

A Scientific Information Extraction Dataset for Nature Inspired Engineering

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

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

*Corresponding author: rk22@hw.ac.uk

Abstract

Nature has inspired various ground-breaking technological developments in applications ranging from robotics to aerospace engineering and the manufacturing of medical devices. However, accessing the information captured in scientific biology texts is a time-consuming and hard task that requires domain-specific knowledge. Improving access for outsiders can help interdisciplinary research like Nature Inspired Engineering. This paper describes a dataset of 1,500 manually-annotated sentences that express domain-independent relations between central concepts in a scientific biology text, such as trade-offs and correlations. The arguments of these relations can be Multi Word Expressions and have been annotated with modifying phrases to form non-projective graphs. The dataset allows for training and evaluating Relation Extraction algorithms that aim for coarse-grained typing of scientific biological documents, enabling a high-level filter for engineers.

Keywords: Scientific Information Extraction, Relation Extraction, Biomimetics, Trade-Offs

1. Introduction

Discovering relevant scientific literature is a hard and time-consuming task that often requires domain expertise (Alper et al., 2004; El-Arini and Guestrin, 2011; Pain, 2016). This difficulty can lead to barriers in inter-disciplinary fields of study, where an expert in a target domain is often a novice in the source domain (Carr et al., 2018). A specific example of an inter-disciplinary field of study is Nature Inspired Engineering, also known as biomimetics (Kruiper et al., 2016). A well known example of biomimetics is Velcro® (Mestral, 1955), an ‘*adhesion*’ method that is inspired by nature - specifically the burrs of burdock plants. In biomimetics an engineer is interested in learning from nature to solve a problem at hand (Hoeller et al., 2016). It has been shown that identifying, selecting and understanding relevant biological information is challenging for engineers (Vattam and Goel, 2013a). Information Extraction (IE) systems can partially ease these challenges by capturing the central concepts and relations in a text (Augenstein et al., 2017; Gábor et al., 2018).

A major issue in inter-disciplinary research is that search terms from a source domain may not retrieve relevant results in a target domain. As an example, in the biological domain the term ‘*bleaching*’ can refer to a separation process between the retina and opsin in vertebrate eyes (Grimm et al., 2000). A non-biologist may expect ‘*bleaching*’ to refer to *cleaning*, *sterilizing*, or *whitening* (Nagel, 2014). As a result, terminology from the engineering domain does not always provide a good starting point for the identification of relevant biological information (Fayemi et al., 2015; Kruiper et al., 2017). A more appropriate approach that allows for cross-domain search, without relying so much on domain-specific semantics, focuses on TRADE-OFF relations (Vincent, 2016; Kruiper et al., 2018). In biology, trade-offs express how fitness is constrained by functional requirements that are mutually exclusive, e.g., because they share a limited resource (Agrawal et al., 2010). Figure 1 provides an example of a trade-off between ‘*safety*’ and ‘*efficiency*’. Trade-offs are important drivers for adapta-

Broad patterns of increasing ‘*safety*’ and decreasing ‘*efficiency*’ of xylem with increasing aridity are evident, and there is some evidence demonstrating **safety-efficiency trade-offs** within conifer species or individuals.

Figure 1: Example of a TRADE-OFF relation between domain-independent abstract concepts, denoted as ARG1 and ARG2, extracted from (Burgess et al., 2006).

tion and speciation, and underpin much of the research in various biology sub-domains (Stearns, 1989; Garland, 2014; Ferenci, 2016). TRADE-OFF relations are of interest to biomimetics because they are able to capture the problem space of a full-text biology document in highly abstract terms. Analyzing these trade-offs enables a domain-independent document classification. Crucially, trade-offs can direct an engineer to the mechanisms that biological systems employ to manipulate such problem spaces.

In this work we present the Focused Open Biology Information Extraction (FOBIE) dataset. FOBIE comprises 1500 manually-annotated sentences taken from full-text scientific biology documents. The dataset enables training and evaluation of Relation Extraction (RE) tools that extract trade-offs from scientific biology texts. Each sentence is annotated with a non-projective graph of n -ary relations between trigger words, argument phrases and modifying phrases. The relations are domain-independent and the argument phrases contain concepts that are central to the text. Most of these phrases are not available in standard knowledge bases, so they cannot be learned through distant supervision. This paper provides a comparison of FOBIE with regards to existing datasets for scientific Information Extraction (IE), a description of the collection and annotation processes, and the results of a strong scientific IE baseline.

