







Coreference Resolution in Research Papers from Multiple Domains

Arthur Brack¹, Daniel Uwe Müller², Anett Hoppe¹,
and Ralph Ewerth^{1,2}

¹ Research Group Visual Analytics, TIB – Leibniz Information Centre for Science and Technology, Hannover, Germany

{[arthur.brack](mailto:arthur.brack@tib.eu), [anett.hoppe](mailto:anett.hoppe@tib.eu), [ralph.ewerth](mailto:ralph.ewerth@tib.eu)}@tib.eu

² L3S Research Center, Leibniz University, Hannover, Germany

Abstract. Coreference resolution is essential for automatic text understanding to facilitate high-level information retrieval tasks such as text summarisation or question answering. Previous work indicates that the performance of state-of-the-art approaches (e.g. based on BERT) noticeably declines when applied to scientific papers. In this paper, we investigate the task of coreference resolution in research papers and subsequent knowledge graph population. We present the following contributions: (1) We annotate a corpus for coreference resolution that comprises 10 different scientific disciplines from Science, Technology, and Medicine (STM); (2) We propose transfer learning for automatic coreference resolution in research papers; (3) We analyse the impact of coreference resolution on knowledge graph (KG) population; (4) We release a research KG that is automatically populated from 55,485 papers in 10 STM domains. Comprehensive experiments show the usefulness of the proposed approach. Our transfer learning approach considerably outperforms state-of-the-art baselines on our corpus with an F1 score of 61.4 (+11.0), while the evaluation against a gold standard KG shows that coreference resolution improves the quality of the populated KG significantly with an F1 score of 63.5 (+21.8).

Keywords: Coreference resolution · Information extraction · Knowledge graph population · Scholarly communication

1 Introduction

Current research papers are generally published in form of PDF files. This makes them hard to handle for retrieval systems, since their content is hidden in human-but not machine-interpretable text. In consequence, current academic search engines are not able to adequately support researchers in their day-to-day tasks. This is further aggravated by the exploding number of published articles [5].

Approaches to automatically structure research papers are thus an active area of research. *Coreference resolution* is the task of identifying mentions in a text

which refer to the same entity or concept. It is an essential step for automatic text understanding and facilitates down-stream tasks such as text summarisation or question answering. For instance, the text ‘*Coreference resolution is... It is used for question answering...*’, has two coreferent mentions ‘*Coreference resolution*’ and ‘*It*’. This allows us to extract the fact \langle coreference resolution, used_for, question answering \rangle .

Current methods for coreference resolution based on deep learning achieve quite impressive results (e.g. an F1 score of 79.6 for the OntoNotes 5.0 dataset [21]) in the general domain, that is data from phone conversations, news, magazines, etc. But results of previous work indicate [11, 23, 34, 44] that general coreference resolution systems perform poorly on scientific text. This is presumably caused by the specific terminology and phrasing used in a scientific domain. Some other studies state that annotating scientific text is costly since it demands certain expertise in the article’s domain [2, 6, 20]. Most corpora for research papers cover only a single domain (e.g. biomedicine [11], artificial intelligence [27]) and are thus limited to these domains. As a result, the annotated corpora are relatively small and overall only a few domains are covered. Datasets for the general domain are usually much larger, but they have not been exploited yet by approaches for coreference resolution in research papers.

Coreference resolution is also one of the main steps in the KG population pipeline [28, 39]. However, to date it is not clear, to which extent (a) coreference resolution can help to reduce the number of scientific concepts in the populated KG, and (b) how coreference resolution influences the quality of the populated KG. Besides, a KG comprising multiple scientific domains has not been populated yet.

In this paper, we address the task of coreference resolution in research papers and subsequent knowledge graph population. Our contributions can be summarised as follows: (1) First, we annotate a corpus for coreference resolution that consists of 110 abstracts from 10 domains from Science, Technology, and Medicine. The systematic annotation resulted in a substantial inter-coder agreement (0.68 κ). We provide and compare baseline results for this dataset by evaluating five different state-of-the-art approaches. Our experimental results confirm that state-of-the-art coreference approaches do not perform well on research papers. (2) Consequently, we propose sequential transfer learning for coreference resolution in research papers. This approach utilises our corpus by fine-tuning a model that is pre-trained on a large corpus from the general domain [37]. Experimental results show that our approach significantly outperforms the best state-of-the-art baseline (F1 score of 61.4, i.e. +11.0). (3) We investigate the impact of coreference resolution on automatic KG population. To evaluate the quality of various KG population strategies, we (i) compile a gold standard KG from our annotated corpus that contains scientific concepts referenced by mentions from text, and (ii) present a procedure to evaluate the clustering results of mentions. (4) We release (i) an automatically populated KG from 55,485 abstracts of the 10 STM domains and (ii) a gold KG (Test-STM-KG) from the annotated STM-corpus. Experimental results show that coreference resolution has only a small

impact on the number of concepts in a populated KG, but it helps to improve the quality of the KG significantly: the population with coreference resolution yields an F1 score of 63.5 evaluated against the gold KG (+21.8 F1). We release the data corpora and the source code to facilitate further research: <https://github.com/arthurbra/stm-coref>.

The remainder of the paper is organised as follows: Sect. 2 summarises related work on coreference resolution. Section 3 describes the annotation procedure and the characteristics of the corpus, and our proposed approaches for coreference resolution, KG population, and KG evaluation. The experimental setup and results are reported in Sect. 4 and 5, while Sect. 6 concludes the paper and outlines areas of future work.

2 Related Work

2.1 Approaches for Coreference Resolution

For a given document, the task of coreference resolution is (a) to extract mentions of scientific concepts, and (b) to cluster those mentions that refer to the same concept. Recent approaches mostly rely on supervised learning and can be categorised into three groups [32]: (1) Mention-pair models [33,45] are binary classifiers that determine whether two mentions are coreferent or not. (2) Entity-mention models [9,41] determine whether a mention is coreferent to a preceding *cluster*. A cluster has more expressive features compared to a mention in mention-pair models. (3) Ranking-based models [12,25,31] simultaneously rank all candidate antecedents (i.e. preceding mention candidates). This enables the model to identify the most probable antecedent.

Lee et al. [25,26] propose an end-to-end neural coreference resolution model. It is a ranking-based model that jointly recognises mentions and clusters. Therefore, the model considers all spans in the text as possible mentions and learns distributions over possible antecedents for each mention. For computational efficiency, candidate spans and antecedents are pruned during training and inference. Joshi et al. [22] enhance Lee et al.’s model with BERT-based word embeddings [14], while Ma et al. [30] improve the model with better attention mechanisms and loss functions.

