

 Open access • Journal Article • DOI:10.1016/J.IPM.2021.102563

DWIE : an entity-centric dataset for multi-task document-level information extraction

— [Source link](#) 

[Klim Zaprojets](#), [Johannes Deleu](#), [Chris Develder](#), [Thomas Demeester](#)

Institutions: [Ghent University](#)

Published on: 01 Jul 2021 - [Information Processing and Management](#) (Pergamon)

Topics: [Entity linking](#), [Relationship extraction](#), [Named-entity recognition](#), [Information extraction and Graph \(abstract data type\)](#)

Related papers:

- [DWIE: an entity-centric dataset for multi-task document-level information extraction.](#)
- [Improving Coreference Resolution by Leveraging Entity-Centric Features with Graph Neural Networks and Second-order Inference.](#)
- [Progressive Joint Framework for Chinese Question Entity Discovery and Linking With Question Representations](#)
- [From unstructured to structured: Context-based named entity mining from text](#)
- [Integrating Named Entities in a Semantic Search Engine.](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/dwie-an-entity-centric-dataset-for-multi-task-document-level-3zxp66j5f>

DWIE: an entity-centric dataset for multi-task document-level information extraction.

Klim Zaporojets^{1,*}, Johannes Deleu¹, Chris Develder¹, Thomas Demeester¹

*Ghent University – imec, IDLab, Department of Information Technology,
Technologiepark Zwijnaarde 15, 9052 Ghent, Belgium*

Abstract

This paper presents DWIE, the ‘Deutsche Welle corpus for Information Extraction’, a newly created multi-task dataset that combines four main Information Extraction (IE) annotation subtasks: (i) Named Entity Recognition (NER), (ii) Coreference Resolution, (iii) Relation Extraction (RE), and (iv) Entity Linking. DWIE is conceived as an *entity-centric* dataset that describes interactions and properties of conceptual entities on the level of the complete document. This contrasts with currently dominant *mention-driven* approaches that start from the detection and classification of named entity mentions in individual sentences. Further, DWIE presented two main challenges when building and evaluating IE models for it. *First*, the use of traditional mention-level evaluation metrics for NER and RE tasks on entity-centric DWIE dataset can result in measurements dominated by predictions on more frequently mentioned entities. We tackle this issue by proposing a new entity-driven metric that takes into account the number of mentions that compose each of the predicted and ground truth entities. *Second*, the document-level multi-task annotations require the models to transfer information between entity mentions located in different parts of the document, as well as between different tasks, in a joint learning setting. To realize this, we propose to use graph-based neural message passing techniques between document-level mention spans. Our experiments show an improvement of up to 5.5 F₁ percentage points when incorporating neural graph propagation into our joint model. This demonstrates DWIE’s potential to stimulate further research in graph neural networks for representation learning in multi-task IE. We make DWIE publicly available at <https://github.com/klimzaporojets/DWIE>.

Keywords: Named Entity Recognition, Entity Linking, Relation Extraction, Coreference Resolution, Joint Models, Graph Neural Networks

1. Introduction

Information Extraction (IE) plays a fundamental role as a backbone component in many downstream applications. For example, an application such as question answering may be improved by relying on relation extraction (RE) (Hu et al., 2019; Yu et al., 2017), coreference resolution (Bhattacharjee et al., 2020; Gao et al., 2019), named entity recognition (NER) (Molla et al., 2006; Singh et al., 2018), and entity linking (EL) (Broscheit, 2019; Chen et al., 2017) components. This also holds for other applications such as personalized news recommendation (Karimi et al., 2018; Wang et al., 2018, 2019), fact checking (Thorne & Vlachos, 2018; Zhang & Ghorbani, 2020), opinion mining (Sun et al., 2017), semantic search (Cifariello et al., 2019), and conversational agents (Roller et al., 2020). The last decade has shown a growing interest in IE datasets suitably annotated for developing multi-task models where each of the tasks (e.g., NER, RE, etc.) would

*Corresponding author

Email addresses: klim.zaporojets@ugent.be (Klim Zaporojets), johannes.deleu@ugent.be (Johannes Deleu), chris.develder@ugent.be (Chris Develder), thomas.demeester@ugent.be (Thomas Demeester)

¹URL: <https://ugent2k.github.io/>

Textual Representation

1 Prince Harry gets engaged to actress Meghan Markle
2 Britain's Prince Harry is engaged to his US partner Meghan Markle, his father Prince Charles has announced.
3 The wedding is due to take place in the spring of 2018 and the couple are to live in Kensington Palace.
4 The Duke and Duchess of Cambridge, Harry's older brother Prince William and Kate Middleton, congratulated the couple.
5 "We are very excited for Harry and Meghan
6 It has been wonderful getting to know Meghan and to see how happy she and Harry are together," Clarence House said in a tweet.
7 Harry spent 10 years in the army and has this year, with his elder brother William, promoted mental health strategies for armed forces in a joint initiative between their Royal Foundation and the Ministry of Defense.
8 Harry is Queen Elizabeth's grandson and fifth-in-line to the British throne.

Entity-Centric Representation

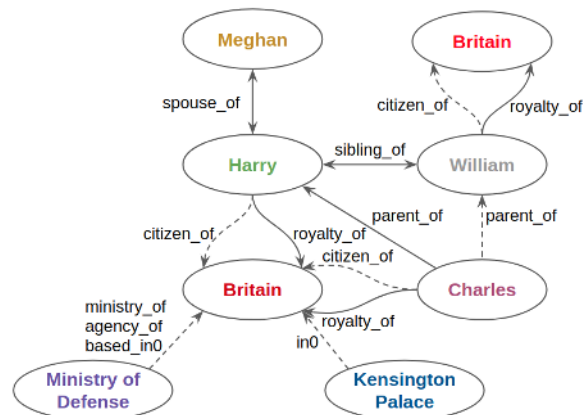


Figure 1: An example from the DWIE dataset with entity mentions underlined. We show 8 of the 29 entities in the graph on the right. It illustrates the relations that can be derived from the content of the article. The relations that are explicitly mentioned in the text (trigger-based) are depicted by solid arrows. Conversely, the relations that are implicit and/or need the whole document context (document-based) to be derived are represented by dashed arrows.

benefit from the interaction with (an)other task(s) (Bekoulis et al., 2018b; Fei et al., 2020; Lee et al., 2017, 2018; Luan et al., 2019), to boost their performance. However, the currently widely used IE datasets to build such multi-task models exhibit three major limitations. *First*, the annotation schema adopted in most of these datasets is mention-driven, focusing on annotating elements (e.g., relations, entity types) that involve specific entity mentions explicitly mentioned in the text. This produces very localized annotations (e.g., sentence-based relations between entity mentions) that do not reflect meaning that can be inferred on a more general document-level. *Second*, the number of annotated extraction tasks in most of the IE datasets is rather limited. Most of them focus on a single or at most a few different tasks. Furthermore, some other datasets, including the well-known TAC-KBPs (Ellis et al., 2015, 2014; Ji et al., 2010, 2015, 2017), use different non-overlapping corpora for each of the tracks that group a few related tasks. Consequently, current models addressing multiple IE tasks together often use multi-tasking (with different datasets per task) rather than really joint modeling approaches. *Finally*, the annotation of currently widely used IE datasets is driven by either relying on a priori defined annotation schemas (Augenstein et al., 2017; Doddington et al., 2004; Hendrickx et al., 2010; Song et al., 2015; Walker et al., 2006; Zhang et al., 2017b) or on distantly supervised labeling techniques (Han et al., 2018; Peng et al., 2017; Quirk & Poon, 2017; Riedel et al., 2010; Yao et al., 2019). In consequence, the resulting annotations are not necessarily representative of the actual information contained in the annotated corpus.

In this work, we tackle the aforementioned limitations of IE datasets by introducing a new dataset named DWIE. It consists of 802 general news articles in English, selected randomly from a corpus collected from Deutsche Welle² between 2002 and 2018, as part of the CPN project.³ We focus on annotating four main IE tasks: (i) Named Entity Recognition (NER), (ii) Coreference Resolution, (iii) Relation Extraction (RE), and (iv) Entity Linking.⁴ Figure 1 shows an example snippet from the DWIE corpus. We adopt an *entity-centric* approach where all annotations (i.e., for NER, RE and Entity Linking tasks) are made on the entity⁵ level. Each of the entities is composed by the coreferenced entity mentions from the entire document (e.g., the entity *Meghan* in Fig. 1 clusters the entity mentions “Meghan Markle” and “Meghan” across the whole document). This entity-centric approach contrasts with mention-driven annotations in widely used IE datasets (Doddington et al., 2004; Han et al., 2018; Hendrickx et al., 2010; Ji et al., 2015; Luan et al.,

²<https://www.dw.com>

³<https://www.projectcpn.eu>

⁴The linking is done to Wikipedia version 20181115.

⁵Also referred to as *entity cluster* or just *cluster*.

2018; Song et al., 2015; Zhang et al., 2017b) where the annotation process is biased towards considering only local explicit textual evidence to annotate elements such as relations and entity types (e.g., the relation `spouse_of`(*Meghan, Harry*) that can be extracted from the 1st sentence in Fig. 1). Consequently, our DWIE dataset paves the way for research on more complex document-level reasoning that goes beyond only the local textual context directly surrounding individual entity mentions. For example, consider the relation `ministry_of`(*Ministry of Defense, Britain*) in Fig. 1: while the text of the document does not directly state such a relation, it can be deduced from a more general document-level entity-centric vision of the article, i.e., combining the information involving the entities *Ministry of Defense* and *Harry* in sentence 7 with the one involving *Britain* and *Harry* in sentence 2. Finally, the entity-centric approach provides entity linking annotations that are consistent across the document: by clustering mentions of the same entity, and then providing links to the Wikidata KB (or NIL if the entity does not appear there) for the whole cluster at once, we limit annotation errors or accidental inconsistencies (in the linking itself, but also in terms of NER labels). To our knowledge, DWIE is the first dataset with this level of conceptual consistency over the considered information extraction tasks. We therefore expect that the dataset will play a key role in advancing research exploring potential benefits of (i) entity-level information extraction in terms of reducing potential inconsistent decisions (within EL across multiple mentions, as well as across multiple tasks), and (ii) using entity-centric information stored in a KB to complement the otherwise exclusively text-dependent IE tasks such as NER, RE, and coreference resolution.

Additionally, we use a bottom-up, data-driven annotation approach where we manually define our annotations (e.g., in terms of the entity and relation types) to maximally reflect the information of the corpus at hand. Currently dominant datasets are driven by distant supervision and executed top-down, by which we mean that the selection of entity and relation types is a priori defined and limited in coverage (i.e., the raw data potentially contains other types that thus remain un-annotated). Conversely, we do not a priori limit the entity and relation types to annotate, but adopt a bottom-up approach driven by the data itself. Our proposed bottom-up approach encompasses a three-pass annotation procedure where we use the first exploratory annotation pass to derive the main annotation types (annotation schema) from the corpus, and the next two passes to perform schema-driven annotations and refine them by carrying out an additional parallel annotation of the corpus for fixing errors inferred from inter-annotator inconsistencies.

Besides the dataset itself, we also contribute empirical modeling results to address the aforementioned IE tasks. Our goal is to study two important properties that are inherent to DWIE. The first key property is the need for *long-range contextual information sharing* to make document-level predictions involving entities whose mentions are located in different parts of the document. The second key property involves the *joint interaction between tasks* where the information obtained in one task can help to solve another task. For example, in Fig. 1 knowing the types of entities (which involves NER and coreference tasks) *Britain* and *Kensington Palace* can boost the performance of the relation extraction task by limiting the number of possible relation types between these two entities (e.g., `ministry_of` but not `citizen_of`). In order to study the impact of these two phenomena inherent to our DWIE dataset on the final results, we experiment with neural graph-based models (Li et al., 2016b; Wu et al., 2020; Xu et al., 2018). These models allow message passing between local contextual encodings, making it possible to measure the impact of local contextual information sharing both on a more general document level and across the tasks. Furthermore, previous work already has shown the positive effect of using graph-based information passing techniques on single tasks (Kantor & Globerson, 2019; Lee et al., 2018), and between tasks (Fei et al., 2020; Fu et al., 2019; Luan et al., 2019; Wadden et al., 2019) on mention-driven datasets. We expand this work even further by extending these models to be used on the entity-centric, document-level DWIE dataset. More specifically, we experiment with both single-task (Section 4.5) as well as joint (Section 4.2) models to study the effect of contextual information propagation in single task and joint settings. Additionally, for the NER and RE tasks, we propose a new entity-centric evaluation metric that not only considers the predictions on separate entity mentions (as is done in related IE datasets), but also accounts for the impact of the predictions on entity cluster level.

In summary, the main objective that we address in the current paper is to introduce an entity-centric multi-task IE dataset that covers different related tasks on a document level as well as provides a connection with external structured knowledge (through the entity linking task). Furthermore, we aim to explore how

neural graph-based models can boost the performance by enabling local contextual information propagation across the document (single-task models) and between different tasks (joint models). The results presented in this paper suggest that, while challenging, DWIE opens up new possibilities of research in the domain of joint entity-centric information extraction methods. The main contributions of our work are that:

- (1) We construct a self-contained dataset (Section 3) with joint annotations for four basic information extraction tasks (NER, entity linking, coreference resolution, and RE), that provide entity-centric document-level annotations (as opposed to typical mention-driven sentence-level annotations for, e.g., RE) connecting unstructured (text) and structured (KB) information sources.
- (2) We introduce a data-driven, bottom-up three-pass annotation approach complemented by context-based logical rules to build such dataset (Section 3).
- (3) We propose a new evaluation metric for the NER and RE tasks (Section 5), in line with the entity-centric nature of DWIE.
- (4) We extend the competitive graph-based neural IE model DyGIE (Luan et al., 2019) for the four IE tasks in DWIE (Section 4) and provide source code for NER, coreference resolution, and RE. Furthermore, we introduce a new latent attention-driven **AttProp** graph propagation method and show its advantages in both single and joint model settings. The experimental results (Section 6) demonstrate the potential of such neural graph based models.

2. Related Work

This section summarizes the overview of related datasets (Section 2.1), and explores the differences between our newly created DWIE and other similar datasets widely used by the scientific community. The main qualitative differences are presented in Table 1, while the quantitative comparison is provided in Table 2. Next, we describe the current trends in IE to solve the tasks included in DWIE, and compare them to our proposed approach (Section 2.2). Finally, we discuss currently used metrics to evaluate model performance on IE datasets and introduce some challenges in applying them to measuring the performance on DWIE (Section 2.3).

2.1. Related Datasets

Most of IE datasets have focused on a single task, making it very challenging to develop systems that jointly train for different annotation subtasks on a single corpus. Well-known single-task datasets include (i) *for NER*: CoNLL-2003 (Sang & De Meulder, 2003) and WNUT 2017 (Derczynski et al., 2017), (ii) *for relation extraction*: Semeval-2010 T8 (Hendrickx et al., 2010), TACRED (Zhang et al., 2017b) and FewRel (Han et al., 2018), (iii) *for entity linking*: IITB (Kulkarni et al., 2009), CoNLL-YAGO (Hoffart et al., 2011), and WikilinksNED (Eshel et al., 2017), and (iv) *for coreference resolution*: CoNLL-2012 (Pradhan et al., 2012) and GAP (Webster et al., 2018). Conversely, in this work we propose a multi-task dataset as a single corpus annotated with different information extraction layers: named entities, mention clustering in entities (i.e., coreference), relations between entity clusters of mentions, and entity linking. We further complement our dataset with additional tasks such as document classification and keyword extraction. It is worth noting that our *coreference* annotations differ from the widely adopted CoNLL-2012 (Pradhan et al., 2012) scheme in two aspects: (i) we retain singleton entities composed by only one mention as a valid entity cluster, (ii) we only cluster proper nouns, leaving out nominal and anaphoric expressions.

Furthermore, most prominent efforts to produce jointly annotated datasets have focused on using a *top-down* annotation approach. This method involves an a priori defined annotation schema that drives the process of selection and labeling of the corpus. The de facto datasets used in most of the joint learning baselines such as ACE 2005 (Doddingon et al., 2004; Walker et al., 2006), TAC-KBPs (Ellis et al., 2015, 2014; Ji et al., 2010, 2015, 2017) and Rich ERE (Song et al., 2015) use this annotation approach. More specifically, during the creation of the ACE 2005 dataset, the annotators initially tagged candidate documents as “good” or “bad” depending on the estimated number and types of entities present in each one. In subsequent

Table 1: Qualitative comparison of the datasets. We divide our comparison in five groups: (i) *Core Tasks* represent the main subtasks covered in DWIE, (ii) *Doc-Based* indicates whether different subtasks are annotated on the document-level, (iii) *Entity-Centric* indicates which annotations are done with respect to entity clusters (✓) as opposed to individual mentions (✗), (iv) *Unaided* specifies whether the annotation process was completely manual (✓) or with some form of distant supervision (✗), and (v) *Open* indicates whether the dataset is freely available.

