after interpretation	89.71%
Sentences with identified trade-off	998
Exact match	68.04%
Match after interpretation	84.47%

Table 4: Manual analysis of extractions from 523 scientific documents that were used in the creation of an ontology of biomimetics (Vincent, 2014, 2016).

5 Semi-Open Relation Extraction

5.1 Task description

We define the aim of SORE as extracting the relations and concepts in a text that capture the most central information. The application of SORE is especially of interest to scientific IE where OIE systems perform poorly and narrow IE systems are unable to cover the wealth of different relations types. One possible approach is to automatically filter out uninformative and incorrect extractions generated by OIE systems. In this approach, SORE relies on the output of both types of systems, providing a middle ground between precise, narrow IE and unbounded, but unreliable, OIE. The resulting extractions are expected to be useful for human readers, but can also be used to collect data for annotation and training of scientific IE systems.

5.2 Experimental setup

We explore SORE on scientific biology texts using the output of the SciIE system trained on FOBIE, predicting trade-offs for the unlabeled 10k open access biology papers (see section 3.1). The narrow IE output consists of 2,216 trade-offs found in 1,279 documents. We pre-process arguments by appending their modifier, removing stop words, and embedding the remaining sequences using ELMo

(PubMed)⁴. We use the K-means algorithm to compute clusters on the IDF-weighted average of the resulting argument representations. A domain expert inspected the centroids qualitatively. Table 5 provides insight into some of the resulting argument clusters and their interrelations. The exact number of clusters does not seem to greatly affect SORE. For the given narrow IE output ± 50 clusters seems to provide a good balance between generic and more fine-grained topics. The IDF weights are computed over the subword units found in the dataset; we use SentencePiece⁵ with a vocabulary of 16K. We then run OpenIE 5, a state-of-the-art OIE system (Saha and Mausam; Saha et al., 2017; Pal and Mausam, 2016; Christensen et al., 2011), on the same 1,279 documents that were found to contain one or more TRADE-OFF relations.

We retain only OIE extractions that contain one or more arguments that are classified into the same cluster as the TRADE-OFF arguments found in that text. Furthermore, we omit OIE arguments that belong to noisy clusters containing mostly math symbols or long nested phrases. We compute a simple IDF-weighted cosine similarity (Galárraga et al., 2014) between the vector representations of the remaining OIE and trade-off arguments.

5.3 Qualitative analysis of SORE output

We notice a striking drop in the number of irrelevant and noisy OIE arguments that remain after applying SORE. The total amount of OIE extractions reduces from 401k before filtering to 140k (34.95%) after filtering. As a result, the number of OIE extractions per document reduces from 314 to 110. The unfiltered OIE extractions are found in 170k sentences, of which 67k (39.55%) are retained after applying SORE.

To test our hypothesis that SORE can reduce the number of uninformative extractions, without limiting RE to a narrow set of relations, we randomly select representative samples of unfiltered and filtered OIE extractions (400 each). A domain expert manually annotated whether each extraction or sentence was thought to be informative, e.g., provides relevant information to understanding a biological text. As an example, consider the sentence “*We have used this approach in a previous study to investigate the molecular factors governing the altered liver regeneration dynamics caused*

⁴<https://allennlp.org/elmo>

⁵<https://github.com/google/sentencepiece>

Cluster name	Immunity	Size	Locomotion
Top-5 arguments	immunity	size	swimming
	immune function	number	sprinting
	the immune system	volume	running
	incompetence	age	locomotion
	immune response	time	diving
Top-3 related clusters	Mating	Temperature	Attribute of Animal
	Reproduction	Sperm Length	Verbs
	Life History Traits	Offspring Number	Capacity/Endurance

Table 5: Examples of clusters found using the K-means algorithm on trade-off arguments from 1279 documents. For the related clusters only TRADE-OFF relations are taken into account.

by ablation of the gene *adiponectin* (*Adn*)” (Cook et al., 2015). OIE extractions such as ‘(We, have used, this [...] study)’ are considered uninformative, in contrast to ‘(the molecular [...] dynamics, caused by, ablation [...] *adiponectin*)’.