2. Related work

2.1. Computer-Aided biomimetics

Biomimetics is an engineering problem solving process during which one draws on analogous biological solutions. While biomimetics has led to the development of leap-frog innovations, the process remains adventitious. A major issue is that engineers know little biology or few details of plants or animals. Therefore, engineers have trouble identifying, selecting and understanding relevant biological information (Vattam and Goel, 2011; Vattam and Goel, 2013b). To overcome these challenges various computational tools have been developed, so-called Computer-Aided biomimetics (CAB) tools. Their biggest limitation is that they rely heavily on cross-domain mapping of information such as functional similarities, e.g., (Vandevenne et al., 2012; Vandevenne et al., 2016; Shu and Cheong, 2010; Cheong and Shu, 2014; Rugaber et al., 2016; Zhao et al., 2018). However, automatically identifying cross-domain relational similarities, requires substantial reasoning beyond what is currently computationally feasible.

Two caveats are lexical variation and polysemy of concepts across domains (Augenstein et al., 2017). Thirdly, the biological and engineering domains are fundamentally different. Biological systems develop through evolution and natural selection, while engineering is based on conscious decision-making to meet functional requirements (Vincent et al., 2006). Consider again the example of ‘bleaching’: the purely biological term does not carry a notion of teleology, in contrast to the non-biologist interpretations. The active verbs that are associated to *engineering functions* are not always used in biological texts (Kaiser et al., 2012; Kaiser et al., 2014). As a result, each function of interest requires a separate classifier (Glier, 2013) or its own set of extraction rules that work for a specific domain of texts (Etzioni, 2007; Christensen et al., 2011). Training data for each function would accordingly require annotation by a domain expert (Luan et al., 2018).

Instead, this work aims to extract trade-off relations that are central to many biological texts and capture information at an abstraction level that is domain-independent (Vincent, 2016). These trade-offs, and the abstract concepts captured in their arguments, provide an initial filter for information retrieval and support for within-domain search for biological information (Kruiper et al., 2018).

2.2. Scientific Information Extraction

Information Extraction (IE) from scientific text can (1) improve access to scientific information, beyond the possibilities of standard search engines (Gábor et al., 2018; Gupta and Manning, 2011), (2) provide valuable insight into research areas (Tsai et al., 2013; Luan et al., 2018), (3) enable to quickly learn facts on unknown concepts, as extractions can provide a summary view for readers (Mausam, 2016), and (4) augment existing Knowledge Bases (KB) from unlabeled text (Quirk and Poon, 2017). However, annotating scientific text is non-trivial, as it requires domain-specific knowledge from experts. Large-scale crowdsourcing (Tratz, 2019) or human-in-the-loop (He et al., 2016) efforts can be unreliable for such tasks.

To alleviate the manual labeling of FOBIE by a domain expert, we developed a simple Rule-Based System (RBS). Like traditional Relation Extraction (RE) systems this RBS relies on matching specific words to find a TRADE-OFF or similar relation (Sarawagi, 2007). Similar to Open IE (OIE) systems it relies on unlexicalized grammatical structures, e.g., syntactic patterns, to determine the argument phrases (Etzioni et al., 2011). It is expected that this enables flexibility of extracting trade-offs from a variety of domains (Etzioni et al., 2005; Banko et al., 2007). However, the handcrafted or learned syntactic patterns rely on how well input sentences are handled (Léchelle and Langlais, 2018; Niklaus et al., 2018). Sources of issues may include the length and complexity of sentences, the use of pronouns as subjects, the use of abbreviations as adjectives, handling prepositional phrases, and dealing with co-reference (Schneider et al., 2017; Groth et al., 2018). Machine learning approaches can improve recall for the long tail of patterns that will have to be identified (Mausam, 2016).

By manually relabeling the output of our RBS system we rectify errors in argument boundaries and determine the correct relation label. Only few similar datasets exist for scientific IE:

SCIENCEIE Semeval 2017 task 10 introduced the SCIENCEIE dataset, consisting of 500 paragraphs taken from full-text scientific documents in the domains of Computer Science, Material Science and Physics journal articles (Augenstein et al., 2017). It contains annotations of keyphrases that are classified as *materials*, *processes* or *tasks*. Furthermore, hyponym- and synonym-relations between the keyphrases are captured.

SEMEVAL 2018 The manually annotated Semeval 2018 task 7 dataset contains 6 relations types that are noted to occur regularly in scientific abstracts; *usage*, *result*, *model*, *part-whole*, *topic* and *comparison* (Gábor et al., 2018). It contains 500 abstracts from the domain of Computational Linguistics and draws on the ACL RD-TEC 2.0 corpus (Qasemizadeh and Schumann, 2016) for entity annotation; *technology and method*, *tool and library*, *language resource*, *language resource product*, *measures and measurements*, *models* and *other*. Augenstein et al. (2017) found that the SCIENCEIE dataset contains a significantly higher proportion of long keyphrases in comparison to the ACL RD-TEC 2.0 corpus. This is likely due to the different characteristics of sentences taken from abstracts and those in the main body.