Dataset	Core Tasks				Doc-Based				Entity-Centric				Unaided				
	NER	Coreference	Relations	Linking	Coreference	Relations	Multi-label Rel	Keywords	Classification	Multi-label Ent	Relations	Linking	NER	Coreference	Relations	Linking	Open
DWIE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
TAC-KBP (Ji et al., 2010, 2015, 2017)	✓	✓	✓	✓	✗	✗	✗	✗	✓	✓	✗	✗	✓	✓	✓	✓	✗
BC5CDR (Li et al., 2016a; Wei et al., 2015)	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✓	✗	✓	✓	✗	✓	✓
MUC-7 (Chinchor & Marsh, 1998)	✓	✓	✓	✗	✓	✓	✗	✗	✓	✗	✓	✗	✓	✓	✓	✗	✗
SciERC (Luan et al., 2018)	✓	✓	✓	✗	✗	✗	✗	✗	✓	✗	✗	✗	✓	✓	✓	✓	✓
DocRED (Yao et al., 2019)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Rich ERE (Aguilar et al., 2014; Song et al., 2015)	✓	✓	✓	✗	✗	✗	✗	✗	✓	✗	✗	✗	✓	✓	✓	✗	✗
ACE 2005 (Walker et al., 2006)	✓	✓	✓	✗	✗	✗	✗	✗	✓	✗	✗	✗	✓	✓	✓	✗	✗
OntoNotes 5.0 (Hovy et al., 2006; Weischedel et al., 2013)	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓
ScienceIE (Augenstein et al., 2017)	✓	✗	✗	✗	✗	✓	✓	✓	✗	✗	✗	✗	✗	✗	✓	✓	✓
FewRel (Han et al., 2018)	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✓	✓	✗	✗	✗	✗	✓
GENIA (Kim et al., 2003)	✓	✓	✓	✗	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✓
AIDA CoNLL-YAGO (Hoffart et al., 2011)	✓	✗	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✓	✓
SemEval 2010 T8 (Hendrickx et al., 2010)	✓	✗	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	✗	✓	✓	✓
NYT (Riedel et al., 2010)	✓	✗	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓
ACEtoWiki (Bentivogli et al., 2010)	✗	✗	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓
WNUT 2017 (Derczynski et al., 2017)	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✓
CoNLL-2003 (Sang & De Meulder, 2003)	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✓
TACRED (Zhang et al., 2017b)	✗	✗	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗

annotation stages, only “good” documents were fully annotated and included in the final dataset. Similarly, during the creation of the TAC-KBP datasets, the annotators focused on producing annotations evenly distributed among three entity types (PERs, ORGs, and GPEs) by annotating only the documents that contained a minimum number of entities related to event types. In the case of Rich ERE, the documents to tag were prioritized by the event trigger word density calculated per 1,000 tokens, thus focusing only on content with a high number of previously defined key event-related tokens. Furthermore, other IE-related datasets (Augenstein et al., 2017; Han et al., 2018; Hendrickx et al., 2010; Yao et al., 2019; Zhang et al., 2017b) use similar pre-filtering techniques in order to select the text to be annotated. As a consequence, the corpus and annotations in these datasets tend to be biased and likely not representative of the language used in the different input domains. Conversely, we adopt a radically different *bottom-up* approach where we derive the annotations (e.g., entity classification types, relation types) from the data itself. This bottom-up data-driven procedure guarantees that the annotations in DWIE are representative of the document corpus information and reflects the particularities of the language used in its journalistic domain. Furthermore, it better represents the properties that are inherently present in written corpora, e.g., the long-tail distribution of different annotation types.

Finally, from the perspective of the necessary evidence to annotate a particular entity type or relation, we propose to make a distinction for the currently existing datasets between *trigger-based* and *document-based* annotations (see *Doc-Based* comparison group in Table 1). The *trigger-based* datasets require that a particular relation or entity type should only be annotated if it is supported by an explicit reference in a text. For example, in Fig. 1 there is a concrete reference of the relation between “Meghan” and “Harry” in form of triggers such as “gets engaged” in sentence 1 and “The wedding” in sentence 2. Most of the

⁶The EDL track only of TAC-KBP 2010.

⁷Numbers based on publicly available train and development sets.

Table 2: Numerical comparison of DWIE and well-known IE datasets. Note that some datasets (including DWIE) use an entity-centric approach, organizing entity mentions in entity clusters, and annotating entities, relations, and linking on the cluster level. Hence, we provide both mention-level as well as cluster-level (if a particular dataset supports it) statistics.

Dataset	# Tokens	Entities			Relations			Linking	
		# Mentions	# Entity clusters	# Entity types	# Relation mentions	# Relation clusters	# Relation types	# Mention KB Links	# Cluster KB Links
NYT	5,765,332	1,388,982	-	-	142,823	-	52	1,388,982	-
TACRED	3,866,863	-	-	-	21,784	-	42	-	-
TAC-KBP ⁶	3,053,336	6,495	3,750	-	-	-	-	3,818	2,094
OntoNotes 5.0	2,088,832	161,783	136,037	-	-	-	-	-	-
FewRel ⁷	1,397,333	114,213	112,000	-	58,267	56,000	80	114,213	112,000
DocRED	1,018,297	132,392	98,610	6	155,535	50,503	96	-	-
MUC-4	717,798	14,196	-	13	-	-	-	-	-
GENIA	554,346	56,743	10,728	5	2,337	-	2	-	-
DWIE	501,095	43,373	23,130	311	317,204	21,749	65	28,482	13,086
BC5CDR	343,175	29,271	10,326	2	47,813	3,116	1	29,562	10,326
CoNLL-2003	301,418	35,089	-	4	-	-	-	-	-
CoNLL-YAGO	301,418	34,929	-	-	-	-	-	34,929	-
ACE 2005	259,889	54,824	37,622	51	8,419	7,786	18	-	-
ACEtoWiki	259,889	-	-	-	-	-	-	16,310	-
SEval 2010 T8	207,307	21,434	-	-	6,674	-	9	-	-
ACE 2004	185,696	29,949	12,507	43	5,976	5,525	24	-	-
WNUT 2017	101,857	3,890	-	6	-	-	-	-	-
ScienceIE	99,580	9,946	9,536	3	638	-	1	-	-
SciERC	65,334	8,094	1,015	6	2,687	-	7	-	-

traditionally used jointly annotated datasets such as ACE 2005 (Dodgington et al., 2004; Walker et al., 2006), TAC-KBPs (Ellis et al., 2015, 2014; Ji et al., 2010, 2015, 2017) and Rich ERE (Song et al., 2015), as well as others, including FewRel (Han et al., 2018), OntoNotes (Hovy et al., 2006; Weischedel et al., 2013), TACRED (Zhang et al., 2017b), SemEval 2010 Task 8 (Hendrickx et al., 2010) and SciERC (Luan et al., 2018), are *trigger-based*. The disadvantage of such an approach is that it only captures the most simple cases of relations and entity types that are explicitly mentioned in the text. As a general rule, this also limits the datasets to cover only the relations between entity mentions (i.e., the annotation process is mention-driven) that appear within a single or at most few adjacent sentences where the relation trigger occurs (see Fig. 2 in Section 3 for a more detailed illustration of this phenomenon). However, as we move to a broader *document-based* interpretation, it is common to find relations that are not explicitly mentioned in text. Thus, in our example of Fig. 1 the relation between “Ministry of Defense” and “Britain” is not explicitly indicated in the text. However, after reading the whole article we can infer relations such as *ministry_of*, *agency_of* and *based_in* between these two entities. This document-level reasoning makes it essential to adopt an entity-centric approach (see *Entity-Centric* comparison group in Table 1) where each *entity* comprises one or more entity *mentions*, and the annotations (i.e., *relations*, *entity tags* and *entity linking* in DWIE) are made on the entity level, thus abstracting from specific *mention-driven* triggers.

2.2. Recent advances in Information Extraction

In the last couple of years, the advances in joint modeling have been accompanied by an ever increasing interest in the use of graph-based neural networks (Li et al., 2016b; Wu et al., 2020; Xu et al., 2018). Initially, this approach has been applied to improve the performance of the single coreference resolution task by transferring document-level contextual information between coreferenced entity mention spans (Kantor & Globerson, 2019; Lee et al., 2018). Most recently, these graph propagation techniques have been successfully used in a joint setting (Fei et al., 2020; Fu et al., 2019; Luan et al., 2019; Wadden et al., 2019) by performing graph message passing updates between the shared spans across different tasks. However, while successful on mention-driven datasets such as ACE 2005 (Walker et al., 2006) and NYT (Riedel et al., 2010), as far as we are aware, the advantages of these techniques have not yet been investigated in an entity-centric document-level setting. We fill this gap by extending the neural graph-based model initially proposed by Luan et al.

(2019) to be used on DWIE (see Section 4). More specifically, we explore the effect of performing document-level coreference (**CorefProp**) (Lee et al., 2018; Luan et al., 2019) and relation-driven (**RelProp**) (Luan et al., 2019) graph message passing updates between the spans. Additionally, we introduce a new latent attention-based graph propagation method (**AttProp**) and compare it to previously proposed task-driven graph propagation methods (**CorefProp** and **RelProp**).

2.3. Metrics and evaluation

Current dominant IE systems consider *mention-level* scoring of NER as well as RE components when reporting on datasets such as CoNLL-2003 (Akbik et al., 2019, 2018; Baevski et al., 2019; Chiu & Nichols, 2016; Lample et al., 2016), OntoNotes (Chiu & Nichols, 2016; Clark et al., 2018; Strubell et al., 2017), ACE 2004 (Bekoulis et al., 2018a; Li & Ji, 2014; Zhang et al., 2017a), ACE 2005 (Fei et al., 2020; Luan et al., 2019; Zhang et al., 2017a), TACRED (Soares et al., 2019; Zhang et al., 2018, 2017b), and SelEval 2010-Task 8 (Guo et al., 2019; Hu et al., 2020; Peters et al., 2019) among others. In contrast, the DWIE dataset is entity-centric where all the annotations are done on the entity cluster level. Consequently, adopting a purely mention-based evaluation approach can lead to a dominance of the score by predictions on entities composed by many mentions as opposed to entities composed by only few ones. Conversely, a purely cluster-level evaluation would be overly strict, requiring correct prediction of relation/entity types as well as an exact match of the predicted entity clusters. To tackle this problem, we propose a new scoring method that combines entity mention-level and cluster-level evaluation, while avoiding the pitfalls of either method alone (see Section 5).

3. Annotation process

In this work we introduce our *bottom-up* data-driven annotation approach. Our main goal is to get an annotation schema that reflects the types of entities and relations that are effectively mentioned throughout the corpus to maximally capture the information it contains. Therefore, we derive the annotation schema from the corpus itself, adopting three annotation passes that are detailed next: (i) *exploratory pass*, (ii) *schema-driven pass*, and (iii) *inter-annotator refinement*. Each pass encompasses substeps to cover all IE subtasks: (i) mention annotation (i.e., the entities and their types), (ii) coreference resolution, (iii) relation extraction on the entity level (i.e., clustering all mentions referring to the same entity), and (iv) entity linking (again, on the entity level, providing the same link for all clustered mentions).

3.1. Exploratory pass

The first annotation pass aims to discover the annotation structure (i.e., annotation schema) to be used on the corpus, in particular the types to use for named entity recognition (NER) and relation extraction (RE) tasks. Three annotators are involved in this step to provide annotations on the mention level: one expert annotator and two paid students. However, no parallel annotation is done and the role of the expert annotator is to annotate part of the corpus, as well as instructing and supervising the paid annotators. No a priori fixed schema is followed, but we ask the annotators to be as consistent as possible during the process. More specifically, the annotators are free to define their own entity and relation types for the NER and RE tasks that reflect the contents of the articles as long as they comply with the following generic guidelines:

1. **Named Entities:** any physical or abstract object (e.g., “Washington”, “Jeff Davis”, “Nobel Prize”, “Lisbon Treaty”, etc.) that can be denoted with a proper noun. Entities are usually upper-cased in the text, although values such as money and time can also be included. Use short and specific entity types (e.g., person, organization, etc.) to classify entities, the types can be overlapping (a single entity can have multiple types).
2. **Relations:** identify meaningful relations between entities. The type of a relation should be specific and reflect the type of the connected entities as well as the semantic meaning of the relation. For example, instead of using a generic “located in” relation for entities located in a particular country, we can divide it in “based in country” for organizations that are based in a country, “city located in

Table 3: *Descriptions* and *Examples* (with entity mentions underlined) of each of the most granular entity classes in DWIE (*ENTITY*, *VALUE* and *OTHER*) in the *type* tag hierarchy. Additionally, for the type *ENTITY*, we describe and give examples of each of its direct subtypes (*location*, *organization*, *person*, *misc*, *event* and *ethnicity*).

Entity Tag	Description	Example
ENTITY	All nominal named entities.	“ <u>UK</u> court rules <u>WikiLeaks</u> ’ <u>Assange</u> should be extradited to <u>Sweden</u> ”
location	Entities referring to a particular geographical location.	“ <u>Libya</u> is one of <u>Germany</u> ’s strongest trading partners in northern <u>Africa</u> .”
organization	Organizations such as companies, governmental organizations, etc.	“According to the report, <u>Amazon</u> would pay the same level of royalty fees as <u>Apple</u> .”
person	Entities referring to people in general such as politicians, artists, sport players, etc.	“With <u>Ramires</u> out, <u>Drogba</u> could start as striker, with <u>Torres</u> moving to the wing.”
misc	Miscellaneous entity types such as names of work of arts, treaties, product names, etc.	“According to the director’s own words, <u>The Post</u> is a ‘patriotic film.’”
event	Events such as sport competitions, summits, etc.	“Last year’s <u>Champions League</u> final drew a crowd of just 14,303.”
ethnicity	Entity type used to identify different ethnic groups.	“Attempt to assimilate <u>Uyghurs</u> into dominant <u>Han Chinese</u> culture.”
VALUE	Values in general such as time, money, etc.	“It ended the <u>2014</u> fiscal year <u>45 million euros</u> (<u>\$51 million</u>) in the red.”
OTHER	Includes the nominal variations of entity types (e.g., includes variations of country names such as “German”, which is a variation of “Germany”).	“ <u>Franco-German</u> ‘war child’ granted <u>German</u> citizenship.”

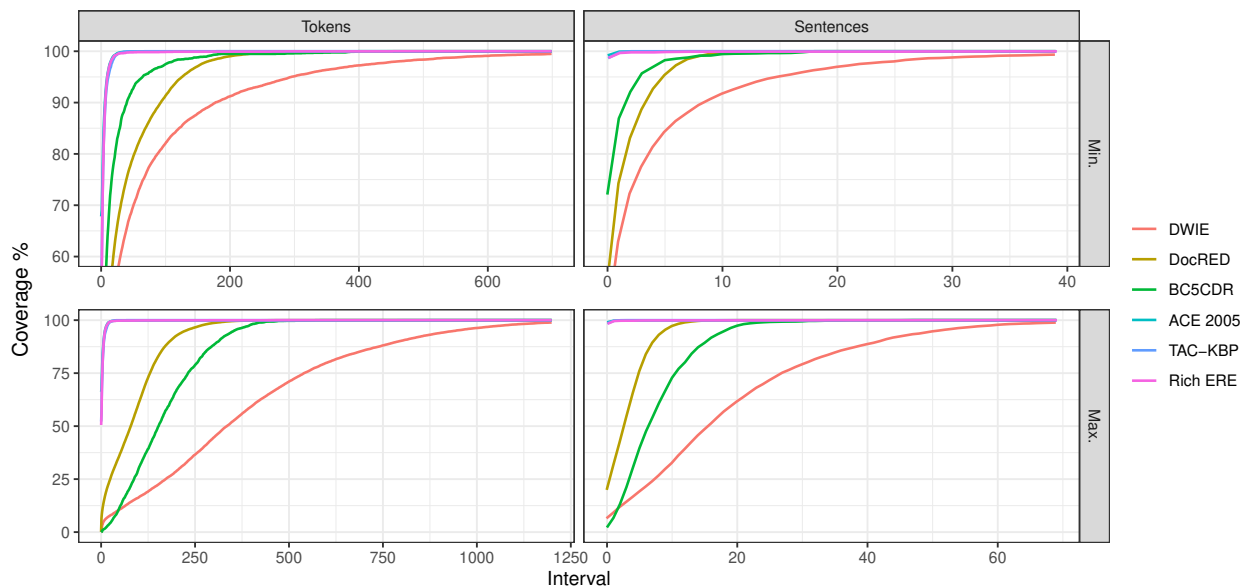


Figure 2: Comparison of the coverage of the % of relations with increasing interval in tokens (left) and interval in sentences (right). The graph at the top illustrates the relations coverage measuring the minimum distance between entities (closest mentions). Conversely, the graph at the bottom shows the coverage measuring the maximum distance between entity mentions. In both graphs, we note that the distance between the related mentions in our dataset is higher than in other widely used datasets.

country” for cities located in the country, etc. The types of the relations should have short names, ideally not exceeding 15 characters.

By not constraining the annotation process to specific entity and relation types, we ensure that our annotations are representative of the actual information contained in the annotated corpus.