Many OIE extractions are found to be poorly structured. Like Groth et al. (2018) we relax the requirement of extractions being well-formed, e.g., we consider extractions that incorrectly identify the boundaries of one or more arguments as possibly capturing relevant information. Different from their evaluation on correctness, we evaluate whether an extraction captures information that is relevant to understanding a text. As a result, we consider poorly structured OIE extractions that contain relevant information to be informative, e.g.:

- (*‘the resumption of respiration’, ‘can lead to an increase of superoxide anions in the cytosol perhaps driving’, ‘increased elevation of Cu-ZnSOD’*).
- (*‘transcriptional coregulation amongst many genes’, ‘will give’, ‘rise to indirect interaction effects in mRNA expression data’*).

The annotation relies on the correctness of the information captured by OIE extractions and whether this information is useful to a reader. However, this does not imply informative extractions are relevant to the central theme of the text captured in a trade-off. We consider OIE extractions uninformative if the extraction:

- contains an uninformative argument class, e.g., (*‘Miller et al. , 2012’, ‘to minimize’, ‘their swimming effort’*).
- contains incomplete arguments, e.g., (*‘the RDME requirement’, ‘reactions’, ‘only fire’*).
- is non-sensible, e.g., (*‘P. magellanicus’, ‘would have resulted’, ‘in a 1.6-fold higher Vmax for the scallop muscle’*).

- is unlikely to help understand a text, e.g., (*‘DeepBind’, ‘was trained’, ‘on data from RNAcompete , CLIP - RIP - seq [10’*) and (*‘microlepidopteran superfamilies’, ‘are heavily entombed’, ‘L:in amber’*).

We also randomly select representative samples from the 170k unfiltered and 67k filtered sentences from which the OIE extractions are sourced. The reason is that erroneous OIE extractions, e.g., not well-formed tuples, can guide a reader to informative passages in a text. We see similar errors as described by Schneider et al. (2017) and Groth et al. (2018), e.g., long sentences lead to incorrect extractions and errors in argument boundaries. To illustrate the complexity of sentences that an OIE system encounters in scientific texts, consider the following examples:

- the arity of relations can be high, e.g., (49 tokens) “A large genome size tends to correlate with delayed mitotic and meiotic division [6–8] decreased plant invasiveness of disturbed sites [9] lower maximum photosynthetic rates in plants [2] and lower metabolic rates in mammals [10] and birds [11, 12].” (Warringer and Blomberg, 2006).
- many phrases are nested and express non-verbal relations, e.g., (45 tokens) “However, for arboreal animals that regularly jump between branches (often when elevated quite high above the ground), jumping **accurately (which we define as the ability to land close to the intended target)** may also be important to fitness.” (Kuo et al., 2011).

Table 6 provides an overview of the annotation results. Filtering is found to increase the informativeness of both OIE extractions ($\chi^2=6.39$, $p<.025$) and sentences ($\chi^2=11.75$, $p<.01$). The percentage of informative OIE extractions increases by 5.75% and of the percentage of informative sentences by

	Filtering	# Nr.	% Informative
<i>OIE extractions</i>	before	401,588	29.25%
	after	140,357	35.00%
<i>Sentences</i>	before	170,551	36.50%
	after	67,460	44.75%

Table 6: Total number of OIE extractions before and after filtering, as well as the sentences that these extractions were found in. The % informative denotes the percentage of extractions and sentences annotated as informative by a domain expert, based on 400 randomly sampled instances from each group (95% confidence interval, margin of error 5%).

8.25%. A second domain expert annotated 25% of each set (400 total), the inter-annotator agreement Cohen k was found to be 0.84.

5.4 Results

Manual inspection of the retained OIE extractions shows that many relevant extractions are retained, e.g., see Table 7. These extractions are useful to a reader in determining whether a document is worth reading in full, and can be used to identify informative sections in a text. The presented approach to SORE shows promising results w.r.t. automatically filtering out a large proportion of irrelevant, incorrect, or uninformative OIE extractions. Considering the poor quality of OIE extractions, however, we propose presenting a reader with the sentences that entail the filtered OIE extractions. Furthermore, SORE provides a method to collect data for annotation and training of scientific OIE systems.

6 Conclusions

We introduce the task of Semi-Open Relation Extraction (SORE) on scientific texts and the Focused Open Biological Information Extraction (FOBIE) dataset. We adapt off-the-shelf IE systems to show that SORE is feasible, and that our approach is worth improving upon – both in terms of performance, as well as reducing the system’s complexity. A strong scientific IE system is used as a baseline, and its output is used to filter the relations found by a state-of-the-art OIE system.