SCIERC The SCIERC dataset consists of 500 abstracts taken from Artificial Intelligence conference and workshop proceedings (Luan et al., 2018). It extends the entity and relation types of the Semeval 2018 task 7 dataset; the relation types are *used-for*, *feature-of*, *hyponym-of*, *part-of*, *compare*, *conjunction* and *coreference*, and the entity types are *task*, *method*, *evaluation metric*, *material*, *other scientific terms*, and *generic terms* that include anaphora and cataphora that are annotated for coreference resolution.

In this work we introduce the FOBIE dataset that targets sentences taken from full-text scientific publications in the Biology domain. Table 1 provides an overview of the sizes




Figure 2: Two examples of sentences that express a trade-off. The indicator denotes the trigger word that connects two trade-off arguments - ARG1 and ARG2. ARGMOD denotes a modifying phrase.

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

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

of FOBIE, SCIERC, SCIENCEIE and the Semeval 2018 dataset. Similar to SCIENCEIE the focus lies on extracting keyphrases, rather than entities. Different from the previously described datasets, the relation arguments are not classified into a type, also see section 3.2. FOBIE enables the training of IE systems that extract both a narrow set of relations between argument phrases and an unbounded set of modifying relations.

3. The FOBIE dataset

We collected a corpus of 10,000 open-access papers from the Journal of Experimental Biology (JEB) and three BioMed Central (BMC) journals: Biology, Systems Biology and Evolutionary Biology. We only retained the abstract, introduction, results, discussion and conclusion section. We used spaCy¹ to split the texts into sentences and tokens, and computed dependency parses².

3.1. Rule-Based System

A set of 200 papers was analyzed to determine the types of verbatim expressions that indicate a TRADE-OFF relation (Kruiper et al., 2018). The paper subjects ranged from cell biology to bio-mechanics and trigger words include *association*, *balance*, *conflict*, *correlation*, *compromise*, *interaction*, *interplay*, *optimization*, *ratio*, *relationship* and *trade-off*. We stem these trigger words and scan the sentences of the documents in our corpus for their presence. If a trigger word is found, we traverse the dependency parse tree to identify sub-trees that comprise argument phrases.

The output of the RBS system indicates that it has difficulty dealing with negation, and that complex long sentences indeed lead to errors in extractions. Figure 2 shows two sentences that express a *trade-off*, where each sentence deals differently with phrase attachment and argument boundaries. Variance in argument and modifier arrangement com-

PLICATE the definition of extraction rules that rely on a dependency parse tree. Therefore, an accurate RBS for the extraction of trade-offs from scientific text requires detailed extraction rules and exceptions. Figure 3 shows two examples of partial dependency parses taken from sentences that are near the average length of 38 tokens. The main source of errors is due to the extraction rules themselves struggling with the large variety of verbatim trade-off expressions (e.g., ‘*negative correlation*’, ‘*balance*’), rather than the quality of the dependency parser output. Improving the extraction rules of the simple RBS is expected to be a time-consuming and complex task. Very often the presence of a trigger word does not automatically lead to the presence of a trade-off. Determining whether the sentence expresses a trade-off requires additional reasoning and often domain-specific knowledge. Considering the large variety of syntactic structures that *do* indicate a trade-off, as well as the required expertise in biology, a RBS is not suited for our task. Instead, the output of our simple RBS is used to speed up manual annotation.

3.2. Dataset annotation

Out of the 10,000 documents, we retain only the sentences that were annotated with a TRADE-OFF relation by the RBS. Using the BRAT³ interface a biology expert manually corrected argument boundaries and relation types (Stenetorp et al., 2012). During manual annotation we correct the relation label for trigger words, handle negation and identify the boundaries of argument and modifier phrases. We annotate binary relations that constitute non-projective graphs of one or more *n*-ary relations in a sentence. Each binary relation is a triple $\langle \text{governor}, \text{relation}, \text{dependent} \rangle$ where:

- *governor* is either a trigger word or a modifying phrase.
- *relation* indicates the type of relation – TRADE-OFF, ARGUMENT-MODIFIER or NOT-A-TRADE-OFF.
- *dependent* is an argument phrase.

Three relation types were used: TRADE-OFF, ARGUMENT-MODIFIER and NOT-A-TRADE-OFF. The latter relation is used to indicate that a trigger word does not express a trade-off. We retain these annotations because their expressions are syntactically similar and, therefore, can provide useful

¹<https://spacy.io/>

²spaCy parse accuracy on the Penn Treebank corpus is 94.48.