3.2. Schema-driven pass

The main goal of this step is to create a consistent annotation schema for (i) named entity types and (ii) relation types based on the annotations made in the *exploratory pass*. As a first step, we identify the classification *tags* to be assigned to **entities**. We divide these tags in five main categories: *type*, *topic*, *iptc*, *slot*, and *gender* (see Table A.4). Our *type* tag is organized in a hierarchical structure (see Table A.1 in Appendix A), making it easier to extend our annotations to more granular subtypes. Table 3 defines and provides examples of each of the top *type* tags in the entity type hierarchy (*ENTITY*, *VALUE* and *OTHER*) as well as the direct subtypes of *ENTITY*. The *topic* tag allows to assign topics (e.g., politics, culture, education, etc.) to the entities and it complements the *type* tag (see Table A.2). The *iptc* tag is used for the universally defined IPTC news categories based on a media taxonomy (<https://iptc.org/standards/subject-codes/>). The *slot* tag is used for additional categorization that is transversal to different entity types. One example of this is the slot *interviewee* that can be assigned to any person (entity of type *person*) interviewed in a particular article.⁸ Finally, the *gender* tag is used to indicate the gender of the entities that refer to people. By defining these multiple overlapping tag types, we realize that the entity classification is multi-label by nature and thus allows different complementary entity tags to be assigned to a particular entity.⁹ This contrasts with prevailing single-label multi-class datasets such as ACE 2005 (Doddington et al., 2004; Walker et al., 2006), TAC-KBPs (Ellis et al., 2015, 2014; Ji et al., 2015, 2017), Rich ERE (Song et al., 2015), WNUT 2017 (Derczynski et al., 2017) and CoNLL-2003 (Sang & De Meulder, 2003).

⁸Other possible slot values are: *keyword*, *head*, *death*, *interviewer* and *expert*.

⁹The average number of labels per entity is 4.0 in our DWIE dataset.

Table 4: *Descriptions* and *Examples* of the top 5 most occurring relation types in DWIE. The entity mentions involved in the relations are underlined.

Relation Type	Description	Example
based_in0	Relations between organizations and the countries they are based in, ex: <code>based_in0</code> (<i>University of Cologne, Germany</i>)	“Now he’s back in <u>Germany</u> carrying on with his cancer research at the <u>University of Cologne</u> .”
in0	Relations between geographic locations and the countries they are located in, ex: <code>in0</code> (<i>Athens, Greece</i>)	“The murder of a left-wing activist in <u>Athens</u> has shaken up <u>Greece</u> and inspired a backlash.”
citizen_of	Relations between people and the country they are citizens of, ex: <code>citizen_of</code> (<i>Guerrero, Peru</i>)	“Even as a teenager, <u>Guerrero</u> played for the national side in his native <u>Peru</u> .”
based_in0-x	Relations between organizations and the nominal variations of the countries they are based in, ex: <code>based_in0-x</code> (<i>SPD, German</i>)	“ <u>SPD</u> denies ‘green light’ for new <u>German</u> government, but keeps options open”
citizen_of-x	Relations between people and the nominal variations of the countries they are citizens of, ex: <code>citizen_of-x</code> (<i>Assange, Australian</i>)	“ <u>Australian</u> national <u>Assange</u> said the accusations were politically motivated.”

For our **relation** annotations, we focus on annotating relations between entities themselves (cf. *document-based entity-centric* approach). Our adopted approach allows us to think concept-wise and come up not only with relations that are explicitly stated, but also those that can be implicitly inferred from the text. As a result, our dataset includes relations whose connected mentions are located further apart in the document. This can be seen in Fig. 2, where we compare the *minimum* (*Min.*) and *maximum* (*Max.*) distances between the mentions of the two entities connected by a relation for various mention-driven (Rich ERE¹⁰, TAC-KBP¹¹, and ACE 2005) and entity-centric (DocRED, BC5CDR, and the final version of our DWIE dataset) RE datasets. We note how other datasets that define the relation in terms of entities (BC5CDR and DocRED) require a higher number of token and sentence spans to cover all the relations in the respective dataset: entity-centric relations very often involve mentions located in different sentences in the document that refer to those entities. This is not the case for mention-driven trigger-based relations as in the TAC-KBP, Rich ERE and ACE 2005 datasets, where the annotation bias is towards finding explicitly mentioned relations, often involving entity mentions in a single sentence.

Similarly as with entity tags, we organize our relation annotations using multi-label types (see Table A.5 for details). Table 4 gives some examples from the DWIE corpus for the top 5 most occurring relation types (a detailed list can be consulted in Table A.6). For reasons of space, the examples only involve relations between entities whose mentions occur in a single sentence; for an example involving document-level relations we refer to Fig. 1.

Additionally, we define logical rules to automatically guarantee the consistency of the relations and their types. The following is an example,

$$\text{based_in2}\langle X, Z \rangle \wedge \text{in0}\langle Z, Y \rangle \implies \text{based_in0}\langle X, Y \rangle \quad (1)$$

reflecting the knowledge that if an organization X is based in a city Z (relation `based_in2`), and that this city Z is located in the country Y (relation `in0`), the fact that company X is also located in that country (relation `based_in0`) is valid as well. The goal of this step is mainly consistency of the annotations, but it

¹⁰We use the Rich ERE dataset from the LDC2015E29 and LDC2015E68 catalogs.

¹¹We use the TAK-KBP 2017 dataset from the LDC2017E54 and LDC2017E55 catalogs.

Table 5: The inter-annotation agreement Cohen’s kappa scores for all the different annotation tasks *before* and *after* the dataset refinement used to analyze and correct the discrepancies between the parallel annotations.

Task	Before Refinement	After Refinement
Named Entity	0.8497	0.8703
Named Entity Detection	0.9665	0.9673
Named Entity Classification	0.8812	0.9026
Coreference	0.9302	0.9324
Entity Linking	0.9280	0.9320
Relation	0.6594	0.8729
Relation Detection	0.7686	0.8727
Relation Classification	0.8118	0.9666

implies that an effective predictor would need to perform some form of reasoning to correctly predict all relations in the dataset. A complete list of logical rules is provided in Appendix C.

3.3. Inter-Annotator Refinement

In order to assess and further improve the quality of our dataset we re-annotate a 100 randomly selected news articles (12.5% of the articles used in the previous annotation rounds) from scratch. This work is done by a second independent expert annotator. The annotations in this pass are performed by following the already defined annotation schema based on the annotation process in the *exploratory* and *schema-driven* passes. We use this second annotated subset to measure the inter-annotator agreement and subsequently determine the parts of the dataset that still need to be improved. Table 5 compares the kappa scores before and after this refinement pass for each of the tasks (see Appendix B for details on how the kappa score is calculated). We observe that, after the refinement, all of the kappa scores are above 0.85, which is considered a ‘strong’ (McHugh, 2012) to ‘almost perfect’ (Landis & Koch, 1977) agreement.

Note that the revisions were seeded by and evaluated on the subset of 100 re-annotated articles. However, we argue that the inter-annotator refinement improved the annotation consistency of the entire dataset, given that the reviewed entity and relation types are used in more than 99.4% of all annotations in DWIE.

4. Model Architecture

In this section we introduce the end-to-end architecture used to compare the performance of models trained on the separate tasks with the models that are trained jointly for multiple tasks on the DWIE dataset. The main component of our approach is the use of Graph Neural Networks (Li et al., 2016b; Scarselli et al., 2008; Wu et al., 2020; Xu et al., 2018), relying on propagation techniques in both single-task and joint setups. More specifically, we implement span-based graph message passing on coreference (**CorefProp**) (Lee et al., 2018; Luan et al., 2019) and relation levels (**RelProp**) (Luan et al., 2019). Additionally, we introduce a latent attentive propagation method (**AttProp**) which is not driven by annotations of any task in particular and, as a result, can be freely applied to any task or joint combination of tasks. The interconnection between the different components of our model architecture is depicted in Fig. 3. It is based on the *span-based architecture* introduced in Lee et al. (2017), which supports training on the space of all entity spans simultaneously, dynamically updating span representations by using the graph propagation approach (further detailed in Section 4.4). Recent works have shown that this idea has the potential for improved effectiveness (albeit at a higher computational cost) (Dixit & Al-Onaizan, 2019; Fei et al., 2020; Lee et al., 2018; Luan et al., 2019), compared to more traditional sequence-labeling approaches (Katiyar & Cardie, 2018; Lample et al., 2016; Luan et al., 2017; Ma & Hovy, 2016). More concretely, the use of a span-based approach where all the spans are shared between the individual task modules avoids the cascading of errors from the entity mention identification module (*entity scorer* in Fig. 3) to the rest of the tasks.

The most similar architecture to ours in using joint span-based neural graph IE is DyGIE (Luan et al., 2019) and its successor DyGIE++ (Wadden et al., 2019). Our model is described in detail below, but here we already list the aspects in which it differs from these models:

1. We introduce the graph propagation technique **AttProp** (see Section 4.4), which is not directly conditioned on a particular task and can be used in single-task (for each of the tasks) as well as joint settings.
2. We define a coreference architecture that, unlike previous work in span-based coreference resolution (Lee et al., 2017, 2018), allows to also account for singleton entities in the DWIE dataset (see Sections 4.2.2 and 4.5.2) by using an additional *pruner loss*, which turns out essential for the single model focusing on end-to-end coreference resolution.
3. Due to the document-level nature of DWIE, we run graph propagations on the whole document. This contrasts with a sentence-based approach adopted initially in the DyGIE/DyGIE++ architectures. It also drives some changes such as the use of a single *pruner* (see Section 4.1) to extract spans used in coreference and RE modules. Similarly, instead of applying the shared BiLSTM sentence by sentence as in Luan et al. (2019) and Wadden et al. (2019), we do it on the entire document, in order to allow capturing cross-sentence dependencies for document-level relations and entity clusters in DWIE.
4. We add an additional decoding step (see Section 4.3) needed to transform mention-based predictions for RE and NER tasks into entity-based ones, as required by the entity-centric nature of DWIE, and propose corresponding evaluation metrics (see Section 5).
5. Finally, we make changes in the loss and prediction components to support multi-label classification (in NER and RE) as required in DWIE.

4.1. Span-Based Representation

The input to our model consists of document-level annotation instances. Each document D from the considered document collection \mathcal{D} is represented by its sequence T of tokens. These tokens are represented internally as a concatenation of GloVe (Pennington et al., 2014) and character embeddings (Ma & Hovy, 2016). We also experiment with additionally concatenating BERT (Devlin et al., 2019) contextualized embeddings. Since BERT is run on a sub-token level, to the representation of each token we only concatenate the BERT-based representation of the first sub-token, as originally proposed by Devlin et al. (2019). This input is fed into a BiLSTM layer in order to obtain the output token representations by concatenating the forward and backward LSTM hidden states. The BiLSTM outputs for the considered document D are written on the token level as $\mathbf{e}_i \in \mathbb{R}^m$ ($i = 1, \dots, |T|$). These are converted into *span representations*. The set of all possible spans for D , up to maximum span width w_{\max} (which is a hyperparameter of the model), is written as $S = \{s_1, \dots, s_{|S|}\}$. The number of spans can be calculated as follows,

$$|S| = \sum_{k=1}^{w_{\max}} |T| - k + 1 = w_{\max} \left(|T| - \frac{w_{\max} - 1}{2} \right) \quad (2)$$

We obtain the representation \mathbf{g}_i^0 for span s_i , ranging from token l to token r , by concatenating their respective BiLSTM states \mathbf{e}_l and \mathbf{e}_r with an embedding ψ_{r-l} for the span width $w_i = r - l$

$$\mathbf{g}_i^0 = [\mathbf{e}_l; \mathbf{e}_r; \psi_{r-l}] \quad (3)$$

As seen from Eq. (2), the number of possible spans scales approximately linearly with the maximum span width w_{\max} , as well as the document length $|T|$ (assuming $w_{\max} \ll |T|$). This leads to a strongly increased set of spans, as compared to previous works where $|S|$ scales with the length of individual sentences rather than entire documents (Luan et al., 2019; Wadden et al., 2019). In order to mitigate the required memory of our model, we use a shared *pruner* to reduce S to a smaller set P of candidate spans to be used by

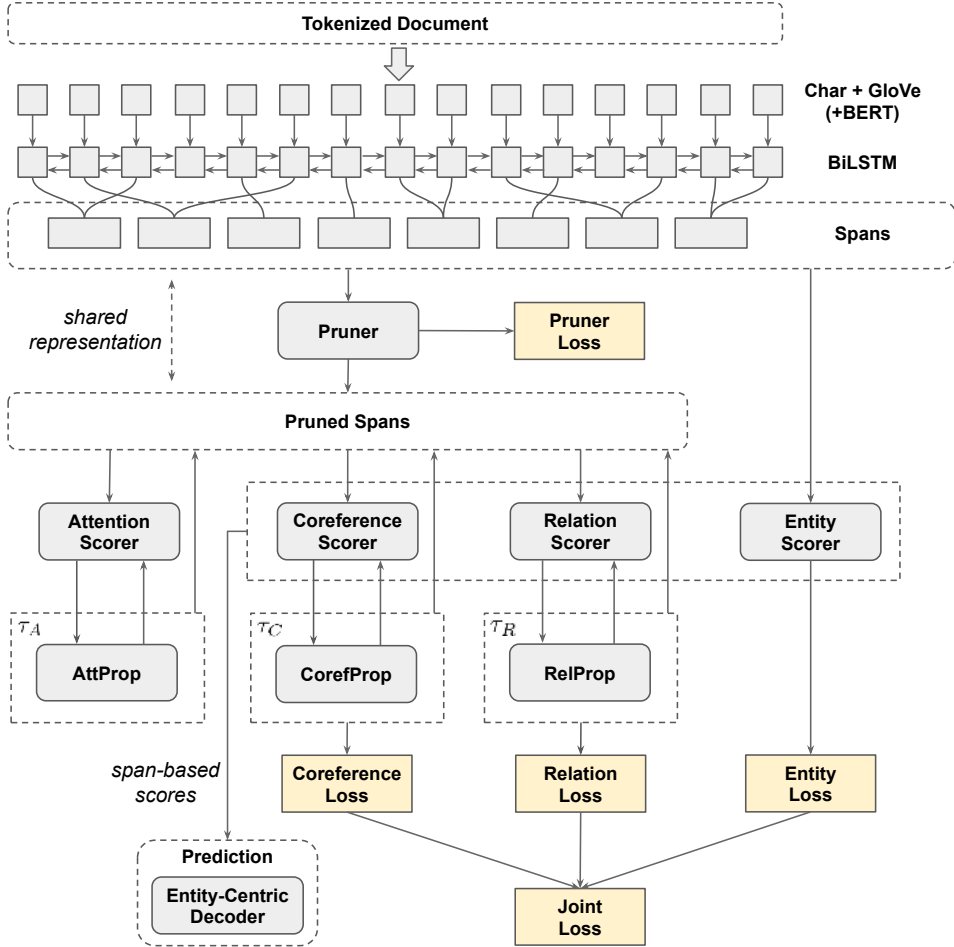


Figure 3: Architecture of our model; the span-oriented approach makes it possible to execute *coreference* (Section 4.2.2) and *relation* (Section 4.2.3) scorers independently from *entity scorer* (Section 4.2.1). However, a pruning step (described in Section 4.1) is needed in order to limit the memory required to perform matrix operations on span representations involved in graph propagation (**AttProp**, **CorefProp**, **RelProp**) (Section 4.4) as well as in the attention, coreference and relation scorer modules. The *pruned spans* share the same representation with the rest of the *spans* (*shared representation*). This way, the update in span representations caused by the graph propagation modules also affects the *entity scorer*. Our **AttProp** graph propagation method runs independently from coreference, relation, and entity scorers, enabling its use in combination with any task. Finally, the *entity-centric decoder* (Section 4.3) uses the entity clusters predicted by the *coreference scorer* to convert the span-based predictions from the *relation* and *entity* scorers to entity-centric ones.

the coreference and RE scorers and in the graph propagation modules (see further). The choice of using a single pruner contrasts with similar work in Luan et al. (2019) and Wadden et al. (2019) where two separate pruners are used, one for the relation task, and another for coreference. Our design choice is based on the fact that both of these tasks use the same document-level entity mentions. This contrasts with datasets used in Luan et al. (2019) and Wadden et al. (2019) where, while the coreference is defined on the document-level, the relations are sentence-based.

Finally, we use graph propagation to iteratively refine the pruned spans representations. Three graph propagation mechanisms are compared in the experiments. Our own contribution is the attention-based graph propagation method **AttProp**, where the span representations are updated in τ_A iterations. Alternatively, τ_C iterations of **CorefProp** (Lee et al., 2018; Luan et al., 2019) can be performed, or τ_R iterations of **RelProp** (Luan et al., 2019).

The span representation of a particular span s_i after iteration t is denoted as \mathbf{g}_i^t in our notation. The

details of graph propagation are explained in Section 4.4. Note that in theory several of these graph propagation techniques could be accumulated, but in our setting the benefits thereof in terms of model effectiveness were minor, at a significantly higher computational cost. Therefore, in our experiments, we only compare models without graph propagation with models applying a single form of graph propagation. To keep the sections introducing the models clear, we will write τ to denote the number of propagations in general (which could be 0, or any of τ_A, τ_C or τ_R , depending on the chosen experiments and considered model components).