OIE from scientific text is a hard task. The large number of errors that we find in OIE extractions from scientific texts render them near-useless to downstream computing tasks. A human reader may, nevertheless, find many incorrect extractions informative. An issue for humans is the sheer amount of OIE extractions and the high proportion of uninformative extractions. We show that our approach

TRADE-OFF relations	
<i>Trade-off arguments</i>	<i>Argument modifiers</i>
sleep	
cognitive abilities	
energy conservation	
memory retention	(the keeping of memory over prolonged periods of time)
memory consolidation	(in bats) (without a food reward) (shift from short- to long-term memory) (using torpor)
Examples filtered OIE extractions	
<i>(A memory; is normally formed; after repeated learning events; sleep enhances this process)</i>	
<i>(learning; is associated; with a food reward)</i>	
<i>(Sleep deprivation; has; negative effects on both memory consolidation)</i>	
<i>(torpor; has; a negative influence on memory consolidation)</i>	
<i>(digestion; prevents; the bats; from falling into torpor quickly)</i>	
<i>(torpor; indeed affects; learning abilities)</i>	

Table 7: SORE extractions from a scientific biology text (Ruczy ski et al., 2014). The TRADE-OFF relations are extracted by a narrow IE system trained on FOBIE. These relations capture the central theme and concepts of the text, and are used to filter the extractions that an OIE system outputs for the same document. The resulting extractions can support discerning the relevance of scientific documents.

to SORE reduces the number of OIE extractions by 65%, while increasing the relative amount of informative extractions by 5.75%. As a result, SORE improves the ability for a reader to quickly skim through the remaining extractions, or sentences that they are sourced from, and analyze how central concepts are related in a scientific text.

The presented approach is currently limited to the domain of biology and the use of trade-off relations, but we expect that central relations can be identified for other scientific domains that enable SORE. We show that creating a dataset for narrow RE can be done relatively cheaply by re-annotating the output of a simple RBS. Similarly, SORE may aid the collection of a dataset for scientific OIE.

Acknowledgments

The authors would like to gratefully acknowledge the financial support of the Engineering and Physical Sciences Research Council (EPSRC) Centre for Doctoral Training in Embedded Intelligence under grant reference EP/L014998/1 and the EPSRC Innovation Placement fund. We also thank Ben Trevett for proof-reading the document.