Furthermore, several approaches propose multi-task learning, such that related tasks may benefit from knowledge in other tasks to achieve better prediction accuracy: Luan et al. [27,49] train a model on three tasks (coreference resolution, entity and relation extraction) using one dataset of research papers. Sanh et al. [43] introduce a multi-task model that is trained on four tasks (mention detection, coreference resolution, entity and relation extraction) using two different datasets in the general domain.

Results of some previous studies [11,23,34,44] revealed that general coreference systems do not work well in the biomedical domain due to the lack of domain knowledge. For instance, on Colorado Richly Annotated Full Text (CRAFT) corpus [11] a coreference resolution system for the news domain achieves only 14.0 F1 (−32.0).

To the best of our knowledge, a transfer learning approach from the general to the scientific domain has not been proposed for coreference resolution yet.

2.2 Corpora for Coreference Resolution in Research Papers

For the general domain, multiple datasets exist for coreference resolution, e.g. Message Understanding Conference (MUC-7) [1], Automatic Content Extraction (ACE05) [15], or OntoNotes 5.0 [37]. The OntoNotes 5.0 dataset [37] is the largest one and is used in many benchmark experiments for coreference resolution systems [22, 25, 30].

Various annotated datasets for coreference resolution exist also for research papers: CRAFT corpus [11] covers 97 papers from biomedicine. The corpus of Schäfer et al. [44] contains 266 papers from computational linguistics and language technology. Chaimongkol et al. [7] annotated a corpus of 284 papers from four subdisciplines in computer science. The SciERC corpus [27] comprises 500 abstracts from the artificial intelligence domain and features annotations for scientific concepts and relations. It was used to generate an artificial intelligence (AI) knowledge graph [13]. Furthermore, several datasets exist for scientific concept extraction [2, 6, 27, 40] and relation extraction [2, 20, 27] that cover various scientific domains.

To the best of our knowledge, a corpus for coreference resolution that comprises a broad range of scientific domains is not available yet.

3 Coreference Resolution in Research Papers

As the discussion of related work reveals, existing corpora for coreference resolution in scientific papers normally cover only a single domain, and coreference resolution approaches do not perform well on scholarly texts. To address these issues, we systematically annotate a corpus with coreferences in abstracts from 10 different science domains. Current approaches for coreference resolution in research papers do not exploit existing annotated datasets from the general domain, which are usually much larger than in the scientific domain. We propose a sequential transfer learning approach that takes advantage from large, annotated datasets. Finally, to the best of our knowledge, the impact of (a) coreference resolution and (b) cross-domain collapsing of mentions to scientific concepts on KG population with multiple science domains has not been investigated yet. Consequently, we present an evaluation procedure for the clustering aspect in the KG population pipeline.

In the sequel, we describe our annotated corpus, our transfer learning approach for coreference resolution, and an evaluation procedure for clustering in KG population.

3.1 Corpus for Coreference Resolution in 10 STM Domains

In this section, we describe the STM corpus [6], which we used as the basis for the annotation, our annotation process, and the characteristics of the resulting corpus.

STM Corpus: The STM corpus [6] comprises 110 articles from 10 domains in Science, Technology and Medicine, namely Agriculture (Agr), Astronomy (Ast), Biology (Bio), Chemistry (Che), Computer Science (CS), Earth Science (ES), Engineering (Eng), Materials Science (MS), Mathematics (Mat), and Medicine (Med). It contains annotated mentions of scientific concepts in abstracts with four domain-independent concept types, namely *Process*, *Method*, *Material*, and *Data*. These concept mentions were later linked to entities in Wikipedia and Wikidata [16]. The 110 articles (11 per domain) were taken from the OA-STM corpus [17] of Elsevier Labs.

We build upon related work and extend the STM corpus with coreference annotations. In particular, we (1) annotate coreference links between existing scientific concept mentions in abstracts using the BRAT annotation tool [46], and (2) annotate further mentions, i.e. pronouns and noun phrases consisting of multiple consecutive mentions.

Annotation Process: Other studies have shown that non-expert annotations are viable for the scientific domain [6, 8, 19, 44, 47], and they are less costly than domain-expert annotations. Therefore, we also annotate the corpus with non-domain experts, i.e. by two students in computer science. Furthermore, we follow mostly the annotation procedure of the STM corpus [6], which consists of the following three phases:

Table 1. Per-domain and overall inter-annotator agreement (Cohen’s κ and MUC) for coreference resolution annotation in our STM corpus.

	<i>Mat</i>	<i>Med</i>	<i>Ast</i>	<i>CS</i>	<i>Bio</i>	<i>Agr</i>	<i>ES</i>	<i>Eng</i>	<i>Che</i>	<i>MS</i>	<i>Overall</i>
κ	0.84	0.80	0.78	0.72	0.70	0.66	0.61	0.58	0.56	0.52	0.68
MUC	0.83	0.69	0.78	0.73	0.70	0.72	0.61	0.66	0.56	0.63	0.69

Table 2. Characteristics of the annotated STM corpus with 110 abstracts per concept type in terms of number of scientific concept mentions, number of coreferent mentions, number of coreference clusters and singleton clusters, and the number of overall clusters. MIXED denotes clusters consisting of mentions with different concept types, NONE denotes coreference mentions and clusters without a scientific concept mention.

	Data	Material	Method	Process	MIXED	NONE	Total
# mentions	1,658	2,099	258	2,112	0	0	6,127
# coreferent mentions	351	910	101	510	0	705	2,577
# coreference clusters	153	339	30	198	50	138	908
# singleton clusters	1,307	1,189	157	1,602	0	0	4,255
# overall clusters	1,460	1,528	187	1,800	50	138	5,163

Table 3. Characteristics of the STM corpus per domain (11 abstracts per domain).

	Agr	Ast	Bio	Che	CS	ES	Eng	MS	Mat	Med	Total
# mentions	741	791	649	553	483	698	741	574	297	600	6,127
# coreferent mentions	276	365	275	282	181	241	318	256	124	259	2,577
# coreference clusters	106	120	98	90	67	93	117	87	48	82	908
# singleton clusters	520	549	443	384	339	525	503	371	210	411	4,255
# clusters	626	669	541	474	406	618	620	458	258	493	5,163

1. *Pre-annotation*: This phase aims at developing annotation guidelines through trial annotations. We adapted the comprehensive annotation guidelines of the OntoNotes 5.0 dataset [38], which were developed for the general domain, to research papers. In particular, we provide briefer and simpler descriptions with examples from the scientific domain. Within three iterations both annotators labelled independently 10, 9 and 7 abstracts (i.e. 26 abstracts), respectively. After each iteration the annotators discussed the outcome and refined the annotation guidelines.
2. *Independent Annotation*: After the annotation guidelines were finalised, both annotators independently re-annotated the previously annotated abstracts and 24 additional abstracts. The final inter-coder agreement was measured on the 50 abstracts (5 per domain) using Cohen’s κ [10,24] and MUC [48]. As shown in Table 1, we achieve a substantial agreement with 0.68 κ and 0.69 MUC.
3. *Consolidation*: Finally, the remaining 60 abstracts were annotated by one annotator and the annotation results of this author were used as the gold standard corpus.