³<https://brat.nlplab.org/>

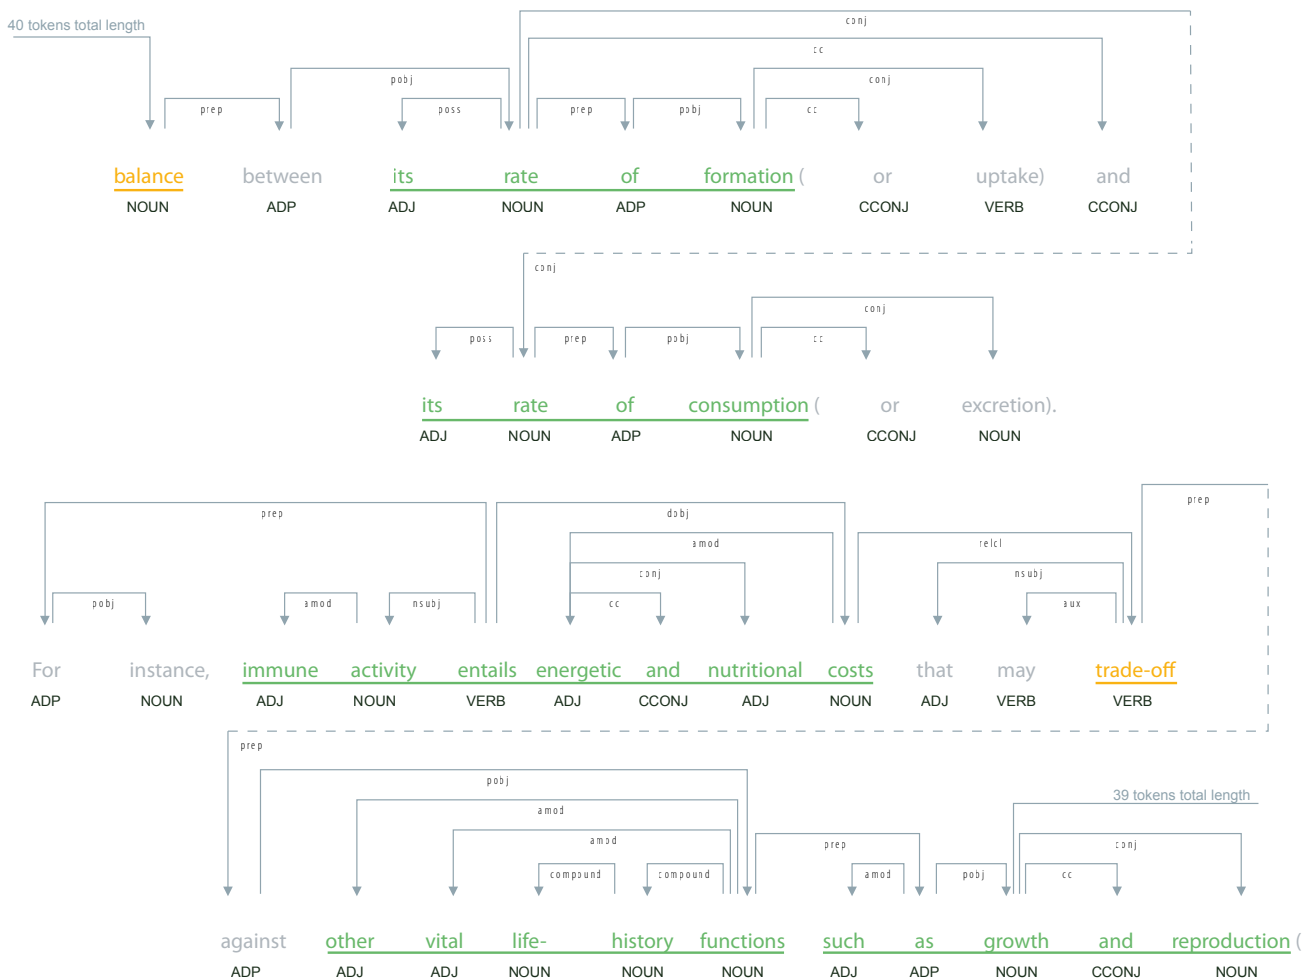


Figure 3: Dependency parses of two sentences that are near the average length of 37.77 tokens. The TRADE-OFF indicators and arguments have been underlined as in Figure 2, as well as the single argument modifier "such as [...]". Defining extraction rules that can traverse a large variety of dependency trees correctly is a time-consuming and hard task that requires expertise in both the linguistics and biology domains.

training signal as negative samples. Furthermore, the trigger words to express these NOT-A-TRADEOFF relations may be contiguous with the trigger words that do express a TRADE-OFF. An example is a sentence that expresses a correlation, which often denotes a positive correlation between two arguments while a TRADE-OFF entails a negative correlation.