References

- Dominique Adriaens. 2019. [Evomimetics: the biomimetic design thinking 2.0](#). In *Bioinspiration, Biomimetics, and Bioreplication IX*, volume 1096509, page 8. SPIE.
- Anurag A Agrawal, Jeffrey K Conner, and Sergio Rasmann. 2010. Tradeoffs and Negative Correlations in Evolutionary Ecology Why Are We Interested in Tradeoffs? In *Evolution since Darwin: the first 150 years*, chapter 10, pages 243–268.
- Brian S Alper, Jason A Hand, Susan G Elliott, Scott Kinkade, Michael J Hauan, Daniel K Onion, and Bernard M Sklar. 2004. [How much effort is needed to keep up with the literature relevant for primary care?](#) *Journal of the Medical Library Association : JMLA*, 92(4):429–37.
- Waleed Ammar, Matthew Peters, Chandra Bhagavathula, and Russell Power. 2018. [The AI2 system at SemEval-2017 Task 10 \(ScienceIE\): semi-supervised end-to-end entity and relation extraction](#). pages 592–596.
- Isabelle Augenstein, Mrinal Das, Sebastian Riedel, Lakshmi Vikraman, and Andrew McCallum. 2017. [SemEval 2017 Task 10: ScienceIE - Extracting Keyphrases and Relations from Scientific Publications](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 546–555, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Michele Banko, M Cafarella, Stephen Soderland, M J Broadhead, and Oren Etzioni. 2007. [Open information extraction from the web](#). *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 2670–2676.
- Robert Bossy, Louise Deléger, Estelle Chaix, Mouhamadou Ba, and Claire Nédellec. 2019. [Bacteria Biotope at BioNLP Open Shared Tasks 2019](#). In *Proceedings of The 5th Workshop on BioNLP Open Shared Tasks*, pages 121–131, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Razvan C Bunescu and Raymond J Mooney. 2005. [A shortest path dependency kernel for relation extraction](#). *Proceedings of the Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, (October):724–731.
- Stephen S. O. Burgess, Jarmila Pittermann, and Todd E. Dawson. 2006. [Hydraulic efficiency and safety of branch xylem increases with height in Sequoia sempervirens \(D. Don\) crowns](#). *Plant, Cell and Environment*, 29(2):229–239.
- Gemma Carr, Daniel P. Loucks, and Günter Blöschl. 2018. [Gaining insight into interdisciplinary research and education programmes: A framework for evaluation](#). *Research Policy*, 47(1):35–48.
- Estelle Chaix, Bertrand Dubreucq, Abdelhak Fatihi, Dialekti Valsamou, Robert Bossy, Mouhamadou Ba, Louise Deléger, Pierre Zweigenbaum, Philippe Bessières, Loïc Lepiniec, and Claire Nédellec. 2016. [Overview of the Regulatory Network of Plant Seed Development \(SeeDev\) Task at the BioNLP Shared Task 2016](#). pages 1–11.
- Yubo Chen, Liheng Xu, Kang Liu, Daojian Zeng, and Jun Zhao. 2015. [Event Extraction via Dynamic Multi-Pooling Convolutional Neural Networks](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, pages 167–176.
- Janara Christensen, Mausam, Stephen Soderland, and Oren Etzioni. 2011. [An analysis of open information extraction based on semantic role labeling](#). In *Proceedings of the sixth international conference on Knowledge capture - K-CAP '11*, page 113.
- K. Bretonnel Cohen, Karin Verspoor, Karèn Fort, Christopher Funk, Michael Bada, Martha Palmer, and Lawrence E. Hunter. 2017. [The Colorado Richly Annotated Full Text \(CRAFT\) Corpus: Multi-Model Annotation in the Biomedical Domain](#). In *Handbook of Linguistic Annotation*, pages 1379–1394. Springer Netherlands.
- Daniel Cook, Babatunde A. Ogunnaike, and Rajanikanth Vadigepalli. 2015. [Systems analysis of non-parenchymal cell modulation of liver repair across multiple regeneration modes](#). *BMC Systems Biology*, 9(1):71.
- Louise Deléger, Robert Bossy, Estelle Chaix, Mouhamadou Ba, Arnaud Ferré, Philippe Bessières, and Claire Nédellec. 2016. [Overview of the Bacteria Biotope Task at BioNLP Shared Task 2016](#). pages 12–22.
- Khalid El-Arini and Carlos Guestrin. 2011. [Beyond keyword search](#). In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '11*, page 439, New York, New York, USA. ACM Press.
- Oren Etzioni, Michael Cafarella, Doug Downey, Ana Maria Popescu, Tal Shaked, Stephen Soderland, Daniel S. Weld, and Alexander Yates. 2005. [Un-supervised named-entity extraction from the Web: An experimental study](#). *Artificial Intelligence*, 165(1):91–134.
- Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011. [Identifying relations for Open Information Extraction](#). In *EMNLP 2011 - Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, pages 1535–1545.
- Kata Gábor, Davide Buscaldi, Anne-Kathrin Schumann, Behrang Qasemizadeh, Hafa Zargayouna, and Thierry Charnois. 2018. [SemEval-2018 Task 7: Semantic Relation Extraction and Classification in Scientific Papers](#). In *SemEval*, pages 679–688.