Corpus Characteristics: Table 2 shows the characteristics of the resulting corpus broken down per concept type, while they are listed per domain in Table 3. The original corpus has in total 6,127 mentions. 2,577 mentions were annotated as coreferent resulting in 908 coreference clusters. Thus, each coreference cluster contains on average 2.84 mentions, while *Method* clusters contain the most (3.4 mentions) and *Data* clusters the least (2.3 mentions). Furthermore, 705 mentions were annotated additionally (referred to as NONE) since they represent pronouns (422 mentions) or noun phrases consisting of multiple consecutive original mentions (283 mentions) such as ‘... [[A], [B], and [C] [treatments]]... [These treatments]...’. Fifty clusters (5%) contain mentions with different concept types (referred to as MIXED) due to disagreements between the annotators of the original concept mentions, and the annotators of coreferences. For instance, non-coreferent mentions were annotated as coreferent, or coreferent mentions have different concept types. Finally, 138 clusters (15%) do not have a concept type (NONE) since they form clusters which are not coreferent with the original concept mentions.

3.2 Transfer Learning for Coreference Resolution

We suggest sequential transfer learning [42] for coreference resolution in research papers. Therefore, we fine-tune a model pre-trained on a large (source) dataset to our (target) dataset. As the source dataset, we use the English portion of the OntoNotes 5.0 dataset [37], since it is a broad corpus that consists of 3,493 documents with telephone conversations, magazine and news articles, web data, broadcast conversations, and the New Testament. Besides, our annotation guidelines were adapted from OntoNotes 5.0.

For the model, we utilise *BERT for Coreference Resolution (BFCR)* [22] with *SpanBERT* [21] word embeddings. This model achieves state-of-the-art results on the OntoNotes dataset [21]. Another advantage is the availability of the pre-trained model and the source code. The BFCR model improves Lee et al.’s approach [26] by replacing the LSTM encoder with the SpanBERT transformer-encoder. SpanBERT [21] has different training objectives than BERT [14] to better represent spans of text.

3.3 Cross-Domain Research Knowledge Graph Population

Let $d \in D$ be an abstract, $M(d) = \{m_1, \dots, m_h\}$ the mentions of scientific concepts in d , and $c_d(m_i) \subseteq M(d)$ the corresponding coreference cluster for mention m_i in d . If mention m_s is not coreferent with other mentions in d , then $c_d(m_s) = \{m_s\}$ is a singleton cluster. The set of all clusters is denoted by C . An equivalence relation *collapsable* $\subseteq C \times C$ defines if two clusters can be collapsed, i.e. if the clusters refer to the same scientific concept. To create the set of all concepts E , we build the quotient set for the set of clusters C with respect to the relation *collapsable*:

$$C := \{c_d(m) | d \in D, m \in M(d)\} \quad (1)$$

$$[c] := \{x \in C | \text{collapsable}(c, x)\} \quad (2)$$

$$E := \{[c] | c \in C\} \quad (3)$$

Now, we can construct the KG: for each paper $d \in D$ and for each scientific concept $e \in E$ we create a node in the KG. The scientific concept type of e is the most frequent concept type of all mentions in e . Then, for each mention $m \in M(d)$ we create a ‘mentions’ link between the paper and the corresponding scientific concept $[m] \in E$.

Cross-Domain vs. In-Domain Collapsing: One commonly used approach to define the *collapsable* relation is to treat two clusters as equivalent, if and only if the ‘label’ of the clusters is the same. The label of a cluster is the longest mention in the cluster normalised by (a) lower-casing, (b) removing articles, possessives and demonstratives, (c) resolving acronyms, and (d) lemmatisation using WordNet [18] to transform plural forms to singular. Other studies [13, 27] used a similar label function for KG population.

However, a research KG that comprises multiple scientific disciplines has not been populated yet. Thus, it is not clear whether it is feasible to collapse clusters across domains. Usually, terms within a scientific domain are unambiguous, but some terms can have different meanings across scientific disciplines (e.g. “neural network” in *CS* and *Med*). Thus, we investigate both cross-domain and in-domain collapsing strategies.

Knowledge Graph Population Approach: We populate a research KG with research papers from multiple scientific domains, i.e. 55,485 abstracts of Elsevier with CC-BY licence from the 10 investigated domains. First, we extract (a) concept mentions from the abstracts using the scientific concept extractor of the STM-corpus [6], and (b) clusters within the abstracts with our transfer learning coreference model. Then, those mention clusters, which contain solely mentions recognised by the coreference resolution model and not by the scientific concept extraction model, are dropped, since the coreference resolution model does not recognise the concept type of the mentions. Finally, the remaining clusters serve for the population of the KG as described above.

3.4 Evaluation Procedure of Clustering in KG Population

One common approach to evaluate the quality of a populated KG is to annotate a (random) subset of statements by humans as true or false and to calculate precision and recall [13, 50]. To evaluate recall, small collections of ground-truth capturing *all* knowledge is necessary, that are usually difficult to obtain [50]. To the best of our knowledge, a common approach to evaluate the clustering aspect of the KG population pipeline does not exist yet. Thus, in the following, we present (1) an annotated test KG, and (2) metrics to evaluate clustering of mentions to concepts in KG population.

Test KG: To enable evaluation of KG population strategies, we compile a test KG, referred to as *Test-STM-KG*. For this purpose, we reuse the STEM-ECR corpus [16], in which 1,221 mentions of the STM corpus are linked to Wikipedia entities. First, we extract all annotated clusters of the STM corpus in which all mentions of the cluster uniquely refer to the same Wikipedia entity. Then, we collapse all clusters which refer to the same Wikipedia entity to concepts. Formally, the Test-STM-KG is a partition of mentions, where each part denotes a concept, i.e. a disjoint set of mentions. A mention is uniquely represented by the tuple (start offset, end offset, concept type, doc id).

Table 4 shows the characteristics of the compiled Test-STM-KG. It consists of 920 clusters, of which 711 are singleton clusters. These clusters were collapsed to 762 concepts, of which 31 concepts are used across multiple domains (referred to as MIX).