4.2. Joint Model for Entity Recognition, Coreference Resolution, and Relation Extraction

In this section, we present the joint model including recognition of entity mentions as belonging to L_T types (introduced as NER), the clustering of the entity mentions into entities (coreference resolution), and identifying relations between entities, all on the document level. The building blocks responsible for the three subtasks are discussed next, as well as the total loss of the joint model. The details of the graph propagation mechanisms are then provided further on (Section 4.4).

4.2.1. Entity Mention Module

All spans s_i (up to width w_{\max}) of the considered document¹² are scored by feeding their representation (starting from Eq. (3) and potentially updated after τ graph propagation iterations) into the feed-forward neural network (FFNN) written as $\mathcal{F}_{\text{mention}}$, with as many outputs as there are entity types:

$$\Phi_{\text{mention}}^\tau(s_i) = \mathcal{F}_{\text{mention}}(\mathbf{g}_i^\tau). \quad (4)$$

Throughout this section, we will maintain the same notation of $\mathcal{F}(\mathbf{x})$ to denote a FFNN that takes as input a vector \mathbf{x} and produces a vector of scores, and $\mathcal{F}(x)$ to refer to a FFNN with a scalar output.

The probability of each label being valid for the considered span is modeled by component-wise application of a sigmoid ($\sigma(x) = 1/(1 + e^{-x})$) to these scores $\Phi_{\text{mention}}^\tau(s_i) \in \mathbb{R}^{L_T}$ (with L_T the number of entity tags). The log probability of the ground truth mention labels for all spans of document D is given by

$$\log P_{\text{mention}}(E^*|G^\tau) = \sum_{i=1}^{|S|} \sum_{l=1}^{L_T} I_{i,l} \log \sigma(\Phi_{\text{mention}}^\tau(s_i)_l) + (1 - I_{i,l}) \log (1 - \sigma(\Phi_{\text{mention}}^\tau(s_i)_l)), \quad (5)$$

in which E^* represents the set of ground truth mention labels for all spans in the document, and $I_{i,l} \in \{0, 1\}$ is the ground truth indicator label for mention tag l of span s_i . G^τ denotes the set of all considered span representations for the current document. The superscript τ reflects the fact that, in case graph propagation is applied, the subset of $|P|$ representations (for the spans retained after pruning) have been updated over τ iterations. By summing over all entity types ($l = 1, \dots, L_T$), we account for the fact that a particular span can have multiple associated entity tags (i.e., the considered NER task is multi-label). At inference time, spans get assigned those entity types for which the corresponding score $\Phi_{\text{mention}}^\tau(s_i) > 0$. Note that not all valid entity mentions necessarily get an entity type assigned: if the relation extractor determines that a span is part of a relation, it effectively becomes an entity mention, even if none of the pre-defined types is considered applicable by the entity scorer.

4.2.2. Coreference Module

While the entity scoring is performed on all span representations S , this is not possible for the coreference and relation scorers, due to memory limitations. The latter scorers predict on *pruned spans*, as shown in Fig. 3. How the pruner is trained jointly with the model, is described in Section 4.2.4. In order to avoid confusion by introducing additional notations, we list the spans in the pruned set P as $s_1, \dots, s_{|P|}$, according to their original order in the text.

¹²For convenience, the subscript D indicating the current document is left out in the equations of this section.

The module for coreference resolution is based on pairwise scoring of the pruned spans from P . Following ideas from Lee et al. (2017, 2018); Luan et al. (2018, 2019), for any span s_j , scores with respect to each of the preceding (also referred to as ‘antecedent’) spans s_i ($i \leq j$) in the document are calculated with a neural network $\mathcal{F}_{\text{coref}}$:

$$\Phi_{\text{coref}}^\tau(s_i, s_j) = \mathcal{F}_{\text{coref}}([\mathbf{g}_i^\tau; \mathbf{g}_j^\tau; \mathbf{g}_i^\tau \odot \mathbf{g}_j^\tau; \varphi_{i,j}]). \quad (6)$$

This expression scores the compatibility between spans s_i and s_j , taking as input the concatenation of their respective span representations (after τ propagation iterations), their component-wise product, and an embedding $\varphi_{i,j}$ representing their distance in terms of the number of ordered candidate spans from s_i to s_j .

In order to deal with non-coreferent or incorrect spans, previous work in span-based coreference (Lee et al., 2017, 2018) defines a dummy antecedent ϵ to which all non-coreferent or invalid spans point. While this approach is effective in datasets that do not contain singleton entity clusters, such as OntoNotes-based CoNLL-2012 (Pradhan et al., 2012), it does not allow to distinguish between valid singleton entity mentions and invalid mention spans. This makes it unsuitable to use on DWIE, since it contains singleton entity clusters, consisting of a single mention. In fact, 66.4% of the entity clusters in DWIE are singletons. Furthermore, the current official CoNLL-2012 evaluation script¹³ based on Pradhan et al. (2014) accounts for scenarios where either the dataset or the predicted mentions are singletons, which has a direct impact on the established B-CUBED (Bagga & Baldwin, 1998) and CEAF_e (Luo, 2005) coreference scores. In order to tackle the singleton entity cluster detection in our coreference model, we propose to start from $\Phi_{\text{coref}}^\tau(s_j, s_j)$ ¹⁴ as a self-coreference span score. By applying the correct target in the coreference loss, it allows indicating that either the span s_j is not a valid mention, or that it is a valid mention that is not co-referenced with any antecedent span.

The log probability of the ground truth coreference labels of document D is given by

$$\log P_{\text{coref}}(C^* | G^\tau) = \sum_{j=1}^{|P|} \log \frac{\sum_{s^* \in S_j^*} \exp(\Phi_{\text{coref}}^\tau(s^*, s_j))}{\sum_{i=1}^j \exp(\Phi_{\text{coref}}^\tau(s_i, s_j))}. \quad (7)$$

The set of ground truth coreference labels is indicated as C^* . The summation over j represents the contribution to the log likelihood of the correct antecedent labels for each span s_j in the pruned set P . The individual terms in the right-hand side correspond to the log probability of the correct antecedent labels for a particular span s_j . In the denominator, the summation ranges from the first span, up to span s_j itself (i.e., for the self-coreference score), but not beyond it (given that only *antecedents* in the sorted sequence of pruned spans are considered). The numerator contains the contributions from the potentially multiple ground truth antecedents for span s_j . This stems from the fact that multiple antecedent mentions may belong to the same cluster as s_j , which all contribute to the probability of the correct antecedent labels. The set of ground truth antecedents corresponding to span s_j is written S_j^* .

At inference time, the highest scoring antecedent for span s_j (including s_j itself) is picked. Due to the idea of only predicting *antecedents*, picking any of the ground truth antecedents leads to the correct mention clusters (Durrett & Klein, 2013; Lee et al., 2017, 2018; Wiseman et al., 2015).

4.2.3. Relation Module

Similar to the coreference module (Eq. (6)), we score span pairs using an FFNN

$$\Phi_{\text{relation}}^\tau(s_i, s_j) = \mathcal{F}_{\text{relation}}([\mathbf{g}_i^\tau; \mathbf{g}_j^\tau; \mathbf{g}_i^\tau \odot \mathbf{g}_j^\tau; \varphi_{i,j}]), \quad (8)$$

where $\varphi_{i,j}$ is again the distance embedding as introduced in Section 4.2.2. $\Phi_{\text{relation}}^\tau(i, j) \in \mathbb{R}^{L_R}$ is a vector representing relation span pair scores for each of the L_R possible relation types between spans s_i and s_j .

¹³<https://github.com/conll/reference-coreference-scorers>

¹⁴This would be replaced with $\Phi_{\text{coref}}^\tau(\epsilon, s_j)$ in the *dummy-based* formulation defined in Lee et al. (2017).

The log probability of the ground truth relation labels of document D is given by

$$\log P_{\text{relation}}(R^*|G^\tau) = \sum_{i,j=1}^{|P|} \sum_{l=1}^{L_R} I_{i,j,l} \log \sigma(\Phi_{\text{relation}}^\tau(s_i, s_j)_l) + (1 - I_{i,j,l}) \log(1 - \sigma(\Phi_{\text{relation}}^\tau(s_i, s_j)_l)), \quad (9)$$

in which R^* represents the set of ground truth relation labels for all combination of pruned span pairs in the document, and $I_{i,j,l} \in \{0, 1\}$ is the ground truth indicator label for relation type l of the span pair (s_i, s_j) . Note that all $|P|^2$ pruned span pairs are considered, since the order of the spans in the relation matters (unlike the coreference case). By summing over all possible relation types L^R , we account for the fact that a particular relation between two spans can be multi-label (which is the case for more than 30% of relations, as shown in Table A.5).

Since this model is run in parallel with the coreference module, it is used to predict relations only between entity mentions and not entity clusters. During inference, candidate relations are accepted when $\Phi_{\text{relation}}^\tau(s_i, s_j)_l > 0$.

4.2.4. Span Pruner

The span pruner is an FFNN, denoted $\mathcal{F}_{\text{pruner}}$, that scores all spans s_i based on their initial representation \mathbf{g}_i^0 , after which only the highest scoring spans are retained in the pruned span set P . In our experiments P contains the top $0.2|T|$ highest scoring spans, which covers more than 98% of all the ground truth mention spans in the DWIE dataset. We represent the pruner score for span s_i as

$$\Phi_{\text{pruner}}(s_i) = \mathcal{F}_{\text{pruner}}(\mathbf{g}_i^0). \quad (10)$$

Several strategies can be used to train the pruner. One option is to directly optimize the probability of the pruner to detect the spans of correct entity mentions. With S^* the set of spans with at least one ground truth entity type, and $I_i \in \{0, 1\}$ an indicator for whether $s_i \in S^*$, the corresponding log likelihood can be written as

$$\log P_{\text{pruner}}(S^*|G^0) = \sum_{i=1}^{|S|} I_i \log \sigma(\Phi_{\text{pruner}}(s_i)) + (1 - I_i) \log(1 - \sigma(\Phi_{\text{pruner}}(s_i))), \quad (11)$$

leading to a separate pruner loss term. Alternatively, the pruner can be trained indirectly by adapting the mention score from Eq. (4), the coreference score from Eq. (6) or the relation score from Eq. (8) as follows:

$$\tilde{\Phi}_{\text{mention}}^\tau(s_i) = \Phi_{\text{mention}}^\tau(s_i) + \Phi_{\text{pruner}}(s_i) \quad (12)$$

$$\tilde{\Phi}_{\text{coref}}^\tau(s_i, s_j) = \Phi_{\text{coref}}^\tau(s_i, s_j) + \Phi_{\text{pruner}}(s_i) \quad (13)$$

$$\tilde{\Phi}_{\text{relation}}^\tau(s_i, s_j) = \Phi_{\text{relation}}^\tau(s_i, s_j) + \Phi_{\text{pruner}}(s_i) \quad (14)$$

for use in the expressions Eq. (5), Eq. (7) and Eq. (9), respectively. As such, higher pruner scores would directly correspond to higher mention or coreference scores, and lead to a meaningful ranking of spans according to pruner scores. All three strategies seem to work on a similar level, but for the presented joint model experiments, we use the indirect training through the coreference module, as in Eq. (13). Note that we did not experiment with training the pruner through the relation module, because it would be trained only on those spans involved in relations, which is a mere subset of all valid mentions.

4.2.5. Joint Model

We perform joint training in order to explore the degree to which the graph propagation techniques (see Section 4.4) affect related tasks in DWIE. For instance, we expect that performing a coreference propagation can have a positive impact on the NER task. We hypothesize that enriching the entity spans with broader contextual information coming from other mention spans in the cluster, can improve the effectiveness of the entity module. Furthermore, given the entity-centric nature of DWIE, the mention-based predictions

for NER and RE have to be grouped in coreference clusters (see section 4.3 for details), which makes it necessary to execute these tasks jointly with the coreference task.

The joint loss for each document D is a weighted sum of the individual loss functions of the subtasks:

$$\mathcal{L}_D^{\text{joint}} = \sum_{(E^*, C^*, R^*)} \lambda_E \log P_{\text{mention}}(E^* | G^T) + \lambda_C \log P_{\text{coref}}(C^* | G^T) + \lambda_R \log P_{\text{relation}}(R^* | G^T), \quad (15)$$

in which λ_E , λ_C , and λ_R are hyperparameters of the joint model.

4.3. Decoding and Prediction

Unlike previous datasets used in span-based predictions (Doddington et al., 2004; Kulkarni et al., 2018; Luan et al., 2018; Walker et al., 2006) where the relation and entity extraction are done on the mention-level, DWIE is an entity-centric dataset. During inference, this requires an additional decoding step to cluster the mention-based span-dependent predictions into entity-centric ones. The component responsible for this decoding in the proposed architecture is the *entity-centric decoder* (see Fig. 3). The pseudo-code in Algorithm 1 summarizes the steps performed by this component. First, the decoder receives as *input* the predicted span clusters (p_cl), entity mentions (p_men) and relations between spans (p_rel) obtained from the scores calculated in Eq. (13), Eq. (4) and Eq. (8), respectively. Next, the predicted entity mentions are connected with the respective clusters by using the dictionary C that maps mention spans to cluster ids (lines 3–12 in Algorithm 1). Specifically, each of the entity clusters is assigned the union of the entity types predicted for any of the mention spans inside the cluster (line 11 in Algorithm 1). If the predicted entity mention can not be located inside the predicted clusters, a new singleton cluster is added (lines 5–6 in Algorithm 1). Finally, all the pairwise predicted relations on the mention level (p_rel) between members of two different clusters are assigned as predicted relations between the (cluster-level) entities (lines 13–20 in Algorithm 1). Similarly as with entity mentions, the dictionary C is used to map the mention spans ($span_h$ and $span_t$) of a particular relation type rel_type to the corresponding cluster ids. Furthermore, the relations added between two clusters are the union of all the relations predicted between any pair of mentions inside these clusters (line 18 in Algorithm 1).

4.4. Graph Propagation Mechanisms

In order to evaluate the impact of graph-based propagation of contextual information between the spans, we propose **AttProp**, and reimplement the **CorefProp** and **RelProp** graph propagation algorithms. Lee et al. (2018) proposed the gated graph propagation update function for use on coreference resolution, which was then successfully applied in a joint multi-task setting by Luan et al. (2019); Wadden et al. (2019). The graph propagation equations are written as:

$$\mathbf{f}_x^t(s_i) = \sigma(\mathcal{F}_x([\mathbf{g}_i^t; \mathbf{u}_x^t(s_i)])), \quad (16)$$

$$\mathbf{g}_i^{t+1} = \mathbf{f}_x^t(s_i) \odot \mathbf{g}_i^t + (1 - \mathbf{f}_x^t(s_i)) \odot \mathbf{u}_x^t(s_i), \quad (17)$$

where in our case $x \in \{A, C, R\}$ denotes **AttProp**, **CorefProp**, and **RelProp**, respectively. The n -dimensional vector $\mathbf{f}_x^t(s_i)$, produced by the single-layer FFNN \mathcal{F}_x can be interpreted as a gating vector that acts as a switch between the current span representations $\mathbf{g}_i^t \in \mathbb{R}^n$, and the update span vector $\mathbf{u}_x^t(s_i) \in \mathbb{R}^n$. The various graph propagation methods differ in how $\mathbf{u}_x^t(s_i)$ is calculated.

CorefProp — The coreference confidence score between span s_i and s_j for propagation iteration t is denoted as $P_C^t(s_i, s_j)$ and calculated as follows,

$$P_C^t(s_i, s_j) = \frac{\exp(\tilde{\Phi}_{\text{coref}}^t(s_i, s_j))}{\sum_{i'=1}^j \exp(\tilde{\Phi}_{\text{coref}}^t(s_{i'}, s_j))}, \quad (18)$$

in which $i' \in \{1, \dots, j\}$ refers to all antecedent spans $s_{i'}$ to span s_j in the pruned span set. Note that the coreference scores according to Eq. (13) are used. This means the confidence scores not only reflect whether

Algorithm 1 Entity-centric decoder for the *Joint* model.