- Luis Galárraga, Jeremy Heitz, Kevin Murphy, and Fabian M. Suchanek. 2014. [Canonicalizing Open Knowledge Bases](#). In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management - CIKM '14*, pages 1679–1688.
- Theodore Garland. 2014. [Trade-offs](#). *Current Biology*, 24(2):R60–R61.
- Paul Groth, Michael Lauruhn, Antony Scerri, and Ron Daniel. 2018. [Open Information Extraction on Scientific Text: An Evaluation](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, page 3414–3423. Association for Computational Linguistics.
- Luheng He, Kenton Lee, Omer Levy, and Luke Zettlemoyer. 2018. [Jointly Predicting Predicates and Arguments in Neural Semantic Role Labeling](#). pages 364–369.
- Lynette Hirschman, Alexander Yeh, Christian Blaschke, and Alfonso Valencia. 2005. [Overview of BioCreAtIvE: critical assessment of information extraction for biology](#). *BMC Bioinformatics*, 6(Suppl 1):S1.
- Rezarta Islamaj Doğan, Sun Kim, Andrew Chatr-aryamontri, Chih-Hsuan Wei, Donald C Comeau, Rui Antunes, Sérgio Matos, Qingyu Chen, Aparna Elangovan, Nagesh C Panyam, Karin Verspoor, Hongfang Liu, Yanshan Wang, Zhuang Liu, Berna Altinel, Zehra Melce Hüsünbeyi, Arzucan Özgür, Aris Fergadis, Chen-Kai Wang, Hong-Jie Dai, Tung Tran, Ramakanth Kavuluru, Ling Luo, Albert Steppi, Jinfeng Zhang, Jinchuan Qu, and Zhiyong Lu. 2019. [Overview of the BioCreative VI Precision Medicine Track: mining protein interactions and mutations for precision medicine](#). *Database*, 2019.
- Robin Jia, Cliff Wong, and Hoifung Poon. 2019. [Document-Level N-ary Relation Extraction with Multiscale Representation Learning](#). In *NAACL HLT*, pages 3693–3704.
- Xiaotian Jiang, Quan Wang, Peng Li, and Bin Wang. 2016. [Relation Extraction with Multi-instance Multi-label Convolutional Neural Networks](#). *Coling*, (89):1471–1480.
- J. D. Kim, T Ohta, Y Tateisi, and J Tsujii. 2003. [GENIA corpus - A semantically annotated corpus for bio-textmining](#). In *Bioinformatics*, volume 19, pages 180–182.
- Jin-Dong Kim, Tomoko Ohta, Sampo Pyysalo, Yoshinobu Kano, and Jun'ichi Tsujii. 2009. [Overview of BioNLP'09 shared task on event extraction](#). In *Proceedings of the Workshop on BioNLP Shared Task - BioNLP '09*, page 1, Morristown, NJ, USA. Association for Computational Linguistics.
- Olga Kononova, Haoyan Huo, Tanjin He, Wenhao Sun, Ziqin Rong, Tiago Botari, Vahe Tshitoyan, and Gerbrand Ceder. 2019. [Text-mined dataset of inorganic materials synthesis recipes](#). *Sci Data*, 6(203).
- Martin Krallinger, Obdulia Rabal, Saber A Akhondi, Martín Pérez Pérez, Jesús Santamaría, Gael Pérez Rodríguez, Georgios Tsatsaronis, Ander Intxaurre, José Antonio López, Umesh Nandal, Erin Van Buel, Akileshwari Chandrasekhar, Marleen Rodenburg, Astrid Laegreid, Marius Doornenbal, Julen Oyarzabal, Analia Lourenço, and Alfonso Valencia. 2017. [Overview of the BioCreative VI chemical-protein interaction Track](#). In *Proceedings of BioCreative VI workshop*, pages 141–146.
- Ruben Kruiper, Jessica Chen-Burger, and Marc P Y Desmulliez. 2016. [Computer-aided biomimetics](#). In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 9793, pages 131–143.
- Ruben Kruiper, Julian Vincent, Eitan Abraham, Rupert Soar, Ioannis Konstas, Jessica Chen-Burger, and Marc Desmulliez. 2018. [Towards a Design Process for Computer-Aided Biomimetics](#). *Biomimetics*, 3(3):14.
- Chi Yun Kuo, Gary B. Gillis, and Duncan J. Irschick. 2011. [Loading effects on jump performance in green anole lizards, Anolis carolinensis](#). *Journal of Experimental Biology*, 214(12):2073–2079.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. [End-to-end Neural Coreference Resolution](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. [Neural Relation Extraction with Selective Attention over Instances](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2124–2133.
- Chunyang Liu, Wenbo Sun, Wenhan Chao, and Wanxiang Che. 2013. [Convolution neural network for relation extraction](#). *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8347 LNAI(PART 2):231–242.
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018a. [Multi-Task Identification of Entities, Relations, and Coreference for Scientific Knowledge Graph Construction](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3219–3232, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Yi Luan, Mari Ostendorf, and Hannaneh Hajishirzi. 2018b. [The UWNLP system at SemEval-2018 Task 7: Neural Relation Extraction Model with Selectively Incorporated Concept Embeddings](#). In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 788–792, Stroudsburg, PA, USA. Association for Computational Linguistics.