Evaluation Procedure: To evaluate the clustering result of a KG population strategy, we use the metrics of coreference resolution. The three popular metrics for coreference resolution are *MUC* [48], B^3 [3] and $CEAF_{e_{\phi_4}}$ [29]. Each of them

Table 4. Characteristics of the *Test-STM-KG*: number of concepts per concept type and per domain. MIX denotes the number of cross-domain concepts.

	Agr	Ast	Bio	CS	Che	ES	Eng	MS	Mat	Med	MIX	Total
Data	5	18	3	20	4	9	28	13	37	8	9	154
Material	27	35	30	20	26	52	32	30	9	40	7	308
Method	1	1	1	21	6	2	4	10	3	8	7	64
Process	17	12	21	34	13	33	20	25	15	38	8	236
Total	50	66	55	95	49	96	84	78	64	94	31	762

represents different evaluation aspects (see [36] for more details). To calculate these metrics, we treat the gold concepts (i.e. a partition of mentions) of the Test-STM-KG as the ‘key’ and the predicted concepts as the ‘response’. We report also the *CoNLL P/R/F1* scores, that is the averages of *MUC*’s, *B³*’s and *CEAF_{e_φ4}*’s respective precision (P), recall (R) and F1 scores. The CoNLL metrics were proposed for the conference on Computational Natural Language Learning (CoNLL) shared tasks on coreference resolution [36].

4 Experimental Setup

Here we describe our experimental setup for coreference resolution and KG population.

4.1 Automatic Coreference Resolution

We evaluate three different state-of-the-art architectures on the STM dataset: (I) *BERT for Coreference Resolution (BFCR)* [22] with *SpanBERT* [21] word embeddings (referred to as *BFCR.Span*), (II) BFCR with *SciBERT* [4] word embeddings (referred to as *BFCR.Sci*), and (III) *Scientific Information Extractor (SCIE)* [27] with ELMo [35] word embeddings (referred to as *SCIE*). The three architectures are evaluated in the following six approaches (#1–#6):

- *Pre-Trained Models*: We evaluate already pre-trained models on the test sets of the STM corpus, i.e. #1 *BFCR.Span* trained on the English portion of the OntoNotes dataset [38], and #2 *SCIE* trained on SciERC [27] from the AI domain.
- *Supervised Learning*: We train a model from scratch with the three architectures using the training data of the STM corpus and evaluate their performance with the test sets of STM: #3 *BFCR.Span*, #4 *BFCR.Sci*, and #5 *SCIE*.
- *Transfer Learning*: This is our proposed approach #6. We fine-tune all parameters of a pre-trained model on the English portion of the OntoNotes dataset [21] with the training data of our STM corpus. For that, we use the *BFCR.Span* architecture.

Table 5. Performance of the baseline approaches #1–#5 and our proposed transfer learning approach #6 on the test sets of the STM corpus across five-fold cross validation.

		Training data	<i>MUC</i>			B^3			<i>CEAFe_{ϕ_4}</i>			<i>CoNLL</i>		
			P	R	F1	P	R	F1	P	R	F1	P	R	F1
#1	BFCR_Span	OntoNotes	57.1	31.1	40.2	55.9	25.7	35.2	50.2	28.1	36.0	54.4	28.3	37.1
#2	SCIIE	SciERC	13.4	4.5	6.8	13.1	4.3	6.5	18.1	6.0	9.0	14.9	4.9	7.4
#3	BFCR_Span	STM	61.6	45.6	52.3	59.8	41.5	48.8	57.9	44.4	50.0	59.8	43.8	50.4
#4	BFCR_Sci	STM	61.9	40.2	48.6	59.7	36.1	44.9	61.7	36.9	46.0	61.1	37.7	46.5
#5	SCIIE	STM	60.3	45.2	51.6	57.6	41.7	48.3	56.6	43.6	49.1	58.1	43.5	49.7
#6	BFCR_Span	Onto→STM	64.5	63.5	63.9	61.0	60.0	60.4	60.5	59.6	60.0	62.0	61.0	61.4

Table 6. Per domain and overall CoNLL F1 results of the best baseline #3 and our transfer learning approach #6 on the STM corpus across five-fold cross validation.

		Training data	Agr	Ast	Bio	Che	CS	ES	Eng	MS	Mat	Med	Overall
#3	BFCR_Span		STM	48.0	50.5	52.2	49.0	59.1	39.6	52.8	47.6	42.5	51.0
#6	BFCR_Span	Onto→STM	62.8	61.1	57.5	56.3	74.9	57.5	59.8	52.1	55.7	62.1	61.4

Evaluation: We use the metrics *MUC* [48], B^3 [3], *CEAF e_{ϕ_4}* [29] and *CoNLL* [36] in compliance with other studies on coreference resolution [22, 25, 30]. To obtain robust results, we apply five-fold cross-validation, according to the data splits given by Brack et al. [6], and report averaged results. For each fold, the dataset is split into train/validation/test sets with 8/1/2 abstracts per domain, respectively, i.e. 80/10/20 abstracts. We reuse the original implementations and default hyperparameters of the above architectures. Hyperparameter-tuning of the best baseline approach #3 according to [22] confirmed that the default hyperparameters of *BFCR_Span* perform best on our corpus.

4.2 Evaluation of KG Population Strategies

We compare four KG population strategies: (1) cross-domain and (2) in-domain collapsing, as well as (3) cross-domain and (4) in-domain collapsing without coreference resolution. To evaluate cross-domain and in-domain collapsing, we take the gold clusters (i.e. mention clusters within the abstracts) of the Test-STM-KG and collapse them to concepts according to the respective strategy. When leaving out the coreference resolution step, we treat all mentions in the Test-STM-KG as singleton clusters and collapse them to concepts according to the respective strategy. Finally, we calculate the metrics as described in Sect. 3.4.

5 Results and Discussion

In this section, we discuss the experimental results for automatic coreference resolution and KG population.

5.1 Automatic Coreference Resolution

Table 5 shows the overall results of the six evaluated approaches and Table 6 the results per domain of the best baseline #3 and our approach #6. Our transfer learning approach #6 *BFCR_Span* from OntoNotes (Onto) [37] to STM significantly outperforms the best baseline approach #3 with an overall CoNLL F1 of 61.4 (+10.0) and a low standard deviation ± 1.5 across the five folds.

Table 7. CoNLL scores on the tests sets of the SciERC corpus [27] across 3 random restarts of the approaches: current state of the art of Luan et al., the best baseline approach (#3), and our transfer learning approach (#6). We report results using the whole and using only $\frac{1}{5}$ th of the training data of SciERC (referred to as $\frac{1}{5}$ SciERC).