Input: predicted clusters (p_cl), entity mentions (p_men) and relations between mentions (p_rel):

1. p_cl is a dictionary (map) that maps cluster ids to mention spans
2. p_men is list of tuples \langle predicted span, predicted tag \rangle
3. p_rel is list of tuples \langle predicted head span, predicted relation, predicted tail span \rangle

Output: clusters (p_cl), decoded entities (d_ent) and relations between entities (d_rel)

```
1: Initialize  $d\_ent, d\_rel \leftarrow$  empty dictionary (map)
2:  $C \leftarrow$  transformed  $p\_cl$  that maps spans to cluster ids
    $\triangleright$  Decode entity mentions ( $p\_men$ ) to entities ( $d\_ent$ ) (lines 3–12)
3: for  $span, tag$  in  $p\_men$  do
4:   if  $span$  not in  $C.keys()$  then
5:      $C[span] \leftarrow$  new concept id
6:      $p\_cl[C[span]] \leftarrow \text{list}([span])$ 
7:   end if
8:   if  $C[span]$  not in  $d\_ent.keys()$  then
9:      $d\_ent[C[span]] \leftarrow$  empty set
10:  end if
11:   $d\_ent[C[span]].add(tag)$ 
12: end for
    $\triangleright$  Decode relations between mentions ( $p\_rel$ ) to relations between entities ( $d\_rel$ ) (lines 13–20)
13: for  $span\_h, rel\_type, span\_t$  in  $p\_rel$  do
14:   if ( $span\_h$  in  $C.keys()$ ) and ( $span\_t$  in  $C.keys()$ ) then
15:     if  $\langle C[span\_h], C[span\_t] \rangle$  not in  $d\_rel.keys()$  then
16:        $d\_rel[\langle C[span\_h], C[span\_t] \rangle] \leftarrow$  empty set
17:     end if
18:      $d\_rel[\langle C[span\_h], C[span\_t] \rangle].add(rel\_type)$ 
19:   end if
20: end for
```

the considered spans are compatible, but also whether the individual spans are likely to be retained by the pruner as potential entity mentions. In order to perform a **CorefProp** graph iteration, the span update vector $\mathbf{u}_C^t(i) \in \mathbb{R}^n$ is first calculated as a weighted average of the current representation of span s_j and all of its antecedents

$$\mathbf{u}_C^t(s_j) = \sum_{i=1}^j P_C^t(s_i, s_j) \mathbf{g}_i^t, \quad (19)$$

in which the weighting coefficients quantify the coreference compatibility of the corresponding span with s_j . After that, the update equations Eq. (16) and Eq. (17) are applied.

RelProp — Similarly as with **CorefProp**, a relation span update vector is calculated as formalized next,

$$\mathbf{u}_R^t(s_j) = \sum_{i=1}^{|P|} (\mathbf{A}_R f(\Phi_{\text{relation}}^t(s_i, s_j))) \odot \mathbf{g}_i^t, \quad (20)$$

where $\mathbf{A}_R \in \mathbb{R}^{n \times L_R}$ is a trainable projection tensor, and f is a non-linear activation function (ReLU). Similarly as in Eq. (19), the update vector can be interpreted as a weighted sum of all span representations, with the additional expressiveness stemming from the projection matrix \mathbf{A}_R in accounting for the relation scores.

AttProp — In order to measure the impact of the ‘supervised’ **CorefProp** and **RelProp** propagation techniques described by equations (18)-(20) above, we introduce a latent attentive propagation. Unlike **CorefProp** and **RelProp** that are driven by the task-specific confidence propagation scores $\Phi_{\text{coref}}^t(s_i, s_j)$ and $\Phi_{\text{relation}}^t(s_i, s_j)$, **AttProp** is influenced only by latent attention weights between all the pruned spans P calculated as follows,

$$\Phi_{\text{att}}^t(s_i, s_j) = \mathcal{F}_{\text{att}}([\mathbf{g}_i^t; \mathbf{g}_j^t; \mathbf{g}_i^t \odot \mathbf{g}_j^t; \varphi_{i,j}]), \quad (21)$$

where $\varphi_{i,j}$ is the distance feature embedding function between spans s_i and s_j , and $\Phi_{\text{att}}^t(s_i, s_j)$ is the attention score between these spans. This score is normalized with a softmax to get the $P_A^t(s_i, s_j)$ confidence score

$$P_A^t(s_i, s_j) = \frac{\exp(\Phi_{\text{att}}^t(s_i, s_j))}{\sum_{j'=1}^{|P|} \exp(\Phi_{\text{att}}^t(s_i, s_{j'}))}. \quad (22)$$

The span update vector $\mathbf{u}_A^t(s_i) \in \mathbb{R}^n$ is calculated as a weighted sum of all the P span representations as opposed to only antecedents in **CorefProp**

$$\mathbf{u}_A^t(s_i) = \sum_{j=1}^{|P|} P_A^t(s_i, s_j) \mathbf{g}_j^t. \quad (23)$$

4.5. Single Task Models

In this section we shortly describe independent baseline models for the three individual core tasks under study in this paper, as training these models not entirely corresponds to merely minimizing the corresponding loss term from the total loss Eq. (15).

4.5.1. Single Entity Recognition Model

The single-task NER model is designed for detecting and correctly labeling the individual entity spans, and is based on Eq. (5). However, even for the single models, the graph propagation mechanism **AttProp** may be useful, but for that the pruner needs to be jointly trained with the model. This is obtained by augmenting the mention loss $-\log P_{\text{mention}}(E^*|G^T)$ with the pruner loss $-\log P_{\text{pruner}}(S^*|G^0)$ according to Eq. (11).

4.5.2. Single Coreference Resolution Model

The single-task end-to-end coreference model needs to detect mentions and correctly cluster them. Here again, the standard coreference loss $-\log P_{\text{coref}}(C^*|G^\tau)$ according to Eq. (13) and Eq. (7) is extended with the pruner loss $-\log P_{\text{pruner}}(S^*|G^0)$. This turned out essential for correctly predicting the singleton clusters.

4.5.3. Single Relation Extraction Model

The single relation extraction model is trained to detect mentions as well as the correct pairwise relations between mentions (i.e., without the coreference step). In order to train the pruner as well, the standard relation score is extended as described in Eq. (14) before calculating the loss $-\log P_{\text{relation}}(R^*|G^\tau)$ based on Eq. (9).

5. Entity-Centric Metrics

Unlike the currently widespread datasets that use a mention-driven approach to annotate named entities (Bekoulis et al., 2017; Derczynski et al., 2017; Sang & De Meulder, 2003; Weischedel et al., 2011), relations (Augenstein et al., 2017; Bekoulis et al., 2017; Doddington et al., 2004; Ji et al., 2017; Kim et al., 2003; Luan et al., 2018; Song et al., 2015) and entity linking (Bentivogli et al., 2010; Hoffart et al., 2011; Riedel et al., 2010), DWIE is entirely entity-centric. As explained before, we group entity mentions s_i referring to the same entity into clusters C_k . While we can, and will, adopt the traditional coreference measures as defined by Pradhan et al. (2014) to judge this cluster formation, the NER and relation extraction (RE) evaluation (using precision, recall and F_1) can be done either on (i) mention level, or (ii) entity (cluster) level. The first option however would have the metrics being dominated by the more frequently occurring entities, while the second would penalize mistakes in the clustering (since partially correctly identified clusters would be seen as completely incorrect). This is illustrated in Fig. 4 and the corresponding performance metrics in Table 6, where scenarios 1 and 2 highlight the effect of making labeling mistakes on the cluster level for different sizes, and scenario 3 highlights the pessimistic view of hard entity-level metrics in case of clustering mistakes. Note that we indicate the mention-level metrics with subscript m , while the (hard) entity-level metrics will have subscript with e .

Because the (hard) entity-level metrics in our opinion overly penalize clustering mistakes (cf. scenario 3), we propose a variant of entity-level evaluation which we term *soft* entity-level metrics (denoted by subscript s). Basically, instead of adopting a binary count of 1 (all mentions correct) or 0 (as soon as a single mention is missed) on an entity cluster level, we rather count the fraction of its mentions that are correctly labeled. This is illustrated in the formula part of Fig. 4(a) for NER, and below we present the adopted formulas in detail. Note that in case clusters are completely predicted correctly, the soft entity-level metrics are the same as hard entity-level metrics (and thus avoid the metric being dominated by frequent mentions, as in the mention-level case).

The formal definition of the metrics depends on counting true positives $tp_p(l)$ and $tp_g(l)$, false positives $fp(l)$, and false negatives $fn(l)$ for a particular NER tag/relation type l , which are specified in Eq. (24)–(25). These and other notation definitions are summarized in Table 8. Further, note that we define two true positives for a particular label l , because of the potential difference between predicted and ground truth clusters: $tp_p(l)$ sums fractions of *predicted* clusters and is used to calculate the precision Pr_s in Eq. (26), while $tp_g(l)$ considers *ground truth* clusters and is used for the recall Re_s in Eq. (26). This allows us to preserve the *cluster-based* relationships between true positives, false positives and false negatives as described for expressions $tp_p(l) + fp(l)$ and $tp_g(l) + fn(l)$ in Table 7. Thus our soft entity-level metrics are still cluster-based, while accounting for the mention-level predictions.

$$tp_p(l) = \sum_{C_p \in P_C(l)} \frac{|C_p \cap G_M(l)|}{|C_p|}, \quad tp_g(l) = \sum_{C_g \in G_C(l)} \frac{|C_g \cap P_M(l)|}{|C_g|} \quad (24)$$

$$fp(l) = |P_C(l)| - tp_p(l), \quad fn(l) = |G_C(l)| - tp_g(l) \quad (25)$$

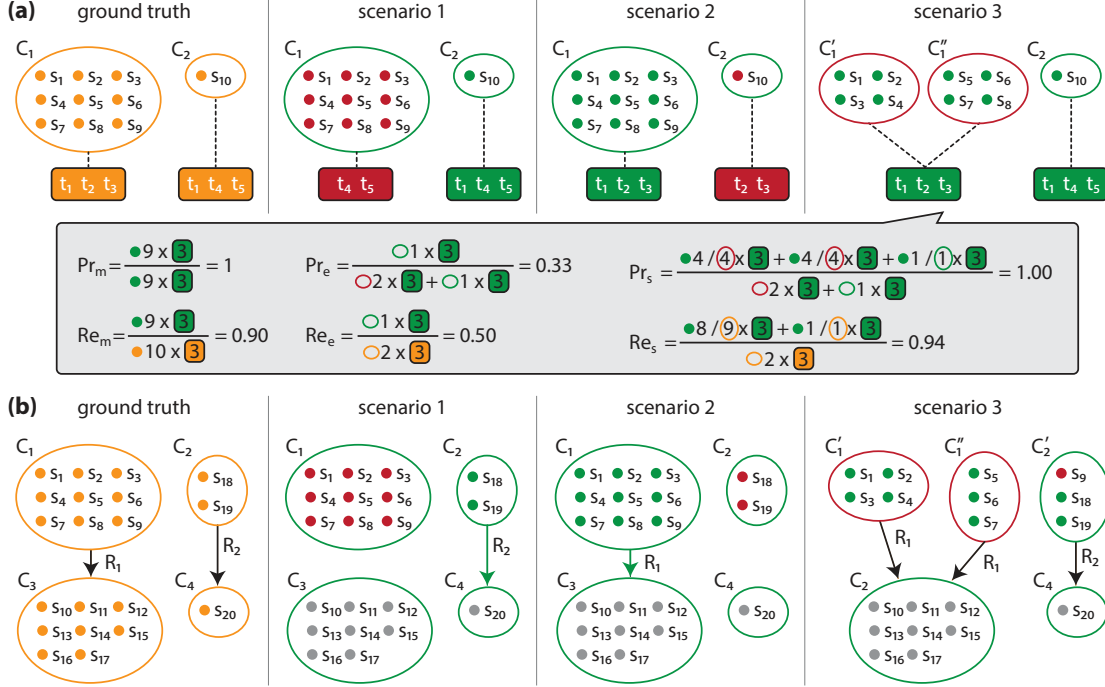


Figure 4: Illustration of entity prediction scenarios for (a) NER and (b) relation extraction, with large clusters (C_1, C_3) and smaller ones (C_2, C_4). *Scenario 1* erroneously labels the large one, *scenario 2* incorrectly labels the small one, *scenario 3* incorrectly splits up the large one and makes a mistake for one of its mentions, s_9 . The formulas in the grey box illustrate the calculation of mention-level (Pr_m, Re_m), hard entity-level (Pr_e, Re_e) and soft entity-level (Pr_s, Re_s) precision and recall for NER in *scenario 3*. Note that in (b), the mention dots are colored for correct (green) and incorrect (red) relation heads only.

Table 6: Comparison of different metrics for the example scenarios depicted in Fig. 4, for (a) NER and (b) relation extraction.

		Mention-Level			Hard Entity-Level			Soft Entity-Level		
		Pr_m	Re_m	$F_{1,m}$	Pr_s	Re_s	$F_{1,s}$	Pr_e	Re_e	$F_{1,e}$
(a) NER	Ground Truth	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	Scenario 1	0.143	0.100	0.118	0.600	0.500	0.545	0.600	0.500	0.545
	Scenario 2	0.931	0.900	0.915	0.600	0.500	0.545	0.600	0.500	0.545
	Scenario 3	1.000	0.900	0.947	0.333	0.500	0.400	1.000	0.944	0.971
(b) RE	Ground Truth	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
	Scenario 1	1.000	0.027	0.053	1.000	0.500	0.667	1.000	0.500	0.667
	Scenario 2	1.000	0.973	0.986	1.000	0.500	0.667	1.000	0.500	0.667
	Scenario 3	0.983	0.783	0.872	0.000	0.000	0.000	0.889	0.889	0.889

Table 7: The relations between the weighted true positives by the size of predicted ($tp_p(l)$) and ground truth ($tp_g(l)$) entity clusters allows us to achieve the constraints needed for the denominators of precision ($tp_p(l) + fp(l)$) and recall ($tp_p(l) + fn(l)$) functions (Eq. (26)) in terms of the number of entity clusters.

Expression	(a) Meaning for NER	(b) Meaning for RE
$tp_p(l) + fp(l)$	Number of <i>predicted</i> entity clusters with tag l .	Number of <i>predicted</i> relations of type l between entity clusters.
$tp_g(l) + fn(l)$	Number of <i>ground truth</i> entity clusters with tag l .	Number of <i>ground truth</i> relations of type l between entity clusters.

Our soft entity-level precision, recall and F_1 metrics are formally defined as follows, where L refers to either the number of all possible tags for NER or the number of all possible relation types for RE:

$$\text{Pr}_s = \frac{\sum_{l=1}^L tp_p(l)}{\sum_{l=1}^L tp_p(l) + fp(l)}, \quad \text{Re}_s = \frac{\sum_{l=1}^L tp_g(l)}{\sum_{l=1}^L tp_g(l) + fn(l)}, \quad F_{1,s} = \frac{2 \cdot \text{Pr}_s \cdot \text{Re}_s}{\text{Pr}_s + \text{Re}_s} \quad (26)$$

6. Experimental results

6.1. Experimental Setup

We train and evaluate our model as described in Section 4 on three tasks: NER, coreference, and relation extraction (RE) independently and jointly. We experiment with three main model variations:

1. **Single:** Experiments on individual tasks by training with the respective loss functions as described in Section 4.5.
2. **Joint:** Experiments jointly on all three tasks using pre-trained *GloVe representations*¹⁵ concatenated to character embeddings in the shared input layer (see Fig. 3). For training we use the joint loss defined in Section 4.2.
3. **Joint+BERT:** as in the *Joint* setting, experiments jointly on all three tasks, but using pre-trained BERT_{BASE} embeddings¹⁶ concatenated to the GloVe and character embeddings. We use an input window size of 250 tokens and concatenate the last 2 hidden layers of BERT to get token representations.

Additionally, for each of the three model setups we experiment with the graph propagation techniques defined in Section 4.4. To maximize result consistency, we train each model 5 times and report the average of these 5 results for each of the experiments.

We use a single-layer BiLSTM with forward and backward hidden states of 200 dimensions each. All our FFNNs used to obtain confidence scores ($\mathcal{F}_{\text{pruner}}$, $\mathcal{F}_{\text{coref}}$, $\mathcal{F}_{\text{mention}}$, $\mathcal{F}_{\text{relation}}$, and \mathcal{F}_{att}) have two 150-dimensional hidden layers trained with a dropout of 0.4. We set the maximum span width w_{max} to 5 and the pruner ratio to 0.2 of the total number of tokens in a document. For training, we use Adam with a learning rate of $1e-3$ for 100 epochs with a linear decay of 0.1 starting at epoch 15.

¹⁵<http://nlp.stanford.edu/data/glove.840B.300d.zip>

¹⁶https://storage.googleapis.com/bert_models/2018.10.18/cased_L-12_H-768_A-12.zip

Table 8: Short definition of the symbols and expressions involved in our *soft-entity level* metric formulation in Eq. (24)–(26) for both NER and RE tasks.

Symbol	(a) Meaning for NER	(b) Meaning for RE
$P_C(l)$	Set of predicted entity clusters with tag l .	Set of predicted relations of type l between the predicted entity clusters.
$C_p \in P_C(l)$	Set of predicted entity mentions for a particular entity cluster in $P_C(l)$.	Set of relations between the predicted entity mentions for a particular pair of related entity clusters in $P_C(l)$.
$G_C(l)$	Set of ground truth entity clusters annotated with tag l .	Set of ground truth relations of type l between the ground truth entity clusters.
$C_g \in G_C(l)$	Set of ground truth entity mentions for a particular entity cluster in $G_C(l)$.	Set of relations between the ground truth entity mentions for a particular pair of related entity clusters in $G_C(l)$.
$P_M(l)$	Set of predicted entity mentions with tag l .	Set of predicted relations of type l between the predicted entity mentions.
$G_M(l)$	Set of ground truth entity mentions annotated with tag l .	Set of ground truth relations of type l between the ground truth entity mentions.
$tp_p(l)$	Number of true positive predictions of tag l on mentions re-weighted by predicted cluster sizes.	Number of true positive predictions of relation type l between mentions re-weighted by the number of mention level relations between the connected pairs of predicted clusters.
$tp_g(l)$	Number of true positive mention level predictions of tag l re-weighted by ground truth cluster sizes.	Number of true positive predictions of relation type l between mentions re-weighted by the number of mention level relations between the connected pairs of ground truth clusters.
$fp(l)$	Number of false positive mention level predictions of tag l re-weighted by predicted cluster sizes.	Number of false positive predictions of relation type l between mentions re-weighted by the number of mention level relations between the connected pairs of predicted clusters.
$fn(l)$	Number of false negative mentions with ground truth tag l re-weighted by ground truth cluster sizes.	Number of false negative relations of type l between mentions re-weighted by the number of mention level relations between the connected pairs of ground truth clusters.