- Mausam. 2016. [Open Information Extraction Systems and Downstream Applications](#). In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 4074–4077.
- Makoto Miwa and Mohit Bansal. 2016. [End-to-End Relation Extraction using LSTMs on Sequences and Tree Structures](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1105–1116, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Claire Nédellec, Robert Bossy, J-D Kim, J J Kim, and Tomoko Ohta. 2013. [Overview of BioNLP shared task 2013](#). In *BioNLP Shared Task 2013 Workshop*, pages 1–7.
- Thien Huu Nguyen and Ralph Grishman. 2015. [Relation Extraction: Perspective from Convolutional Neural Networks](#). In *NAACL-HLT*, pages 39–48.
- Christina Niklaus, Matthias Cetto, André Freitas, and Siegfried Handschuh. 2018. [A Survey on Open Information Extraction](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3866–3878. Association for Computational Linguistics.
- Elisabeth Pain. 2016. [How to keep up with the scientific literature](#). *Science*.
- Harinder Pal and Mausam. 2016. [Demonyms and Compound Relational Nouns in Nominal Open IE](#). In *Proceedings of the 5th Workshop on Automated Knowledge Base Construction*, page 35–39.
- Nanyun Peng, Hoifung Poon, Chris Quirk, Kristina Toutanova, and Wen-tau Yih. 2017. [Cross-Sentence N -ary Relation Extraction with Graph LSTMs](#). *Transactions of the Association for Computational Linguistics*, 5:101–115.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [Glove: Global Vectors for Word Representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep Contextualized Word Representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 19, pages 2227–2237, Stroudsburg, PA, USA. Association for Computational Linguistics.
- I. Ruczy ski, T. M. A. Clarin, and B. M. Siemers. 2014. [Do greater mouse-eared bats experience a trade-off between energy conservation and learning?](#) *Journal of Experimental Biology*, 217(22):4043–4048.
- Swarnadeep Saha and Mausam. [Open Information Extraction from Conjunctive Sentences](#). In *To be published in COLING 2018*.
- Swarnadeep Saha, Harinder Pal, and Mausam. 2017. [Bootstrapping for Numerical Open IE](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 317–323.
- Cicero Nogueira dos Santos, Bing Xiang, and Bowen Zhou. 2015. [Classifying Relations by Ranking with Convolutional Neural Networks](#). pages 626–634.
- 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.
- Pontus Stenetorp, Sampo Pyysalo, Goran Topic, Tomoko Ohta, Sophia Ananiadou, and Jun’ichi Tsujii. 2012. [BRAT: a Web-based Tool for NLP-Assisted Text Annotation](#). In *13th Conference of the European Chapter of the ACL*, Figure 1, pages 102–107.
- S.S. Vattam and A.K. Goel. 2013. [An Information Foraging Model of Interactive Analogical Retrieval](#). In *Proceedings of the Annual Meeting of the Cognitive Science Society*, pages 3651–3656.
- Julian F. V. Vincent. 2014. [An Ontology of Biomimetics](#). In Ashok K Goel, Daniel A McAdams, and Robert B. Stone, editors, *Biologically Inspired Design*, pages 269–285. Springer London, London.
- Julian F. V. Vincent. 2016. [The trade-off: a central concept for biomimetics](#). *Bioinspired, Biomimetic and Nanobiomaterials*, 6(2):67–76.
- Jonas Warringer and Anders Blomberg. 2006. [Evolutionary constraints on yeast protein size](#). *BMC Evolutionary Biology*, 6(1):61.
- Dian Yu, Lifu Huang, and Heng Ji. 2017. [Open Relation Extraction and Grounding](#). *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 1:854–864.
- Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. [Distant Supervision for Relation Extraction via Piecewise Convolutional Neural Networks](#). In *EMNLP*, pages 1753–1762.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. [Relation Classification via Convolutional Deep Neural Network](#). In *COLING*, pages 2335–2344.
- Deyu Zhou, Dayou Zhong, and Yulan He. 2014. [Biomedical Relation Extraction: From Binary to Complex](#). *Computational and Mathematical Methods in Medicine*, pages 1–18.