		Training data	P	R	F1
Luan et al. [27]		SciERC	52.0	44.9	48.2
#3	<i>BFCR_Span</i>	SciERC	63.3	55.7	59.3
#6	<i>BFCR_Span</i>	OntoNotes→SciERC	63.9	57.1	60.1
#3	<i>BFCR_Span</i>	$\frac{1}{5}$ SciERC	63.1	39.1	47.1
#6	<i>BFCR_Span</i>	OntoNotes→$\frac{1}{5}$SciERC	52.8	56.7	54.2

The approaches #1 *BFCR_Span* pre-trained on OntoNotes [37], and #2 *SCIIE* pre-trained on SciERC [27] achieve a CoNLL F1 score of 37.1 and 7.4, respectively. These scores are quite low compared to the approaches #3–#6 that use training data of the STM corpus. This indicates that models pre-trained on existing datasets do not generalise sufficiently well for coreference resolution in research papers. Models trained only on the STM corpus (i.e. #3–#5) achieve better results. However, they have quite low recall scores indicating that the size of the training data might not be sufficient to enable the model to generalise well. SciBERT #4, although pre-trained on scientific texts, performs worse than SpanBERT #3. Presumably the reason is that SpanBERT has approximately 3 times more parameters than SciBERT. Our transfer learning approach #6 achieves the best results with quite balanced precision and recall scores.

Furthermore, to evaluate the effectiveness of our transfer learning approach, we compare the best baseline #3 and our transfer learning approach #6 also with the SciERC corpus [27]. The SciERC corpus comprises 500 abstracts from the AI domain. Since SciERC has around 5 times more training data than STM, we compare the approaches #3 and #6 also using only $\frac{1}{5}$ th of the training data in SciERC while keeping the original validation and test sets. It can be seen in Table 7 that our transfer learning approach #6 improves slightly the baseline result using the whole training data with 60.1 F1 (+0.8). When using only $\frac{1}{5}$ th of the training data, our transfer learning approach noticeably outperforms the baseline with 54.2 F1 (+7.1). Thus, our transfer learning approach can help significantly to improve the performance of coreference resolution in research papers with few labelled data.

5.2 Cross-Domain Research KG

In this subsection, we describe the characteristics of our populated KG and discuss the evaluation results of various KG population strategies.

Characteristics of the Research KG: Table 8 shows the characteristics of the populated KGs per domain. The resulting KGs with cross-domain and in-domain collapsing have more than 994,000 and 1.1 Mio. scientific concepts, respectively, obtained from 55,485 abstracts with more than 2,1 Mio. concept mentions and 726,000 coreferent mentions. *Ast* and *Bio* are the most represented domains, while *CS* and *Mat* are the most underrepresented.

Table 8. Characteristics of the populated research KGs per domain: (1) number of abstracts, number of extracted scientific concept mentions and coreferent mentions, (2) the number of scientific concepts for the KG with cross-domain collapsing, (3) in-domain collapsing, (4) cross-domain collapsing but without coreference resolution, and (5) in-domain collapsing but without coreference resolution. Reduction denotes the percentual reduction of mentions to scientific concepts and MIX the cross-domain concepts.

	Agr	Ast	Bio	CS	Che	ES	Eng	MS	Mat	Med	MIX	Total
# abstracts	7,731	15,053	11,109	1,216	1,234	2,352	3,049	2,258	665	10,818	–	55,485
# mentions	332,983	370,311	423,315	45,388	46,203	129,288	127,985	86,490	20,466	586,019	–	2,168,448
# coref. men.	108,579	120,942	143,292	17,674	14,059	40,974	42,654	25,820	8,510	203,884	–	726,388
Cross-domain collapsing												
KG concepts	138,342	173,027	177,043	20,474	21,298	62,674	55,494	39,211	9,275	227,690	70,044	994,572
- Data	27,132	64,537	32,946	5,380	5,124	19,542	17,053	10,629	2,982	66,473	19,715	271,513
- Material	69,534	45,296	83,627	6,242	10,154	24,322	19,689	17,276	2,406	68,141	20,812	367,499
- Method	2,992	8,819	6,135	2,001	1,055	1,776	2,953	1,605	685	9,363	1,627	39,011
- Process	38,684	54,375	54,335	6,851	4,965	17,034	15,799	9,701	3,202	83,713	27,890	316,549
Reduction	58%	53%	58%	55%	54%	52%	57%	55%	55%	61%	–	54%
In-domain collapsing												
KG concepts	180,135	197,605	229,201	30,736	32,191	81,584	78,417	55,358	14,567	278,686	–	1,178,480
Reduction	46%	47%	46%	32%	30%	37%	39%	36%	29%	52%	–	46%
Cross-domain collapsing without coreference resolution												
KG concepts	146,894	182,479	187,557	21,950	22,555	66,600	59,689	41,776	9,939	242,797	77,493	1,059,729
Reduction	56%	51%	56%	52%	51%	48%	53%	52%	51%	59%	–	51%
In-domain collapsing without coreference resolution												
KG concepts	184,218	199,894	234,399	31,525	32,937	83,445	80,476	56,690	14,911	284,547	–	1,203,042
Reduction	45%	46%	45%	31%	29%	35%	37%	34%	27%	51%	–	45%

Evaluation of KG Population Strategies: Next, we discuss the different KG population strategies. For each strategy, Table 8 reports the number of concepts in the populated KG and the percentage reduction of mentions to concepts, and in Table 9 the evaluation results of KGs against the Test-STM-KG.

Cross-Domain vs. In-Domain Collapsing: Cross-domain collapsing achieves a higher CoNLL F1 score of 64.8 than in-domain collapsing with a score of 63.5

(see Table 9). However, in-domain collapsing yields (as expected) a higher precision (CoNLL P 85.5), since some terms have different meanings across domains (e.g. *Measure_(mathematics)* vs. *Measurement* in <https://en.wikipedia.org>). Furthermore, the Test-STM-KG has only 31 cross-domain concepts due to its small size. Thus, we expect that cross-domain collapsing would yield worse results on a larger test set.

Furthermore, as shown in Table 8, cross-domain collapsing yields less concepts than in-domain collapsing (more than 994,000 versus 1.1 Mio. concepts). We can also observe that only 70,044 (7%) of the concepts are used across multiple domains. This indicates that every scientific domain mostly uses its own terminology. However, the concepts used across domains can have different meanings. Thus, when precision is more important than recall in downstream tasks, in-domain collapsing should be the preferred choice.