Many argument phrases are found to be nested and can be broken down into an argument and a modifier, e.g., 'ontogenetic trajectories' modifies 'tolerance' and 'resistance' in Figure 2. We do not indicate specific types of modification, such as temporal expressions or phrases that indicate a location, used in Semantic Role Labeling (SRL). The reason is that trigger words can be either nouns or verbs, e.g., 'balance/NOUN' and 'trade-off/VERB' in Figure 3, hence the classes of modifier labels found in common SRL frameworks such as Propbank (Kingsbury and Palmer, 2002) are not always appropriate. Instead, as a general rule we include the words that indicate the type of modifier during annotation of modifying phrases.

We adopt the heuristic that prepositional phrases (PPs) heading a coordinating clause, are treated as modifying phrases when they apply to word-level arguments of a

trade-off, as in the top example in Figure 2. In the case of coordination of PPs that contain arguments of a trade-off, we consider each PP as a whole argument, as in the bottom example in Figure 2. Similarly, when nested phrases in an argument can be distinctly separated by punctuation, such phrases are treated as an argument-modifier pair. We do not annotate phrases that modify the relation directly, e.g., 'throughout plant development' modifies 'correlation' in the top of Figure 2. Regarding the direction of relations, the trigger words and modifying phrases are treated as governors of a relation.

Nevertheless, determining exact rules for the boundaries of arguments was found to be challenging. As an example, consider the second TRADE-OFF relation in Figure 3. In this sentence, 'immune activity' is a 'life-history function' that trades off with 'growth' and 'reproduction', two other examples of 'vital life-history functions'. Specifically, 'energetic and nutritional costs' need to be shared between these 'life-history functions'. The first argument 'immune activity [...] costs' is treated as a single phrase following the heuristics above. This verb phrase could also be split up, leaving the informative phrase 'entails energetic and nutritional costs' outside of the annotation scope. In such am-

	Train	Dev	Test	Total
# Sentences	1248	150	150	1548
Avg. sent. length	37.28	37.78	37.82	37.77
% of sents \geq 25 tokens	-	-	-	79.26
Relations:				
- TRADE-OFF	639	54	72	765
- NOT-A-TRADE-OFF	2004	258	240	2502
- ARG-MODIFIER	1247	142	132	1521
Triggers	1292	155	153	1600
Arguments	3435	401	398	4234
Spans	5137	596	576	6309
Unique spans	2701			
Unique triggers	41			
Max triggers/sent	2			
Max spans/sent	7			
Spans w/ multiple relations	2075			
# single-word arguments	498 (11.8%)			
Avg. tokens per argument	3.44			

Table 2: The aggregated statistics for FOBIE.

		Relations			Boundaries		
		P	R	F1	P	R	F1
RBS	<i>dev</i>	-	-	-	45.17	35.84	39.97
	<i>test</i>	-	-	-	44.31	35.32	39.31
SCIIE	<i>dev</i>	66.59	66.88	66.74	84.91	80.94	82.87
	<i>test</i>	68.53	65.48	66.97	86.76	79.39	82.91

Table 3: Overview of results on the development (dev) and test set of FOBIE for the RBS and SCIIE. The relations columns refers to the Relation Extraction (RE) setting. The RBS does not classify relations and is, therefore, not evaluated on this task. Boundaries refers to determining the correct boundaries of arguments and trigger words.

biguous cases, we opted for always annotating the longest possible phrase spans, in order to capture more comprehensively the information in the sentence.

A random sample of 250 sentences (16.1%), which were annotated by using the RBS, was re-annotated by a second domain-expert. The inter-annotator agreement Cohen k for both relations and span boundaries was found to be 92.93.

3.3. Dataset description

Table 2 provides an overview of the statistics on FOBIE. The percentage of singleton keyphrases in FOBIE is only 12%, a considerable difference with regard to SCIENCEIE (31%) and ACL RD-TEC 2.0 (83%) (Augenstein and Søgaard, 2017). The number of single token entities in SCIERC is (31%), with an average token length of 2.36 in comparison to 3.44 in FOBIE. The reason is that arguments of trade-offs are often phrases or Multi-Word Expressions (MWE), such as ‘immune response’.

In FOBIE only 21% of the sentences is shorter than 25 tokens, while in SCIERC the average sentence length is 24.31 tokens. The much longer sentences in FOBIE is influenced by the presence of citations, but in general the sentences in full-text documents are longer. The final dataset comprises 1548 single sentences taken from 1292 unique documents and is randomly split into 1248 training, 150 development and 150 test instances. There is no source document overlap between the training, development and test set.

	TRADE-OFF			NOT-A-TRADE-OFF			ARG-MODIFIER		
	P	R	F1	P	R	F1	P	R	F1
<i>dev</i>	70.91	72.22	71.56	72.12	75.19	73.62	54.01	50.34	52.11
<i>test</i>	80.00	83.33	81.63	72.41	70.00	71.19	54.10	48.18	50.97

Table 4: Overview of SCIIE results for each of the relations in FOBIE. Note that the amount of gold relations for the development (dev) and test sets are not equal.