Table 9: Main results of the experiments grouped in three model setups: (i) *Single* models trained individually, (ii) *Joint* model trained using as input GloVe and character embeddings, and (iii) *Joint+BERT* model trained on BERT_{BASE} embeddings. To report the results, we use MUC, CEAF_e, B³ as well as the average (Avg.) of these three metrics for *coreference resolution*. For NER and RE we use mention-level (F_{1,m}), hard entity-level (F_{1,e}), and soft entity-level (F_{1,s}) metrics described in Section 5. In bold we mark the best results for each model setup, the best overall results are underlined. Note that the metrics are expressed in percentage points.

Model Setup	Coreference F ₁				NER F ₁			RE F ₁		
	MUC	CEAF _e	B ³	Avg.	F _{1,m}	F _{1,e}	F _{1,s}	F _{1,m}	F _{1,e}	F _{1,s}
Single	92.8	90.9	88.2	90.6	85.7	-	-	68.2	-	-
+AttProp	93.2	91.5	88.7	91.1	87.1	-	-	71.3	-	-
+CorefProp	92.8	90.9	88.3	90.7	-	-	-	-	-	-
+RelProp	-	-	-	-	-	-	-	68.2	-	-
Joint	92.5	90.5	87.3	90.1	85.4	71.7	84.4	68.1	46.8	66.5
+AttProp	92.3	90.4	87.3	90.0	87.1	72.9	86.1	72.1	50.4	72.1
+CorefProp	92.3	90.3	87.2	89.9	87.2	73.2	86.0	71.6	50.2	71.0
+RelProp	92.6	90.2	86.8	89.9	86.7	72.4	85.2	69.5	48.2	68.8
Joint+BERT	93.8	92.1	89.0	91.6	87.6	74.2	86.4	70.6	48.7	68.9
+AttProp	93.2	91.4	88.6	91.1	88.8	74.2	87.7	72.3	50.4	73.0
+CorefProp	93.5	91.8	88.7	91.3	88.7	74.4	87.4	72.7	50.0	71.9
+RelProp	93.7	91.8	88.7	91.4	88.4	74.8	87.0	72.0	49.9	71.4

6.2. Results and Analyses

Table 9 gives an overview of the results achieved in *Single* as well as *Joint* and *Joint + BERT* setups. Additionally, Fig. 5 illustrates the impact of the number of graph propagation iterations for each of the span graph propagation methods on the final results.

First, we observe a general improvement in all our *Single* tasks when using graph propagation techniques. More specifically, our proposed latent **AttProp** achieves superior results compared to the relation (**RelProp**) and coreference (**CorefProp**) propagations when added to the *Single* setup. The biggest improvement across iterations (see Fig. 5) is for the single RE task mention-level F_{1,m} score with a boost of ~ 3 percentage points when incorporating **AttProp**. We also observe an improvement of ~ 1.5 percentage points in F_{1,m} for the NER task and a consistent but smaller improvement of 0.5 F₁ percentage points for the coreference task. These results illustrate the effectiveness of **AttProp** when applied to single task models.

A further improvement in results is achieved by training our model *jointly* (see the *Joint* setup in Table 9 and graphs in Fig. 5) for NER and RE tasks. This illustrates that, besides the positive effect of neural graph propagation on single task models, training our model jointly has an additional benefit by exploiting the interaction between tasks. In particular, this effect can be seen for RE, where our *Joint* model achieves a boost in performance of 0.8 percentage points for the mention-level F_{1,m} metric compared to the best result for the *Single* setup. Furthermore, our **AttProp** graph propagation method achieves the best performance on all the metrics for the RE task in the *Joint* setting with up to ~ 5.5 percentage points improvement in our newly proposed F_{1,s} metric. Additionally, we observe a beneficial effect of graph propagation for the NER task in the *Joint* setup with slightly better results for the F_{1,m} metric compared to the *Single* setting. Our **AttProp** technique performs on par with **CorefProp**, outperforming the latter by a small margin in terms of F_{1,s} metric.

Similarly to the *Joint* model variation, we observe benefits when using graph propagation techniques in the *Joint+BERT* models. Table 10 illustrates the deltas in performance for the NER and relation extraction tasks. This way, we can see more clearly the difference in impact of our neural message passing methods grouped by the model setup and metric type. First, we observe that the general performance boost from using graph propagation techniques is lower in *Joint+BERT* than in the *Joint* setup. We hypothesize that this effect is due to the fact that BERT itself has a better long-range context extraction due to the

Table 10: Deltas of improvement in performance for each of the graph propagation methods (**AttProp**, **CorefProp**, **RelProp**) in F₁ scores for (a) NER and (b) relation extraction tasks.

		Joint			Joint+BERT		
		F _{1,m}	F _{1,e}	F _{1,s}	F _{1,m}	F _{1,e}	F _{1,s}
(a) NER	Δ AttProp	1.69	1.18	1.67	1.16	-0.02	1.31
	Δ CorefProp	1.78	1.50	1.54	1.05	0.20	1.02
	Δ RelProp	1.33	0.70	0.75	0.78	0.56	0.60
(b) RE	Δ AttProp	3.97	3.62	5.56	1.66	1.69	4.05
	Δ CorefProp	3.48	3.45	4.47	2.02	1.29	2.95
	Δ RelProp	1.35	1.47	2.32	1.37	1.20	2.48

attention-based mechanism, which spans the input window as opposed to purely local (non-contextualized) GloVe embeddings used in the *Joint* setting. This is in line with the findings in Han & Wang (2020), Wadden et al. (2019), and Wu & He (2019) that show the advantage of using large BERT input window sizes to produce better IE results. Second, we observe that our **AttProp** method achieves consistently superior performance on our proposed soft entity-level metric F_{1,s}, capturing thus better the mention-based predictions as weighted by their cluster sizes. Finally, from Table 10(b) we notice that adding BERT to our joint model does not affect the boost in performance caused by the **RelProp** method for relation extraction. We hypothesize that this is due to the fact that **RelProp** propagation can capture relational semantics that goes beyond BERT’s contextual span representation similarity (which mainly drives the positive impact of *Joint+BERT*).

Unlike for the NER and RE tasks, where we observe a consistent positive impact of span graph propagation and joint modeling across all our experiments, the impact on the *coreference* task is not clear. Our experiments on *Single* setup show small, but constant improvement of the Avg.-F₁ score with the number of **AttProp** propagation iterations (see Fig. 5). However, in our *Joint* and *Joint+BERT* setups the graph propagation appears to not have any positive impact on Avg.-F₁ coreference scores. We hypothesize that the main reason for this phenomenon lies in the coreference annotations in DWIE: since we only annotate clusters of proper nouns, leaving out the nominal (e.g., “the prime minister”) and anaphoric expressions (e.g., “he”, “she”, “they”, etc), there might be little to no additional benefit in propagating information between co-referenced entity mentions, since the representation of proper nouns likely is not much influenced by textual context (e.g., the span “Merkel” can have very similar span representation to “Angela Merkel”, gaining nothing in adding contextual graph propagation).

Additionally, we explore in more detail the effect of the number of **AttProp**, **CorefProp**, **RelProp** graph propagation iterations on the final F₁ score of all the tasks in Fig. 5. We observe that the number of iterations have a decreasing effect on the improvement of performance for the NER and RE tasks. Furthermore, the positive effect of **CorefProp** and **RelProp** tends to saturate or even become negative after 1 or 2 iterations. This is in line with findings of Luan et al. (2019) on other datasets, where the performance peak is usually achieved at 2 graph propagation iterations. For our **AttProp** however, we observe that the positive effect of additional iterations tends to persist longer, particularly on the *Joint* setup where the positive effect of **AttProp** seems to be still growing after the last iteration (3) in our experiments.

7. Conclusions and Future Work

In this work we introduced DWIE, a manually annotated multi-task dataset that comprises Named Entity Recognition, Coreference, Relation Extraction and Entity Linking as main tasks. We highlight how DWIE is different from the mainstream datasets by focusing on document-level and entity-centric annotations. This also makes the predictions on this dataset more challenging by having not only to consider explicit, but also implicit document-level interactions between entities. Furthermore, we showed how Graph Neural Networks can help to tackle this issue by propagating local contextual mention span information on a

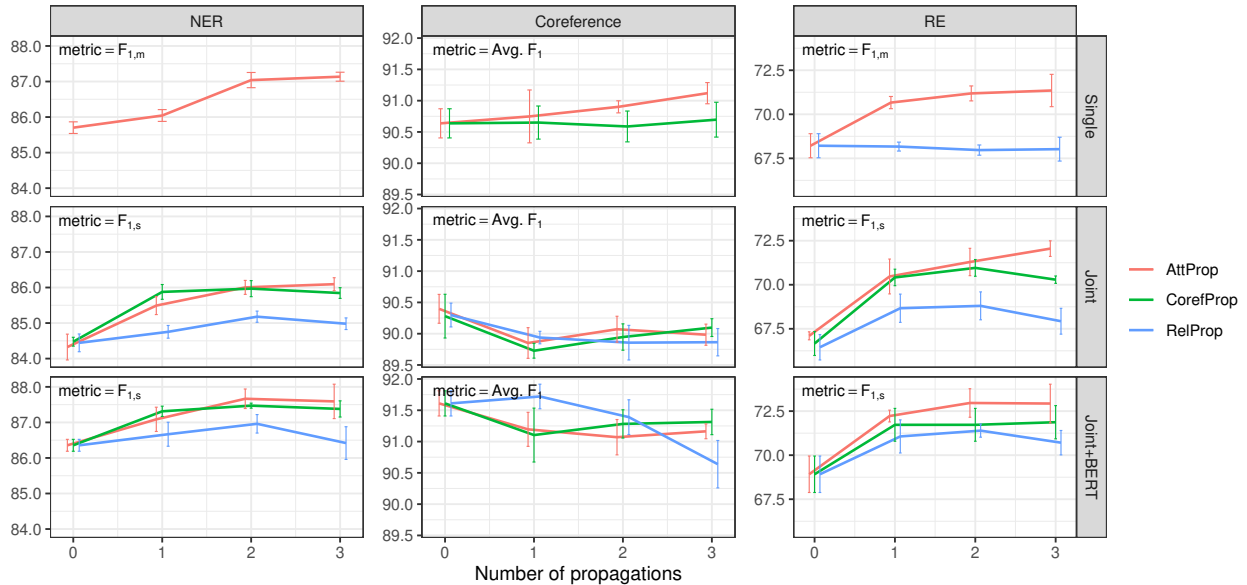


Figure 5: Impact of **AttProp**, **CorefProp** and **RelProp** graph propagations on performance metrics for each of the *Single*, *Joint* and *Joint+BERT* model setups. Note the different Y-axis scales.

document level for a single task as well as across the tasks on the DWIE dataset. We experiment with known graph propagation techniques driven by the scores of the coreference resolution (**CorefProp**) and relation extraction (**RelProp**) components, as well as introduced a new latent task-independent attention-based graph propagation method (**AttProp**). We demonstrated that, without relying on the task-specific scorers, **AttProp** can boost the performance of single-task as well as joint models, performing on par and even outperforming significantly in some scenarios the **RelProp** and **CorefProp** graph propagations.

In future work we will aim to integrate an entity linking component into our joint architecture. As a consequence, we expect to obtain a further boost in performance of different tasks included in DWIE by taking advantage of the information coming from Wikipedia 2018, the reference knowledge base for the entity linking annotations. Conversely, we conjecture that the results of the entity linking component can be improved when training it jointly with other tasks, such as NER and coreference resolution. Finally, we plan extending the coreference annotations to include nominal and anaphoric expressions. We expect that including these diverse mention types, whose initial span embedding representation can be different from coreferenced named entities, will make our coreference resolution task more challenging, allowing to investigate further the potential benefits of using graph-based neural networks.

Acknowledgements

Part of the research leading to these results has received funding from (i) the European Union’s Horizon 2020 research and innovation programme under grant agreement no. 761488 for the CPN project,¹⁷ and (ii) the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme.

¹⁷<https://www.projectcpn.eu/>

Appendix A. Dataset Insights

Table A.1: Statistics depicting the hierarchical structure of entity types described in Section 3.2. Only the most frequent entity types/subtypes are shown (% Mentions > 0.5%)

Entity Type	# Entities	% Entities	# Mentions	% Mentions
<i>ENTITY</i>	<i>13,151</i>	<i>56.9%</i>	<i>30,719</i>	<i>70.8%</i>
location	4,957	21.4%	11,548	26.6%
gpe	3,965	17.1%	9,830	22.7%
gpe0	2,225	9.6%	6,559	15.1%
gpe2	1,497	6.5%	2,873	6.6%
gpe1	244	1.1%	406	0.9%
regio	479	2.1%	916	2.1%
facility	259	1.1%	385	0.9%
organization	3,434	14.8%	8,165	18.8%
media	659	2.8%	984	2.3%
igo	547	2.4%	1,992	4.6%
so	171	0.7%	912	2.1%
party	381	1.6%	949	2.2%
company	368	1.6%	932	2.1%
sport_team	367	1.6%	1,106	2.5%
governmental_organization	342	1.5%	636	1.5%
agency	228	1.0%	444	1.0%
armed_movement	108	0.5%	374	0.9%
person	3,390	14.7%	8,259	19.0%
politician	1,184	5.1%	3,326	7.7%
head_of_state	380	1.6%	1,271	2.9%
head_of_gov	247	1.1%	673	1.6%
minister	217	0.9%	458	1.1%
sport_player	405	1.8%	844	1.9%
artist	260	1.1%	586	1.4%
politics_per	209	0.9%	457	1.1%
manager	104	0.4%	297	0.7%
offender	75	0.3%	347	0.8%
misc	823	3.6%	1,646	3.8%
work_of_art	174	0.8%	247	0.6%
event	354	1.5%	701	1.6%
sport_competition	183	0.8%	410	0.9%
ethnicity	84	0.4%	242	0.6%
<i>VALUE</i>	<i>5,903</i>	<i>25.5%</i>	<i>7,104</i>	<i>16.4%</i>
time	2,907	12.6%	3,608	8.3%
role	2,390	10.3%	2,865	6.6%
money	606	2.6%	631	1.5%
<i>OTHER</i>	<i>2,724</i>	<i>11.8%</i>	<i>5,482</i>	<i>12.6%</i>
gpe0-x	1,596	6.9%	3,827	8.8%
footer	413	1.8%	413	1.0%
loc-x	353	1.5%	585	1.3%
religion-x	235	1.0%	486	1.1%
TOTAL	23,130	100.0%	43,373	100.0%

Table A.2: Illustration of NER entity types in DWIE. Each cells contains possible entity subtypes (of different hierarchy levels) corresponding to the respective parent entity type (column) and topic (row).

	person	organization	event	location	misc
politics	head_of_gov, head_of_state, minister, politician_regional, politician_local, politician_national, candidate, politician, politics_per, activist, gov_per	politics_institution, politics_org, party, ngo, igo, so, policy_institute, movement, agency, ministry, military_alliance	summit_meeting, scandal, politics_event	politics_facility	politics_misc, project, treaty, report
culture	character, culture_per, artist, writer, actor, filmmaker, musician, photographer	music_band, culture_org, theatre_org, dance_org	festival, film_festival	culture_facility	art_title, culture_title, exhibition_title, culture_misc, work_of_art, book_title, film_title, tv_title, music_title, theatre_title, musical_title, film_award, book_award, music_award, tv_award, column_title, game, comic, radio_title, dance_title, opera
education	teacher, education_per, education_student	education_org		education_facility	education_study
religion	deity, clergy	religion_org	religious_event	religion_facility	religion, religion_misc
human	royalty				film_award, book_award, award, music_award, tv_award, sport_award
conflict	military_personnel, military_rebel	army, military_alliance, armed_movement	war, protest	military_facility	military_equipment, military_mission
media	journalist	media			
science	researcher, science_per	research_center			species, research_journal, technology
sport	sport_player, sport_coach, sport_head, sport_referee, sport_person	sport_team, sport_org	sport_competition	sport_facility	sport_award
labor	union_head, union_member, union_rep, union_per	union			
business	manager, employee, business_per	company, business_org, brand, trade_fair, market_exchange, advocacy		business_facility	product, market_index, business_misc
health	health_per	health_org		health_facility	health_disease, health_drug
justice	offender, advisor, victim, judge, police_per, justice_per	court, criminal_org, police_org, justice_org		prison	justice_misc, case
weather			storm		

Table A.3 describes the statistics of linked entities with respect to the total number of entities in each of the *Entity* subtypes. The columns *% Linked Entities* and *% Linked Mentions* indicate the percentage of annotated linked entities and mentions with respect to the total number of annotated entities/mentions in a particular *Entity* type category. Furthermore, we calculate two accuracies on test split when linking the entity mention with the most frequent entity link used either in DWIE: (i) training set of DWIE dataset (“Acc. Prior Train”), or (ii) Wikipedia corpus (“Acc. Prior Wiki”). Overall, using prior linking annotations from Wikipedia gives 9 percentage points better performance (79.0%) than when using train set (70.0%). This difference is explained by the fact that Wikipedia has much larger corpus to calculate the prior linking information from. Nevertheless, we still observe that for some entity types such as *sport_team* and *media* the accuracy based on DWIE training set prior is higher. This suggests the use of domain-specific language to refer to some entities in DWIE not used in a more general Wikipedia domain.