Effect of Coreference Resolution: Coreference resolution has only a small impact on the number of resulting concepts in a populated KG (see Table 8). However, as shown in Table 9, leaving out the coreference resolution step during KG population yields only low CoNLL F1 scores, i.e. 41.7 (−21.8) F1 and 43.5 (−21.3) F1. Thus, coreference resolution significantly improves the quality of a populated KG .

Table 9. Performance of the collapsing strategies evaluated against the *Test-STM-KG*: in-domain and cross-domain collapsing with and without coreference resolution.

	#concepts in KG	MUC			B ³			CEAF _{ϕ_4}			CoNLL		
		P	R	F1	P	R	F	P	R	F1	P	R	F1
In-domain collapsing	859	86.3	70.6	77.7	86.0	69.0	76.6	84.1	23.1	36.2	85.5	54.2	63.5
- Without coreferences	900	75.5	38.8	51.2	75.2	37.9	50.4	71.1	14.0	23.4	73.9	30.2	41.7
Cross-domain collapsing	837	85.0	73.0	78.5	84.5	72.1	77.8	84.7	24.6	38.1	84.7	56.6	64.8
- Without coreferences	876	73.5	41.0	52.6	72.2	15.5	25.5	72.2	15.5	25.5	73.0	32.4	43.5

Qualitative Analysis: We also inspected the top-five frequent domain-specific concepts in the populated KG (a list of these concepts can be found in our public repository). As far as we can judge with our computer science background, we consider the extracted top frequent concepts to be reasonable and useful for the domains. For instance, in *Ast*, the method ‘standard model’ is frequently mentioned, while in *CS* the process ‘cyber attack’ appears most often. The frequency of the top concepts differs significantly between the domains: In *Med*, *Ast*, *Eng*, *ES* and *Agr*, a top frequent concept is referenced 10.8, 10.2, 4.9, 3.8, and 3.1 times per 1000 abstracts, respectively. In *Che*, *MS*, *Mat*, *Bio*, and *CS*, a top frequent concept is referenced only by few abstracts (0.3, 0.4, 1.0, 1.4, and 2.3, respectively, per 1000 abstracts).

6 Conclusions

In this paper, we have investigated the task of coreference resolution in research papers across 10 different scientific disciplines. We have annotated a corpus that

comprises 110 abstracts with coreferences with a substantial inter-coder agreement. Our baseline results with current state-of-the-art approaches for coreference resolution demonstrate that current approaches perform poorly on our corpus. The proposed approach, which uses sequential transfer learning and exploits annotated datasets from the general domain, outperforms noticeably the state-of-the-art baselines. Thus, our transfer learning approach can help to reduce annotation costs for scientific papers, while obtaining high-quality results at the same time.

Furthermore, we have investigated the impact of coreference resolution on KG population. For this purpose, we have compiled a gold KG from our annotated corpus and propose an evaluation procedure for KG population strategies. We have demonstrated that coreference resolution has a small impact on the number of resulting concepts in the KG, but improved significantly its quality. Finally, we have generated a research KG from 55,485 abstracts of the 10 investigated domains. We show that every domain mostly uses its own terminology and that the populated KG contains useful concepts.

In future work, we plan to evaluate multi-task learning approaches, and to populate and evaluate a much larger research KG to get more insights in scientific language use.

References

1. Seventh Message Understanding Conference: Proceedings of a Conference Held in Fairfax, Virginia, USA, MUC 1998, 29 April–1 May 1998. ACL (1998). <https://www.aclweb.org/anthology/volumes/M98-1/>
2. Augenstein, I., Das, M., Riedel, S., Vikraman, L., McCallum, A.: Semeval 2017 task 10: Scienceie - extracting keyphrases and relations from scientific publications. In: Bethard, S., Carpuat, M., Apidianaki, M., Mohammad, S.M., Cer, D.M., Jurgens, D. (eds.) Proceedings of the 11th International Workshop on Semantic Evaluation, SemEval@ACL 2017, Vancouver, Canada, 3–4 Aug 2017, pp. 546–555. Association for Computational Linguistics (2017). <https://doi.org/10.18653/v1/S17-2091>
3. Bagga, A., Baldwin, B.: Algorithms for scoring coreference chains. In: The First International Conference on Language Resources and Evaluation Workshop on Linguistics Coreference, pp. 563–566 (1998)
4. Beltagy, I., Lo, K., Cohan, A.: SciBERT: a pretrained language model for scientific text. In: Inui, K., Jiang, J., Ng, V., Wan, X. (eds.) Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, 3–7 Nov 2019, pp. 3613–3618. Association for Computational Linguistics (2019). <https://doi.org/10.18653/v1/D19-1371>
5. Bornmann, L., Mutz, R.: Growth rates of modern science: a bibliometric analysis based on the number of publications and cited references. *J. Assoc. Inf. Sci. Technol.* **66**(11), 2215–2222 (2015). <https://doi.org/10.1002/asi.23329>
6. Brack, A., D’Souza, J., Hoppe, A., Auer, S., Ewerth, R.: Domain-independent extraction of scientific concepts from research articles. In: Jose, J.M., et al. (eds.) ECIR 2020, Part I. LNCS, vol. 12035, pp. 251–266. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45439-5_17