4. Baseline Evaluation

We train a state-of-the-art scientific RE system on FOBIE called SCIIE (Luan et al., 2018). The SCIIE model is developed for the SCIERC dataset to jointly extract the entities, relations and coreference annotations (Luan et al., 2018). Furthermore, the model achieves state-of-the-art performance on the SCIENCEIE dataset, which includes the tasks of span identification and keyphrase extraction. For FOBIE we only train SCIIE on entities and relations, where the entities are our annotated keyphrases; trigger spans and argument spans.

SCIIE is a span-based IE model that takes as input an unlabeled sentence, e.g., no POS-tags or dependency labels have to be provided. All sequences of consecutive words in the sentence – up to a given length – are considered as a candidate span, while only a few represent the gold annotated keyphrases. Each token is represented as the concatenation of its word, character and ELMo embedding (Peters et al., 2018). Each candidate span is represented as a concatenation of the bi-LSTM outputs for the first token in the span and the last token in the span, as well as an attention mechanism over the span that is thought to represent the syntactic head of span (Lee et al., 2017). Span classification determines whether a span is actually a span, where a dummy non-span class is assigned to incorrect candidate spans. Pairs of candidate spans are fed to the relation classifier, where a dummy non-relation class is introduced. This multi-task setup enables weight sharing that improves performance on both tasks. See (Luan et al., 2018) for more details on SCIIE.

We compare RBS and SCIIE on detecting the correct boundaries of trigger words and argument phrases, see Table 3. As expected the neural approach outperforms the simple RBS by a large margin. We also evaluate how well SCIIE performs on FOBIE with regards to extracting relations. The Relation Extraction (RE) setting is evaluated as a joint task – the presence of spans has to be predicted correctly, as well as the relation types that may hold between them. We do not change the default hyper parameters, apart from setting the maximum span length to 15 tokens. This is required to deal with the longer keyphrases in FOBIE.

Table 4 provides RE results per relation type. The model performs better on the test set with regards to extracting TRADE-OFF relations. Due to the relatively large amount of NOT-A-TRADE-OFF (52.26%) and ARGUMENT-MODIFIER (31.77%) relations this difference becomes negligible in the overall RE setting, also see Table 2. The model performs notably worse on ARGUMENT-MODIFIER relations. This is likely the result of the large syntactic variety in argument modification.

5. Conclusions

We presented FOBIE, the first scientific IE dataset that focuses on the domain of biomimetics. FOBIE supports the extraction of TRADE-OFFS and syntactically similar relations from scientific biological texts. These relations enable cross-domain discovery of relevant scientific literature during biomimetics. However, the manners in which trade-off are expressed, and their arguments modified, varies a lot. This large syntactic variation and the long sentences in scientific text decrease the accuracy of extraction rules. Therefore, Rule-Based Systems (RBS) may not be able to deal with the long tail of patterns. The alternative machine learning approach requires time-consuming manual annotation of datasets. Combining the two approaches, a simple RBS can speed up the annotation of a small dataset by a domain expert considerably. The use of pretrained embeddings provides reliable flexibility with regards to syntactic variance.

The manual annotation of FOBIE instances enables the training of neural Relation Extraction systems. The size of FOBIE is comparable to existing dataset for scientific Information Extraction. Unlike existing datasets the keyphrases in FOBIE are not classified into entity types. We make FOBIE and the annotation guidelines publicly available at <https://github.com/rubenkruiper/FOBIE>.

Acknowledgments

The authors would like to thank the reviewers for their useful comments, and 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.