Table A.3: Entity linking statistics, only the top 5 types and subtypes with largest number of linked entities are showed. The *total* is calculated on all the entity types. The accuracy (both for most likely prior links on train and Wiki corpora) is computed on test set.

Entity Type	# Linked Entities	% Linked Entities	# Linked Mentions	% Linked Mentions	Acc. Prior Train	Acc. Prior Wiki
<i>LOCATION</i>	4,863	98.1%	11,496	99.5%	85.7%	92.9%
gpe	3,938	99.3%	9,810	99.8%	89.8%	95.6%
regio	456	95.2%	889	97.1%	83.3%	76.3%
facility	229	88.4%	381	99.0%	19.7%	73.8%
waterbody	90	98.9%	145	100.0%	83.3%	91.7%
district	37	94.9%	45	100.0%	33.3%	33.3%
<i>ORGANIZATION</i>	3,145	91.6%	8,029	98.3%	69.8%	70.8%
media	622	94.4%	979	99.5%	81.8%	59.5%
igo	525	96.0%	1,952	98.0%	76.4%	78.8%
party	358	94.0%	897	94.5%	77.5%	66.7%
company	320	87.0%	923	99.0%	67.6%	89.7%
sport_team	366	99.7%	1,105	99.9%	71.0%	47.5%
<i>PERSON</i>	2,627	77.5%	8,217	99.5%	45.7%	69.4%
politician	1,162	98.1%	3,324	99.9%	66.0%	78.1%
sport_player	404	99.8%	843	99.9%	34.4%	71.3%
artist	246	94.6%	567	96.8%	0.0%	29.4%
politics_per	126	60.3%	456	99.8%	23.7%	42.1%
manager	58	55.8%	296	99.7%	22.2%	33.3%
<i>MISC</i>	607	73.8%	1,532	93.1%	58.4%	73.4%
work_of_art	142	81.6%	246	99.6%	0.0%	100.0%
award	72	80.0%	186	94.9%	63.6%	81.8%
treaty	60	74.1%	149	99.3%	66.7%	50.0%
product	50	76.9%	146	98.6%	52.0%	92.0%
species	10	25.0%	14	18.4%	0.0%	100.0%
<i>EVENT</i>	320	90.4%	683	97.4%	49.4%	67.1%
sport_competition	163	89.1%	397	96.8%	64.6%	87.5%
summit_meeting	15	68.2%	37	92.5%	100.0%	100.0%
holiday	21	95.5%	39	97.5%	100.0%	100.0%
history	17	89.5%	30	100.0%	100.0%	100.0%
protest	14	100.0%	22	100.0%	80.0%	100.0%
TOTAL	13,086	56.6%	28,482	65.7%	70.0%	79.0%

Table A.4: Main named entity tag categories with statistics of the number and % of covered entities and mentions as well as the number of classes in each and average number of labels per entity cluster.

Entity Tag Category	# Entities	% Entities	# Mentions	% Mentions	# Classes	Labels per Entity
type	21,745	94.0%	43,122	99.4%	174	2.9
topic	7,843	33.9%	18,359	42.3%	14	1.0
iptc	7,059	30.5%	17,195	39.6%	114	1.3
gender	3,352	14.5%	8,200	18.9%	2	1.0
slot	3,232	14.0%	14,983	34.5%	7	1.2
TOTAL	23,130	100.0%	43,373	100.0%	311	4.0

Table A.4 illustrates the number of annotated entities and mentions per each tag category (type, topic, iptc, gender and slot). It also showcases the multi-label nature of entity classification task in DWIE, with the average number of labels per entity of 4.0.

Table A.5 illustrates the number and percentage of related entities and mentions of our dataset grouped by the number of relation labels. It also compares with other entity-centric RE datasets, namely BC5CDR (Li et al., 2016a; Wei et al., 2015) and DocRED (Yao et al., 2019) datasets.

Table A.5: This table groups the number of related pairs in DWIE by the number of assigned relation labels to each of these pairs. We compare with other two entity-centric datasets: BC5CDR and DocRED.

# Relation labels	DWIE				BC5CDR	DocRED
	# Related ent. pairs	% Related ent. pairs	# Related mention pairs	% Related mention pairs	% Related ent. pairs	% Related ent. pairs
1	12,856	76.32%	112,708	69.40%	100%	92.89%
2	3,101	18.41%	34,948	21.52%	0%	6.82%
3	884	5.25%	14,650	9.02%	0%	0.26%
4	3	0.02%	100	0.06%	0%	0.03%
TOTAL	16,844	100.0%	162,406	100.0%	100.0%	100.0%

Table A.6: Relation type statistics. We compare the number of related entity and mention pairs per relation type. Only the most frequent relation types are shown (% Related Mention Pairs > 0.1%)

Relation Type	# Related Ent. Pairs	% Related Ent. Pairs	# Related Men. Pairs	% Related Men. Pairs
based_in0	2,361	14.0%	18,771	11.6%
in0	2,120	12.6%	15,810	9.7%
citizen_of	1,969	11.7%	25,752	15.9%
based_in0-x	1,882	11.2%	12,211	7.5%
citizen_of-x	1,844	10.9%	17,049	10.5%
member_of	1,616	9.6%	19,953	12.3%
gpe0	1,569	9.3%	18,110	11.2%
in0-x	1,474	8.8%	8,784	5.4%
agent_of	954	5.7%	15,776	9.7%
head_of	564	3.3%	7,710	4.7%
agency_of	435	2.6%	4,775	2.9%
player_of	401	2.4%	5,692	3.5%
agency_of-x	382	2.3%	2,108	1.3%
head_of.state	380	2.3%	7,986	4.9%
head_of.state-x	343	2.0%	3,853	2.4%
appears_in	294	1.7%	4,555	2.8%
vs	281	1.7%	7,187	4.4%
head_of.gov	273	1.6%	4,015	2.5%
head_of.gov-x	247	1.5%	2,383	1.5%
minister_of	234	1.4%	2,280	1.4%
minister_of-x	213	1.3%	1,629	1.0%
based_in2	185	1.1%	971	0.6%
event_in0	181	1.1%	843	0.5%
part_of	164	1.0%	2,858	1.8%
in2	157	0.9%	1,055	0.6%
created_by	134	0.8%	945	0.6%
agent_of-x	125	0.7%	897	0.6%
award_received	111	0.7%	969	0.6%
institution_of	105	0.6%	2,113	1.3%
ministry_of	81	0.5%	666	0.4%
coach_of	65	0.4%	1,211	0.7%
won_vs	61	0.4%	1,531	0.9%
spouse_of	55	0.3%	599	0.4%
directed_by	44	0.3%	318	0.2%
is_meeting	41	0.2%	968	0.6%
event_in2	40	0.2%	259	0.2%
spokesperson_of	39	0.2%	177	0.1%
plays_in	38	0.2%	330	0.2%
gpe1	35	0.2%	135	0.1%
product_of	31	0.2%	334	0.2%
parent_of	22	0.1%	281	0.2%
child_of	22	0.1%	281	0.2%
based_in1	22	0.1%	376	0.2%
signed_by	20	0.1%	521	0.3%
law_of	16	0.1%	286	0.2%
TOTAL	16,844	100.0%	162,406	100.0%

Appendix B. Inter-annotator agreement calculations

In order to measure the agreement we use Cohen’s kappa coefficient (Cohen, 1960), defined as

$$\kappa = \frac{p_o - p_e}{1 - p_e} \tag{B.1}$$

where p_o represents the observed agreement between the two annotators and p_e is the expected agreement between the annotators (i.e., agreement by chance). More specifically, in our case we calculate the observed probability p_o as in Eq. (B.2) where N is the number of annotated items, $A_{i,j}$ is the annotation made by annotator i for item j , and $\mathbb{1}\{A_{1,j} = A_{2,j}\}$ returns 1 if $A_{1,j}$ is equal to $A_{2,j}$ and 0 otherwise. Thus, p_o can be interpreted as the fraction of the labels two annotators agree, also called *percent agreement* (McHugh, 2012; Scott, 1955).

$$p_o = \frac{\sum_{j=1}^N \mathbb{1}\{A_{1,j} = A_{2,j}\}}{N} \tag{B.2}$$

To calculate the expected agreement probability we use the formulation in Eq. (B.3). It can be interpreted as the probability that both annotators, when randomly distributing all of their label annotations among the items to be annotated, assign the same label to a given item.

$$p_e = \sum_{l=1}^L \frac{n_{1,l}}{N} \frac{n_{2,l}}{N} \tag{B.3}$$

In this context, $n_{i,l}$ is the number of items the annotator i annotated with label l and L is the total number of labels. For multi-label annotations where it is possible to assign multiple classes for a particular annotation item (i.e., named entity and relation types), we report a weighted kappa score.

Appendix C. Relation consistency rules

This appendix enumerates the logical predicates used as a consistency check in our dataset.

- $$\begin{aligned} \text{spouse_of}\langle Y, X \rangle &\implies \text{spouse_of}\langle X, Y \rangle && \text{(C.1)} \\ \text{vs}\langle Y, X \rangle &\implies \text{vs}\langle X, Y \rangle && \text{(C.2)} \\ \text{won_vs}\langle X, Y \rangle &\implies \text{vs}\langle X, Y \rangle && \text{(C.3)} \\ \text{won_vs}\langle X, Y \rangle &\implies \text{vs}\langle Y, X \rangle && \text{(C.4)} \\ \text{child_of}\langle Y, X \rangle &\implies \text{parent_of}\langle X, Y \rangle && \text{(C.5)} \\ \text{parent_of}\langle Y, X \rangle &\implies \text{child_of}\langle X, Y \rangle && \text{(C.6)} \\ \text{ministry_of}\langle X, Y \rangle &\implies \text{agency_of}\langle X, Y \rangle && \text{(C.7)} \\ \text{agency_of-x}\langle X, Z \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{agency_of}\langle X, Y \rangle && \text{(C.8)} \\ \text{agency_of}\langle X, Y \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{agency_of-x}\langle X, Z \rangle && \text{(C.9)} \\ \text{agent_of-x}\langle X, Z \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{agent_of}\langle X, Y \rangle && \text{(C.10)} \\ \text{agent_of}\langle X, Y \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{agent_of-x}\langle X, Z \rangle && \text{(C.11)} \\ \text{minister_of}\langle X, Y \rangle &\implies \text{agent_of}\langle X, Y \rangle && \text{(C.12)} \\ \text{head_of_gov}\langle X, Y \rangle &\implies \text{agent_of}\langle X, Y \rangle && \text{(C.13)} \\ \text{head_of_state}\langle X, Y \rangle &\implies \text{agent_of}\langle X, Y \rangle && \text{(C.14)} \\ \text{citizen_of-x}\langle X, Z \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{citizen_of}\langle X, Y \rangle && \text{(C.15)} \\ \text{citizen_of}\langle X, Y \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{citizen_of-x}\langle X, Z \rangle && \text{(C.16)} \\ \text{minister_of-x}\langle X, Z \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{minister_of}\langle X, Y \rangle && \text{(C.17)} \\ \text{minister_of}\langle X, Y \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{minister_of-x}\langle X, Z \rangle && \text{(C.18)} \\ \text{head_of_state-x}\langle X, Z \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{head_of_state}\langle X, Y \rangle && \text{(C.19)} \\ \text{head_of_state}\langle X, Y \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{head_of_state-x}\langle X, Z \rangle && \text{(C.20)} \\ \text{head_of_gov-x}\langle X, Z \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{head_of_gov}\langle X, Y \rangle && \text{(C.21)} \\ \text{head_of_gov}\langle X, Y \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{head_of_gov-x}\langle X, Z \rangle && \text{(C.22)} \\ \text{in0-x}\langle X, Z \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{in0}\langle X, Y \rangle && \text{(C.23)} \\ \text{in0}\langle X, Y \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{in0-x}\langle X, Z \rangle && \text{(C.24)} \\ \text{in2}\langle X, Z \rangle \wedge \text{in0}\langle Z, Y \rangle &\implies \text{in0}\langle X, Y \rangle && \text{(C.25)} \\ \text{in1}\langle X, Z \rangle \wedge \text{in0}\langle Z, Y \rangle &\implies \text{in0}\langle X, Y \rangle && \text{(C.26)} \\ \text{based_in2}\langle X, Z \rangle \wedge \text{in0}\langle Z, Y \rangle &\implies \text{based_in0}\langle X, Y \rangle && \text{(C.27)} \\ \text{based_in1}\langle X, Z \rangle \wedge \text{in0}\langle Z, Y \rangle &\implies \text{based_in0}\langle X, Y \rangle && \text{(C.28)} \\ \text{agency_of}\langle X, Y \rangle \wedge \text{gpe0}\langle Y \rangle &\implies \text{based_in0}\langle X, Y \rangle && \text{(C.29)} \\ \text{event_in2}\langle X, Z \rangle \wedge \text{in0}\langle Z, Y \rangle &\implies \text{event_in0}\langle X, Y \rangle && \text{(C.30)} \\ \text{event_in1}\langle X, Z \rangle \wedge \text{in0}\langle Z, Y \rangle &\implies \text{event_in0}\langle X, Y \rangle && \text{(C.31)} \\ \text{head_of}\langle X, Y \rangle &\implies \text{member_of}\langle X, Y \rangle && \text{(C.32)} \\ \text{coach_of}\langle X, Y \rangle &\implies \text{member_of}\langle X, Y \rangle && \text{(C.33)} \\ \text{spokesperson_of}\langle X, Y \rangle &\implies \text{member_of}\langle X, Y \rangle && \text{(C.34)} \\ \text{member_of}\langle X, Y \rangle \wedge \text{sport_player}\langle X \rangle &\implies \text{player_of}\langle X, Y \rangle && \text{(C.35)} \\ \text{mayor_of}\langle X, Y \rangle &\implies \text{head_of_gov}\langle X, Y \rangle && \text{(C.36)} \\ \text{directed_by}\langle X, Y \rangle &\implies \text{created_by}\langle X, Y \rangle && \text{(C.37)} \\ \text{character_in}\langle X, Y \rangle \wedge \text{played_by}\langle X, Z \rangle &\implies \text{plays_in}\langle Z, Y \rangle && \text{(C.38)} \\ \text{institution_of}\langle X, Y \rangle &\implies \text{part_of}\langle X, Y \rangle && \text{(C.39)} \\ \text{based_in0-x}\langle X, Z \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{based_in0}\langle X, Y \rangle && \text{(C.40)} \\ \text{based_in0}\langle X, Y \rangle \wedge \text{gpe0}\langle Z, Y \rangle &\implies \text{based_in0-x}\langle X, Z \rangle && \text{(C.41)} \end{aligned}$$