7. Chaimongkol, P., Aizawa, A., Tateisi, Y.: Corpus for coreference resolution on scientific papers. In: Calzolari, N., et al. (eds.) Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014, Reykjavik, Iceland, 26–31 May 2014, pp. 3187–3190. European Language Resources Association (ELRA) (2014). <http://www.lrec-conf.org/proceedings/lrec2014/summaries/286.html>
8. Chambers, A.: Statistical Models for Text Classification and Clustering: Applications and Analysis. Ph.D. thesis, University of California, Irvine (2013)
9. Clark, K., Manning, C.D.: Entity-centric coreference resolution with model stacking. In: ACL (1), pp. 1405–1415. The Association for Computer Linguistics (2015)
10. Cohen, J.: A coefficient of agreement for nominal scales. *Educ. Psychol. Measur.* **20**(1), 37–46 (1960)
11. Cohen, K.B., et al.: Coreference annotation and resolution in the colorado richly annotated full text (CRAFT) corpus of biomedical journal articles. *BMC Bioinform.* **18**(1), 372:1–372:14 (2017). <https://doi.org/10.1186/s12859-017-1775-9>
12. Denis, P., Baldridge, J.: Specialized models and ranking for coreference resolution. In: 2008 Conference on Empirical Methods in Natural Language Processing, EMNLP 2008, Proceedings of the Conference, 25–27 Oct 2008, Honolulu, Hawaii, USA. A meeting of SIGDAT, a Special Interest Group of the ACL, pp. 660–669. ACL (2008). <https://www.aclweb.org/anthology/D08-1069/>
13. Dessì, D., Osborne, F., Reforgiato Recupero, D., Buscaldi, D., Motta, E., Sack, H.: AI-KG: an automatically generated knowledge graph of artificial intelligence. In: Pan, J.Z., et al. (eds.) ISWC 2020, Part II. LNCS, vol. 12507, pp. 127–143. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-62466-8_9
14. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: Burstein, J., Doran, C., Solorio, T. (eds.) Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, 2–7 June 2019, vol. 1 (Long and Short Papers), pp. 4171–4186. Association for Computational Linguistics (2019). <https://doi.org/10.18653/v1/n19-1423>
15. Doddington, G.R., Mitchell, A., Przybocki, M.A., Ramshaw, L.A., Strassel, S.M., Weischedel, R.M.: The automatic content extraction (ACE) program - tasks, data, and evaluation. In: Proceedings of the Fourth International Conference on Language Resources and Evaluation, LREC 2004, 26–28 May 2004, Lisbon, Portugal. European Language Resources Association (2004). <http://www.lrec-conf.org/proceedings/lrec2004/summaries/5.htm>
16. D’Souza, J., Hoppe, A., Brack, A., Jaradeh, M.Y., Auer, S., Ewerth, R.: The STEM-ECR dataset: grounding scientific entity references in STEM scholarly content to authoritative encyclopedic and lexicographic sources. In: Calzolari, N., et al. (eds.) Proceedings of The 12th Language Resources and Evaluation Conference, LREC 2020, Marseille, France, 11–16 May 2020, pp. 2192–2203. European Language Resources Association (2020), <https://www.aclweb.org/anthology/2020.lrec-1.268/>
17. Elsevier Labs: Elsevier OA STM corpus. <https://github.com/elsevierlabs/OA-STM-Corpus> (2017). Accessed 15 July 2020
18. Fellbaum, C. (ed.): WordNet: An Electronic Lexical Database. Language, Speech, and Communication. MIT Press, Cambridge (1998)

19. Fisas, B., Saggion, H., Ronzano, F.: On the discursive structure of computer graphics research papers. In: Meyers, A., Rehbein, I., Zinsmeister, H. (eds.) Proceedings of The 9th Linguistic Annotation Workshop, LAW@NAACL-HLT 2015, 5 June 2015, Denver, Colorado, USA, pp. 42–51. The Association for Computer Linguistics (2015). <https://doi.org/10.3115/v1/w15-1605>
20. Gábor, K., Buscaldi, D., Schumann, A., QasemiZadeh, B., Zargayouna, H., Charnois, T.: Semeval-2018 task 7: semantic relation extraction and classification in scientific papers. In: Apidianaki, M., Mohammad, S.M., May, J., Shutova, E., Bethard, S., Carpuat, M. (eds.) Proceedings of The 12th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2018, New Orleans, Louisiana, USA, 5–6 June 2018, pp. 679–688. Association for Computational Linguistics (2018). <https://doi.org/10.18653/v1/s18-1111>
21. Joshi, M., Chen, D., Liu, Y., Weld, D.S., Zettlemoyer, L., Levy, O.: SpanBERT: improving pre-training by representing and predicting spans. *Trans. Assoc. Comput. Linguistics* **8**, 64–77 (2020). <https://transacl.org/ojs/index.php/tacl/article/view/1853>
22. Joshi, M., Levy, O., Zettlemoyer, L., Weld, D.S.: BERT for coreference resolution: Baselines and analysis. In: Inui, K., Jiang, J., Ng, V., Wan, X. (eds.) Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, 3–7 Nov 2019, pp. 5802–5807. Association for Computational Linguistics (2019). <https://doi.org/10.18653/v1/D19-1588>
23. Kim, J., Nguyen, N.L.T., Wang, Y., Tsujii, J., Takagi, T., Yonezawa, A.: The gena event and protein coreference tasks of the BioNLP shared task 2011. *BMC Bioinform.* **13**(S-11), S1 (2012). <https://doi.org/10.1186/1471-2105-13-S11-S1>
24. Kopeć, M., Ogrodniczuk, M.: Inter-annotator agreement in coreference annotation of polish. In: Sobecki, J., Boonjing, V., Chittayasothorn, S. (eds.) *Advanced Approaches to Intelligent Information and Database Systems*. SCI, vol. 551, pp. 149–158. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-05503-9_15
25. Lee, K., He, L., Lewis, M., Zettlemoyer, L.: End-to-end neural coreference resolution. In: Palmer, M., Hwa, R., Riedel, S. (eds.) Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, 9–11 Sept 2017, pp. 188–197. Association for Computational Linguistics (2017). <https://doi.org/10.18653/v1/d17-1018>
26. Lee, K., He, L., Zettlemoyer, L.: Higher-order coreference resolution with coarse-to-fine inference. In: Walker, M.A., Ji, H., Stent, A. (eds.) Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, 1–6 June 2018, vol. 2 (Short Papers), pp. 687–692. Association for Computational Linguistics (2018). <https://doi.org/10.18653/v1/n18-2108>
27. Luan, Y., He, L., Ostendorf, M., Hajishirzi, H.: Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction. In: Riloff, E., Chiang, D., Hockenmaier, J., Tsujii, J. (eds.) Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, 31 Oct – 4 Nov 2018, pp. 3219–3232. Association for Computational Linguistics (2018). <https://doi.org/10.18653/v1/d18-1360>
28. Lubani, M., Noah, S.A.M., Mahmud, R.: Ontology population: approaches and design aspects. *J. Inf. Sci.* **45**(4), 502–515 (2019). <https://doi.org/10.1177/0165551518801819>