6. Bibliographical References

- Agrawal, A. A., Conner, J. K., and Rasman, S. (2010). Tradeoff s 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.
- Alper, B. S., Hand, J. A., Elliott, S. G., Kinkade, S., Huan, M. J., Onion, D. K., and Sklar, B. M. (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, 10.
- Augenstein, I. and Søgaard, A. (2017). Multi-Task Learning of Keyphrase Boundary Classification. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 341–346, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Augenstein, I., Das, M., Riedel, S., Vikraman, L., and McCallum, A. (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.
- Banko, M., Cafarella, M., Soderland, S., Broadhead, M. J., and Etzioni, O. (2007). Open information extraction from the web. *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 2670–2676.
- Burgess, S. S. O., Pittermann, J., and Dawson, T. E. (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, 2.
- Carr, G., Loucks, D. P., and Blöschl, G. (2018). Gaining insight into interdisciplinary research and education programmes: A framework for evaluation. *Research Policy*, 47(1):35–48, 2.
- Cheong, H. and Shu, L. (2014). Retrieving Causally Related Functions From Natural-Language Text for Biomimetic Design. *Journal of Mechanical Design*, 136(8):081008.
- Christensen, J., Mausam, Soderland, S., and Etzioni, O. (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.
- El-Arini, K. and Guestrin, C. (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.
- Etzioni, O., Cafarella, M., Downey, D., Popescu, A. M., Shaked, T., Soderland, S., Weld, D. S., and Yates, A. (2005). Unsupervised named-entity extraction from the Web: An experimental study. *Artificial Intelligence*, 165(1):91–134, 6.
- Etzioni, O., Fader, A., Christensen, J., Soderland, S., and Mausam. (2011). Open information extraction: The second generation. In *IJCAI International Joint Conference on Artificial Intelligence*, volume 11, pages 3–10.
- Etzioni, O. (2007). Machine reading of web text. In *Proceedings of the 4th international conference on Knowledge capture - K-CAP '07*, pages 1–4, New York, New York, USA. ACM Press.
- Fayemi, P.-E., Maranzana, N., Aoussat, A., Chekchak, T., and Bersano, G. (2015). Modeling Biological Systems to facilitate their Selection during a Bio-Inspired Design Process. In *ICED15*, pages 225–234, Milan, Italy.
- Ferenci, T. (2016). Trade-off Mechanisms Shaping the Diversity of Bacteria. *Trends in Microbiology*, 24(3):209–223, 3.
- Gábor, K., Buscaldi, D., Schumann, A.-K., Qasemizadeh, B., Zargayouna, H., and Charnois, T. (2018). SemEval-2018 Task 7: Semantic Relation Extraction and Classification in Scientific Papers. In *SemEval*, pages 679–688.
- Garland, T. (2014). Trade-offs. *Current Biology*, 24(2):R60–R61.
- Glier, M. W. (2013). *Machine Learning based classification of textual stimuli to promote ideation in BioInspired Design*. Ph.D. thesis, Texas A&M University.
- Grimm, C., Reme, C. E., Rol, P. O., and Williams, T. P. (2000). Blue light’s effects on rhodopsin: Photoreversal

- of bleaching in living rat eyes. *Investigative Ophthalmology and Visual Science*, 41(12):3984–3990, 11.
- Groth, P., Lauruhn, M., Scerri, A., and Daniel, R. (2018). Open Information Extraction on Scientific Text: An Evaluation. In *Proceedings of the 27th International Conference on Computational Linguistics*, page 3414â3423. Association for Computational Linguistics, 2.
- Gupta, S. and Manning, C. D. (2011). Analyzing the Dynamics of Research by Extracting Key Aspects of Scientific Papers. In *Proceedings of 5th international joint conference on natural language processing*, pages 1–9.
- He, L., Michael, J., Lewis, M., and Zettlemoyer, L. (2016). Human-in-the-Loop Parsing. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2337–2342, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Hoeller, N., Farnsworth, M., Jacobs, S., Chirazi, J., Mead, T., Goel, A., and Salustri, F. (2016). A Systems View of Bio-Inspiration: Bridging the Gaps. *INSIGHT*, 19(1):36–40.
- Kaiser, M. K., Hashemi Farzaneh, H., and Lindemann, U. (2012). An approach to support searching for biomimetic solutions based on system characteristics and its environmental interactions. *Proceedings of International Design Conference, DESIGN, DS 70:969–978*.
- Kaiser, M. K., Hashemi Farzaneh, H., and Lindemann, U. (2014). Bioscrabble - The role of different types of search terms when searching for biological inspiration in biological research articles. In *Proceedings of International Design Conference, DESIGN*, volume 2014-Janua, pages 241–250, Dubrovnik, Croatia.
- Kingsbury, P. and Palmer, M. (2002). From treebank to propbank. In *Proceedings of the 3rd International Conference on Language Resources and Evaluation, LREC 2002*, pages 1989–1993.
- Kruiper, R., Chen-Burger, J., and Desmulliez, M. P. Y. (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.
- Kruiper, R., Vincent, J. F. V., Chen-Burger, J., and Desmulliez, M. P. Y. (2017). Towards Identifying Biological Research Articles in Computer-Aided Biomimetics. In *Living Machines 2017.*, pages 242–254.
- Kruiper, R., Vincent, J., Abraham, E., Soar, R., Konstas, I., Chen-Burger, J., and Desmulliez, M. (2018). Towards a Design Process for Computer-Aided Biomimetics. *Biomimetics*, 3(3):14, 6.
- Léchelle, W. and Langlais, P. (2018). Revisiting the Task of Scoring Open IE Relations. In *LREC 2018 - 11th International Conference on Language Resources and Evaluation*, pages 2052–2058.
- Lee, K., He, L., Lewis, M., and Zettlemoyer, L. (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, 7. Association for Computational Linguistics.
- Luan, Y., He, L., Ostendorf, M., and Hajjishirzi, H. (2018). 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, 8. 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.
- Mestral, G. D. (1955). Velvet type fabric and method of producing same. *US Patent Number 2,717,437*, pages 1–3.
- Nagel, J. K. (2014). A Thesaurus for Bioinspired Engineering Design. In *Biologically Inspired Design: Computational Methods and Tools*, number January 2014, pages 245–267. Springer London, London.
- Niklaus, C., Cetto, M., Freitas, A., and Handschuh, S. (2018). A Survey on Open Information Extraction. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3866–3878. Association for Computational Linguistics.
- Pain, E. (2016). How to keep up with the scientific literature. *Science*, 11.
- Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. (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, 7. Association for Computational Linguistics.
- Qasemzadeh, B. and Schumann, A. K. (2016). The ACL RD-TEC 2.0: A language resource for evaluating term extraction and entity recognition methods. In *Proceedings of the 10th International Conference on Language Resources and Evaluation, LREC 2016*, pages 1862–1868.
- Quirk, C. and Poon, H. (2017). Distant Supervision for Relation Extraction beyond the Sentence Boundary. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, volume 1, pages 1171–1182, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Rugaber, S., Bhati, S., Goswami, V., Spiliopoulou, E., Azad, S., Koushik, S., Kulkarni, R., Kumble, M., Sarathy, S., and Goel, A. (2016). Knowledge Extraction and Annotation for Cross-Domain Textual Case-Based Reasoning in Biologically Inspired Design. In *International Conference on Case-Based Reasoning*, number November, pages 342–355. Springer.
- Sarawagi, S. (2007). Information Extraction. *Foundations and Trends® in Databases*, 1(3):261–377.
- Schneider, R., Oberhauser, T., Klatt, T., Gers, F. A., and Löser, A. (2017). Analysing Errors of Open Information Extraction Systems. In *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Sys-*