References

- Aguilar, J., Beller, C., McNamee, P., Van Durme, B., Strassel, S., Song, Z., & Ellis, J. (2014). A comparison of the events and relations across ace, ere, tac-kbp, and framenet annotation standards. In *Proceedings of the 2nd Workshop on EVENTS: Definition, Detection, Coreference, and Representation* (pp. 45–53).
- Akbik, A., Bergmann, T., & Vollgraf, R. (2019). Pooled contextualized embeddings for named entity recognition. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 724–728).
- Akbik, A., Blythe, D., & Vollgraf, R. (2018). Contextual string embeddings for sequence labeling. In *Proceedings of the 2018 International Conference on Computational Linguistics* (pp. 1638–1649).
- Augenstein, I., Das, M., Riedel, S., Vikraman, L., & 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)* (pp. 546–555).
- Baevski, A., Edunov, S., Liu, Y., Zettlemoyer, L., & Auli, M. (2019). Cloze-driven pretraining of self-attention networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing* (pp. 5363–5372).
- Bagga, A., & Baldwin, B. (1998). Algorithms for scoring coreference chains. In *Proceedings of the 1998 International Conference on Language Resources and Evaluation Workshop on Linguistics Coreference* (pp. 563–566).
- Bekoulis, G., Deleu, J., Demeester, T., & Develder, C. (2017). Reconstructing the house from the ad: Structured prediction on real estate classifieds. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers* (pp. 274–279).
- Bekoulis, G., Deleu, J., Demeester, T., & Develder, C. (2018a). Adversarial training for multi-context joint entity and relation extraction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (pp. 2830–2836).
- Bekoulis, G., Deleu, J., Demeester, T., & Develder, C. (2018b). Joint entity recognition and relation extraction as a multi-head selection problem. *Expert Systems with Applications*, 114, 34–45.
- Bentivogli, L., Forner, P., Giuliano, C., Marchetti, A., Pianta, E., & Tymoshenko, K. (2010). Extending english ace 2005 corpus annotation with ground-truth links to wikipedia. In *Proceedings of the 2nd Workshop on The People’s Web Meets NLP: Collaboratively Constructed Semantic Resources* (pp. 19–27).
- Bhattacharjee, S., Haque, R., de Buy Wenniger, G. M., & Way, A. (2020). Investigating query expansion and coreference resolution in question answering on bert. In *International Conference on Applications of Natural Language to Information Systems* (pp. 47–59). Springer.
- Broscheit, S. (2019). Investigating entity knowledge in bert with simple neural end-to-end entity linking. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)* (pp. 677–685).
- Chen, D., Fisch, A., Weston, J., & Bordes, A. (2017). Reading wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 1870–1879).
- Chinchor, N., & Marsh, E. (1998). Muc-7 information extraction task definition. In *Proceeding of the 1998 Message Understanding Conference (MUC-7)* (pp. 359–367).
- Chiu, J. P., & Nichols, E. (2016). Named entity recognition with bidirectional lstm-cnns. *Transactions of the Association for Computational Linguistics*, 4, 357–370.
- Cifariello, P., Ferragina, P., & Ponza, M. (2019). Wiser: A semantic approach for expert finding in academia based on entity linking. *Information Systems*, 82, 1–16.
- Clark, K., Luong, M.-T., Manning, C. D., & Le, Q. (2018). Semi-supervised sequence modeling with cross-view training. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (pp. 1914–1925).
- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20, 37–46.
- Derczynski, L., Nichols, E., van Erp, M., & Limsopatham, N. (2017). Results of the wnut2017 shared task on novel and emerging entity recognition. In *Proceedings of the 3rd Workshop on Noisy User-generated Text* (pp. 140–147).
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 4171–4186).
- Dixit, K., & Al-Onaizan, Y. (2019). Span-level model for relation extraction. In *Proceedings of the 2019 Annual Meeting of the Association for Computational Linguistics* (pp. 5308–5314).
- Doddington, G. R., Mitchell, A., Przybocki, M. A., Ramshaw, L. A., Strassel, S. M., & Weischedel, R. M. (2004). The automatic content extraction (ace) program - tasks, data, and evaluation. In *Proceedings of the 2004 International Conference on Language Resources and Evaluation Workshop on Linguistics* (pp. 837–840).
- Durrett, G., & Klein, D. (2013). Easy victories and uphill battles in coreference resolution. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing* (pp. 1971–1982).
- Ellis, J., Getman, J., Fore, D., Kuster, N., Song, Z., Bies, A., & Strassel, S. M. (2015). Overview of linguistic resources for the TAC KBP 2015 evaluations: Methodologies and results. In *Proceedings of the 2015 Text Analysis Conference*.
- Ellis, J., Getman, J., & Strassel, S. M. (2014). Overview of linguistic resources for the tac kbp 2014 evaluations: Planning, execution, and results. In *Proceedings of TAC KBP 2014 Workshop, National Institute of Standards and Technology* (pp. 17–18).
- Eshel, Y., Cohen, N., Radinsky, K., Markovitch, S., Yamada, I., & Levy, O. (2017). Named entity disambiguation for noisy text. In *Proceedings of the 2017 Conference on Computational Natural Language Learning* (pp. 58–68).
- Fei, H., Ren, Y., & Ji, D. (2020). Boundaries and edges rethinking: An end-to-end neural model for overlapping entity relation extraction. *Information Processing & Management*, 57, 102311.

- Fu, T.-J., Li, P.-H., & Ma, W.-Y. (2019). Graphrel: Modeling text as relational graphs for joint entity and relation extraction. In *Proceedings of the 2019 Annual Meeting of the Association for Computational Linguistics* (pp. 1409–1418).
- Gao, Y., Li, P., King, I., & Lyu, M. R. (2019). Interconnected question generation with coreference alignment and conversation flow modeling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 4853–4862).
- Guo, Z., Zhang, Y., & Lu, W. (2019). Attention guided graph convolutional networks for relation extraction. In *Proceedings of the 2019 Annual Meeting of the Association for Computational Linguistics* (pp. 241–251).
- Han, X., & Wang, L. (2020). A novel document-level relation extraction method based on bert and entity information. *IEEE Access*, .
- Han, X., Zhu, H., Yu, P., Wang, Z., Yao, Y., Liu, Z., & Sun, M. (2018). Fewrel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (pp. 4803–4809).
- Hendrickx, I., Kim, S. N., Kozareva, Z., Nakov, P., Séaghdha, D. Ó., Padó, S., Pennacchiotti, M., Romano, L., & Szpakowicz, S. (2010). Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. In *Proceedings of the 5th International Workshop on Semantic Evaluation* (pp. 33–38).
- Hoffart, J., Yosef, M. A., Bordino, I., Fürstenau, H., Pinkal, M., Spaniol, M., Taneva, B., Thater, S., & Weikum, G. (2011). Robust disambiguation of named entities in text. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing* (pp. 782–792).
- Hovy, E., Marcus, M., Palmer, M., Ramshaw, L., & Weischedel, R. (2006). Ontonotes: the 90% solution. In *Proceedings of the 2006 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 57–60).
- Hu, R., Rohrbach, A., Darrell, T., & Saenko, K. (2019). Language-conditioned graph networks for relational reasoning. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 10294–10303).
- Hu, W., Ma, B., Li, Z., Li, Y., & Wang, Y. (2020). A cross-media deep relationship classification method using discrimination information. *Information Processing & Management*, 57, 102344.
- Ji, H., Grishman, R., Dang, H. T., Griffitt, K., & Ellis, J. (2010). Overview of the tac 2010 knowledge base population track. In *Proceedings of the 2010 Text Analysis Conference* (pp. 3–3).
- Ji, H., Nothman, J., Hachey, B., & Florian, R. (2015). Overview of tac-kbp2015 tri-lingual entity discovery and linking. In *Proceedings of the 2015 Text Analysis Conference*.
- Ji, H., Pan, X., Zhang, B., Nothman, J., Mayfield, J., McNamee, P., Costello, C., & Hub, S. I. (2017). Overview of tac-kbp2017 13 languages entity discovery and linking. In *Proceedings of the 2017 Text Analysis Conference*.
- Kantor, B., & Globerson, A. (2019). Coreference resolution with entity equalization. In *Proceedings of the 2019 Annual Meeting of the Association for Computational Linguistics* (pp. 673–677).
- Karimi, M., Jannach, D., & Jugovac, M. (2018). News recommender systems—survey and roads ahead. *Information Processing & Management*, 54, 1203–1227.
- Katihar, A., & Cardie, C. (2018). Nested named entity recognition revisited. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 861–871).
- Kim, J.-D., Ohta, T., Tateisi, Y., & Tsujii, J. (2003). Genia corpus - a semantically annotated corpus for bio-textmining. *Bioinformatics*, 19, 180–182.
- Kulkarni, C., Xu, W., Ritter, A., & Machiraju, R. (2018). An annotated corpus for machine reading of instructions in wet lab protocols. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 97–106).
- Kulkarni, S., Singh, A., Ramakrishnan, G., & Chakrabarti, S. (2009). Collective annotation of wikipedia entities in web text. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 457–466).
- Lample, G., Ballesteros, M., Subramanian, S., Kawakami, K., & Dyer, C. (2016). Neural architectures for named entity recognition. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 260–270).
- Landis, J., & Koch, G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, 33, 159–174.
- Lee, K., He, L., Lewis, M., & Zettlemoyer, L. (2017). End-to-end neural coreference resolution. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 188–197).
- Lee, K., He, L., & Zettlemoyer, L. (2018). Higher-order coreference resolution with coarse-to-fine inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 687–692).
- Li, J., Sun, Y., Johnson, R. J., Sciaky, D., Wei, C.-H., Leaman, R., Davis, A. P., Mattingly, C. J., Wieggers, T. C., & Lu, Z. (2016a). Biocreative v cdr task corpus: a resource for chemical disease relation extraction. *Database*, 2016.
- Li, Q., & Ji, H. (2014). Incremental joint extraction of entity mentions and relations. In *Proceedings of the 2014 Annual Meeting of the Association for Computational Linguistics* (pp. 402–412).
- Li, Y., Tarlow, D., Brockschmidt, M., & Zemel, R. (2016b). Gated graph sequence neural networks. In *Proceedings of the 2016 International Conference on Learning Representations*.
- Luan, Y., He, L., Ostendorf, M., & Hajishirzi, 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* (pp. 3219–3232).
- Luan, Y., Ostendorf, M., & Hajishirzi, H. (2017). Scientific information extraction with semi-supervised neural tagging. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 2641–2651).
- Luan, Y., Wadden, D., He, L., Shah, A., Ostendorf, M., & Hajishirzi, H. (2019). A general framework for information extraction using dynamic span graphs. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for*

- Computational Linguistics: Human Language Technologies* (pp. 3036–3046).
- Luo, X. (2005). On coreference resolution performance metrics. In *Proceedings of the 2005 Conference on Human Language Technology and Empirical Methods in Natural Language Processing* (pp. 25–32).
- Ma, X., & Hovy, E. (2016). End-to-end sequence labeling via bi-directional lstm-cnns-crf. In *Proceedings of the 2016 Annual Meeting of the Association for Computational Linguistics* (pp. 1064–1074).
- McHugh, M. L. (2012). Interrater reliability: the kappa statistic. *Biochemia medica: Biochemia medica*, 22, 276–282.
- Molla, D., van Zaanen, M., & Smith, D. (2006). Named entity recognition for question answering. In *Proceedings of the Australasian Language Technology Workshop 2006* (pp. 51–58).
- Peng, N., Poon, H., Quirk, C., Toutanova, K., & Yih, W.-t. (2017). Cross-sentence n-ary relation extraction with graph lstms. *Transactions of the Association for Computational Linguistics*, 5, 101–115.
- Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing* (pp. 1532–1543).
- Peters, M. E., Neumann, M., Logan, R., Schwartz, R., Joshi, V., Singh, S., & Smith, N. A. (2019). Knowledge enhanced contextual word representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Machine Learning and Knowledge Discovery in Databases* (pp. 43–54).
- Pradhan, S., Luo, X., Recasens, M., Hovy, E., Ng, V., & Strube, M. (2014). Scoring coreference partitions of predicted mentions: A reference implementation. In *Proceedings of the 2014 Annual Meeting of the Association for Computational Linguistics* (pp. 30–35).
- Pradhan, S., Moschitti, A., Xue, N., Uryupina, O., & Zhang, Y. (2012). Conll-2012 shared task: Modeling multilingual unrestricted coreference in ontonotes. In *Proceedings of the 2012 Conference on Computational Natural Language Learning* (pp. 1–40).
- Quirk, C., & Poon, H. (2017). Distant supervision for relation extraction beyond the sentence boundary. In *Proceedings of the 2017 Conference of the European Chapter of the Association for Computational Linguistics* (pp. 1171–1182).
- Riedel, S., Yao, L., & McCallum, A. (2010). Modeling relations and their mentions without labeled text. In *Proceedings of the 2010 European Conference on Machine Learning and Knowledge Discovery in Databases* (pp. 148–163).
- Roller, S., Dinan, E., Goyal, N., Ju, D., Williamson, M., Liu, Y., Xu, J., Ott, M., Shuster, K., Smith, E. M. et al. (2020). Recipes for building an open-domain chatbot. *arXiv preprint arXiv:2004.13637*, .
- Sang, E. F. T. K., & De Meulder, F. (2003). Introduction to the conll-2003 shared task: Language-independent named entity recognition. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 142–147).
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., & Monfardini, G. (2008). The graph neural network model. *IEEE Transactions on Neural Networks*, 20, 61–80.
- Scott, W. A. (1955). Reliability of content analysis: The case of nominal scale coding. *Public opinion quarterly*, (pp. 321–325).
- Singh, K., Radhakrishna, A. S., Both, A., Shekarpour, S., Lytra, I., Usbeck, R., Vyas, A., Khikmatullaev, A., Punjani, D., Lange, C. et al. (2018). Why reinvent the wheel: Let’s build question answering systems together. In *Proceedings of the 2018 World Wide Web Conference* (pp. 1247–1256).
- Soares, L. B., FitzGerald, N., Ling, J., & Kwiatkowski, T. (2019). Matching the blanks: Distributional similarity for relation learning. In *Proceedings of the 2019 Annual Meeting of the Association for Computational Linguistics* (pp. 2895–2905).
- Song, Z., Bies, A., Strassel, S., Riese, T., Mott, J., Ellis, J., Wright, J., Kulick, S., Ryant, N., & Ma, X. (2015). From light to rich ere: annotation of entities, relations, and events. In *Proceedings of the the 3rd Workshop on EVENTS: Definition, Detection, Coreference, and Representation* (pp. 89–98).
- Strubell, E., Verga, P., Belanger, D., & McCallum, A. (2017). Fast and accurate entity recognition with iterated dilated convolutions. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 2670–2680).
- Sun, S., Luo, C., & Chen, J. (2017). A review of natural language processing techniques for opinion mining systems. *Information fusion*, 36, 10–25.
- Thorne, J., & Vlachos, A. (2018). Automated fact checking: Task formulations, methods and future directions. In *Proceedings of the 27th International Conference on Computational Linguistics* (pp. 3346–3359).
- Wadden, D., Wennberg, U., Luan, Y., & Hajishirzi, H. (2019). Entity, relation, and event extraction with contextualized span representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing* (pp. 5788–5793).
- Walker, C., Strassel, S., Medero, J., & Maeda, K. (2006). Ace 2005 multilingual training corpus. *Linguistic Data Consortium, Philadelphia*, 57.
- Wang, H., Zhang, F., Xie, X., & Guo, M. (2018). Dkn: Deep knowledge-aware network for news recommendation. In *Proceedings of the 2018 world wide web conference* (pp. 1835–1844).
- Wang, H., Zhang, F., Zhao, M., Li, W., Xie, X., & Guo, M. (2019). Multi-task feature learning for knowledge graph enhanced recommendation. In *The World Wide Web Conference* (pp. 2000–2010).
- Webster, K., Recasens, M., Axelrod, V., & Baldridge, J. (2018). Mind the gap: A balanced corpus of gendered ambiguous pronouns. *Transactions of the Association for Computational Linguistics*, 6, 605–617.
- Wei, C.-H., Peng, Y., Leaman, R., Davis, A. P., Mattingly, C. J., Li, J., Wiegers, T. C., & Lu, Z. (2015). Overview of the biocreative v chemical disease relation (cdr) task. In *Proceedings of the 5th BioCreative Challenge Evaluation Workshop*.
- Weischedel, R., Hovy, E., Marcus, M., Palmer, M., Belvin, R., Pradhan, S., Ramshaw, L., & Xue, N. (2011). Ontonotes: A large training corpus for enhanced processing. *Handbook of Natural Language Processing and Machine Translation*. Springer, (p. 59).
- Weischedel, R., Palmer, M., Marcus, M., Hovy, E., Pradhan, S., Ramshaw, L., Xue, N., Taylor, A., Kaufman, J., Franchini,

- M. et al. (2013). Ontonotes release 5.0 ldc2013t19. *Linguistic Data Consortium, Philadelphia, PA*, 23.
- Wiseman, S., Rush, A. M., Shieber, S. M., & Weston, J. (2015). Learning anaphoricity and antecedent ranking features for coreference resolution. In *Proceedings of the 2015 Annual Meeting of the Association for Computational Linguistics and International Joint Conference on Natural Language Processing* (pp. 1416–1426).
- Wu, S., & He, Y. (2019). Enriching pre-trained language model with entity information for relation classification. In *Proceedings of the 2019 ACM International Conference on Information and Knowledge Management* (pp. 2361–2364).
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Philip, S. Y. (2020). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, (pp. 1–21).
- Xu, K., Hu, W., Leskovec, J., & Jegelka, S. (2018). How powerful are graph neural networks? In *Proceedings of the 2018 International Conference on Learning Representations*.
- Yao, Y., Ye, D., Li, P., Han, X., Lin, Y., Liu, Z., Liu, Z., Huang, L., Zhou, J., & Sun, M. (2019). Docred: A large-scale document-level relation extraction dataset. In *Proceedings of the 2019 Annual Meeting of the Association for Computational Linguistics* (pp. 764–777).
- Yu, M., Yin, W., Hasan, K. S., dos Santos, C., Xiang, B., & Zhou, B. (2017). Improved neural relation detection for knowledge base question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 571–581).
- Zhang, M., Zhang, Y., & Fu, G. (2017a). End-to-end neural relation extraction with global optimization. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 1730–1740).
- Zhang, X., & Ghorbani, A. A. (2020). An overview of online fake news: Characterization, detection, and discussion. *Information Processing & Management*, 57, 102025.
- Zhang, Y., Qi, P., & Manning, C. D. (2018). Graph convolution over pruned dependency trees improves relation extraction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing* (pp. 2205–2215).
- Zhang, Y., Zhong, V., Chen, D., Angeli, G., & Manning, C. D. (2017b). Position-aware attention and supervised data improve slot filling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 35–45).