29. Luo, X.: On coreference resolution performance metrics. In: HLT/EMNLP 2005, Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference, 6–8 Oct 2005, Vancouver, British Columbia, Canada, pp. 25–32. The Association for Computational Linguistics (2005). <https://www.aclweb.org/anthology/H05-1004/>
30. Ma, J., et al.: Jointly optimized neural coreference resolution with mutual attention. In: Caverlee, J., Hu, X.B., Lalmas, M., Wang, W. (eds.) WSDM 2020: The Thirteenth ACM International Conference on Web Search and Data Mining, Houston, TX, USA, 3–7 Feb 2020, pp. 402–410. ACM (2020). <https://doi.org/10.1145/3336191>.
31. Marasovic, A., Born, L., Opitz, J., Frank, A.: A mention-ranking model for abstract anaphora resolution. In: Palmer, M., Hwa, R., Riedel, S. (eds.) Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, 9–11 Sept 2017, pp. 221–232. Association for Computational Linguistics (2017). <https://doi.org/10.18653/v1/d17-1021>
32. Ng, V.: Machine learning for entity coreference resolution: a retrospective look at two decades of research. In: Singh, S.P., Markovitch, S. (eds.) Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, 4–9 Feb 2017, San Francisco, California, USA, pp. 4877–4884. AAAI Press (2017). <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14995>
33. Ng, V., Cardie, C.: Identifying anaphoric and non-anaphoric noun phrases to improve coreference resolution. In: 19th International Conference on Computational Linguistics, COLING 2002, Howard International House and Academia Sinica, Taipei, Taiwan, 24 Aug – 1 Sept 2002 (2002). <https://www.aclweb.org/anthology/C02-1139/>
34. Nguyen, N.L.T., Kim, J., Miwa, M., Matsuzaki, T., Tsujii, J.: Improving protein coreference resolution by simple semantic classification. *BMC Bioinform.* **13**, 304 (2012). <https://doi.org/10.1186/1471-2105-13-304>
35. Peters, M.E., et al.: Deep contextualized word representations. In: Walker, M.A., Ji, H., Stent, A. (eds.) Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, 1–6 June 2018, vol. 1 (Long Papers), pp. 2227–2237. Association for Computational Linguistics (2018). <https://doi.org/10.18653/v1/n18-1202>
36. Pradhan, S., Luo, X., Recasens, M., Hovy, E.H., Ng, V., Strube, M.: Scoring coreference partitions of predicted mentions: a reference implementation. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, 22–27 June 2014, Baltimore, MD, USA, vol. 2: Short Papers, pp. 30–35. The Association for Computer Linguistics (2014). <https://doi.org/10.3115/v1/p14-2006>
37. Pradhan, S., et al.: Towards robust linguistic analysis using ontonotes. In: Hockenmaier, J., Riedel, S. (eds.) Proceedings of the Seventeenth Conference on Computational Natural Language Learning, CoNLL 2013, Sofia, Bulgaria, 8–9 Aug 2013, pp. 143–152. ACL (2013). <https://www.aclweb.org/anthology/W13-3516/>

38. Pradhan, S., Moschitti, A., Xue, N., Uryupina, O., Zhang, Y.: Conll-2012 shared task: modeling multilingual unrestricted coreference in OntoNotes. In: Pradhan, S., Moschitti, A., Xue, N. (eds.) Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning - Proceedings of the Shared Task: Modeling Multilingual Unrestricted Coreference in OntoNotes, EMNLP-CoNLL 2012, 13 July 2012, Jeju Island, Korea, pp. 1–40. ACL (2012). <https://www.aclweb.org/anthology/W12-4501/>
39. Pujara, J., Singh, S.: Mining knowledge graphs from text. In: Chang, Y., Zhai, C., Liu, Y., Maarek, Y. (eds.) Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM 2018, Marina Del Rey, CA, USA, 5–9 Feb 2018, pp. 789–790. ACM (2018). <https://doi.org/10.1145/3159652.3162011>
40. Q. Zadeh, B., Handschuh, S.: The ACL RD-TEC: a dataset for benchmarking terminology extraction and classification in computational linguistics. In: Proceedings of the 4th International Workshop on Computational Terminology (Computerm), pp. 52–63. Association for Computational Linguistics and Dublin City University, Dublin, Ireland (Aug 2014). <https://doi.org/10.3115/v1/W14-4807>
41. ur Rahman, M.A., Ng, V.: Supervised models for coreference resolution. In: Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing, EMNLP 2009, 6–7 Aug 2009, Singapore, A meeting of SIGDAT, a Special Interest Group of the ACL, pp. 968–977. ACL (2009). <https://www.aclweb.org/anthology/D09-1101/>
42. Ruder, S.: Neural Transfer Learning for Natural Language Processing. Ph.D. thesis, National University of Ireland, Galway (2019)
43. Sanh, V., Wolf, T., Ruder, S.: A hierarchical multi-task approach for learning embeddings from semantic tasks. In: The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, 27 Jan – 1 Feb 2019, pp. 6949–6956. AAAI Press (2019). <https://doi.org/10.1609/aaai.v33i01.33016949>
44. Schäfer, U., Spurk, C., Steffen, J.: A fully coreference-annotated corpus of scholarly papers from the ACL anthology. In: Kay, M., Boitet, C. (eds.) COLING 2012, 24th International Conference on Computational Linguistics, Proceedings of the Conference: Posters, 8–15 Dec 2012, Mumbai, India, pp. 1059–1070. Indian Institute of Technology Bombay (2012). <https://www.aclweb.org/anthology/C12-2103/>
45. Soon, W.M., Ng, H.T., Lim, C.Y.: A machine learning approach to coreference resolution of noun phrases. *Comput. Linguist.* **27**(4), 521–544 (2001). <https://doi.org/10.1162/089120101753342653>
46. Stenetorp, P., et al.: BRAT: a web-based tool for NLP-assisted text annotation. In: Daelemans, W., Lapata, M., Màrquez, L. (eds.) EACL 2012, 13th Conference of the European Chapter of the Association for Computational Linguistics, Avignon, France, 23–27 Apr 2012, pp. 102–107. The Association for Computer Linguistics (2012). <https://www.aclweb.org/anthology/E12-2021/>
47. Teufel, S., Siddharthan, A., Batchelor, C.R.: Towards domain-independent argumentative zoning: evidence from chemistry and computational linguistics. In: Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing, EMNLP 2009, 6–7 Aug 2009, Singapore, A meeting of SIGDAT, a Special Interest Group of the ACL, pp. 1493–1502. ACL (2009). <https://www.aclweb.org/anthology/D09-1155/>

48. Vilain, M.B., Burger, J.D., Aberdeen, J.S., Connolly, D., Hirschman, L.: A model-theoretic coreference scoring scheme. In: Proceedings of the 6th Conference on Message Understanding, MUC 1995, Columbia, Maryland, USA, 6–8 Nov 1995, pp. 45–52. ACL (1995). <https://doi.org/10.3115/1072399.1072405>
49. Wadden, D., Wennberg, U., Luan, Y., Hajishirzi, H.: Entity, relation, and event extraction with contextualized span representations. In: Inui, K., Jiang, J., Ng, V., Wan, X. (eds.) Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, 3–7 Nov 2019, pp. 5783–5788. Association for Computational Linguistics (2019). <https://doi.org/10.18653/v1/D19-1585>
50. Weikum, G., Dong, L., Razniewski, S., Suchanek, F.M.: Machine knowledge: Creation and curation of comprehensive knowledge bases. CoRR abs/2009.11564 (2020). <https://arxiv.org/abs/2009.11564>