- tems, pages 11–18, 7.
- Shu, L. H. and Cheong, H. (2010). A natural-language approach to biomimetic design. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing: AIEDAM*, 24(4):507–519.
- Stearns, S. C. (1989). Trade-Offs in Life-History Evolution. *Functional Ecology*, 3(3):259–268.
- Stenetorp, P., Pyysalo, S., Topic, G., Ohta, T., Ananiadou, S., and Tsujii, J. (2012). BRAT: a Web-based Tool for NLP-Assisted Text Annotation. In *13th Conference of the European Chapter of the ACL*, number Figure 1, pages 102–107.
- Tratz, S. (2019). Dependency Tree Annotation with Mechanical Turk. *Proceedings of the First Workshop on Aggregating and Analysing Crowdsourced Annotations for NLP*, pages 1–5.
- Tsai, C.-T., Kundu, G., and Roth, D. (2013). Concept-based analysis of scientific literature. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management - CIKM '13*, pages 1733–1738, New York, New York, USA. ACM Press.
- Vandevenne, D., Verhaegen, P.-A., Dewulf, S., and Dufloy, J. R. (2012). Automated classification into the biomimicry taxonomy. In *Proceedings of International Design Conference, DESIGN*, volume DS 70, pages 1161–1166.
- Vandevenne, D., Pieters, T., and Dufloy, J. R. (2016). Enhancing novelty with knowledge-based support for Biologically-Inspired Design. In *Design Studies*.
- Vattam, S. and Goel, A. (2011). Foraging for Inspiration: Understanding and Supporting the Online Information Seeking Practices of Biologically Inspired Designers. *Proceedings of the ASME 2011 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference IDETC/CIE 2011*, pages 177–186.
- Vattam, S. and Goel, A. (2013a). An Information Foraging Model of Interactive Analogical Retrieval. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, pages 3651–3656.
- Vattam, S. S. and Goel, A. K. (2013b). Biological Solutions for Engineering Problems: A Study in Cross-Domain Textual Case-Based Reasoning. In *International Conference on Case Based Reasoning*, volume 7969 LNAI, pages 343–357.
- Vincent, J. F. V., Bogatyreva, O. A., Bogatyrev, N. R., Bowyer, A., and Pahl, A.-K. (2006). Biomimetics: its practice and theory. *Journal of the Royal Society*, 3(April):471–482.
- Vincent, J. F. V. (2016). The trade-off: a central concept for biomimetics. *Bioinspired, Biomimetic and Nanobiomaterials*, 6(2):67–76, 6.
- Zhao, Y., Baldini, I., Sattigeri, P., Padhi, I., Lee, Y. K., and Smith, E. (2018). Data Driven Techniques for Organizing Scientific Articles Relevant to Biomimicry. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society - AIES '18*, pages 347–353, New York, New York, USA. ACM Press.