

WEC: Deriving a Large-scale Cross-document Event Coreference dataset from Wikipedia

Alon Eirew^{1,2} Arie Cattan¹ Ido Dagan¹

¹Bar Ilan University, Ramat-Gan, Israel ²Intel Labs, Israel
alon.eirew@intel.com arie.cattan@gmail.com
dagan@cs.biu.ac.il

Abstract

Cross-document event coreference resolution is a foundational task for NLP applications involving multi-text processing. However, existing corpora for this task are scarce and relatively small, while annotating only modest-size clusters of documents belonging to the same topic. To complement these resources and enhance future research, we present **Wikipedia Event Coreference (WEC)**, an efficient methodology for gathering a large-scale dataset for cross-document event coreference from Wikipedia, where coreference links are not restricted within predefined topics. We apply this methodology to the English Wikipedia and extract our large-scale **WEC-Eng** dataset. Notably, our dataset creation method is generic and can be applied with relatively little effort to other Wikipedia languages. To set baseline results, we develop an algorithm that adapts components of state-of-the-art models for within-document coreference resolution to the cross-document setting. Our model is suitably efficient and outperforms previously published state-of-the-art results for the task.

1 Introduction

Cross-Document (CD) Event Coreference resolution is the task of identifying clusters of text mentions, across multiple texts, that refer to the same event. Successful identification of such corefering mentions is beneficial for a broad range of applications at the multi-text level, which are gaining increasing interest and need to match and integrate information across documents, such as multi-document summarization (Falke et al., 2017; Liao et al., 2018), multi-hop question answering (Dhingra et al., 2018; Wang et al., 2019) and Knowledge Base Population (KBP) (Lin et al., 2020).

Unfortunately, rather few datasets of reasonable scale exist for CD event coreference. Notable datasets include ECB+ (Cybulska and Vossen,

2014), MEANTIME (Minard et al., 2016) and the Gun Violence Corpus (GVC) (Vossen et al., 2018) (described in Section 2), where recent work has been evaluated solely on ECB+. When addressed in a direct manner, manual CD coreference annotation is very hard due to its worst-case quadratic complexity, where each mention may need to be compared to all other mentions in all documents. Indeed, ECB+ contains less than 7000 event mentions in total (train, dev, and test sets). Further, effective corpora for CD event coreference are available mostly for English, limiting research opportunities for other languages. Partly as a result of this data scarcity, rather little effort was invested in this field in recent years, compared to dramatic recent progress in modeling within-document coreference.

Furthermore, most existing cross-document coreference datasets are restricted in their scope by two inter-related characteristics. First, these datasets annotate sets of documents, where the documents in each set all describe the *same topic*, mostly a news event (consider the Malaysia Airlines crash as an example). While such topic-focused document sets guarantee a high density of corefering event mentions, facilitating annotation, in practical settings the same event might be mentioned across an entire corpus, being referred to in documents of varied topics. Second, we interestingly observed that event mentions may be (softly) classified into two different types. One type, which we term a *descriptive* mention, pertains to a mention involved in presenting the event or describing new information about it. For example, news about the Malaysian Airline crash will include mostly descriptive mentions of the event and its sub-events, such as *shot-down*, *crashed* and *investigated*. Naturally, news documents about a topic, as in prior event coreference datasets, include mostly descriptive event mentions. The other type, which we term a *referential* mention, pertains to

mentions of the event in sentences that do not focus on presenting new information about the event but rather mention it as a point of reference. For example, mentions referring to the airplane crash, such as *the Malaysian plane crash*, *Flight MH17* or *disaster* may appear in documents about the war in Donbass or about flight safety. Since referential event mentions are split across an entire corpus, they are less trivial to identify for coreference annotation, and are mostly missing in current news-based datasets. As we demonstrate later, these two mention types exhibit different lexical distributions and seem to require corresponding training data to be properly modeled.

In this paper, we present the **Wikipedia Event Coreference methodology (WEC)**, an efficient method for automatically gathering a large-scale dataset for the cross-document event coreference task. Our methodology effectively complements current datasets in the above-mentioned respects: data annotation is boosted by leveraging available information in Wikipedia, practically applicable for any Wikipedia language; mentions are gathered across the entire Wikipedia corpus, yielding a dataset that is not partitioned by topics; and finally, our dataset consists mostly of referential event mentions.

In its essence, our methodology leverages the coreference relation that often holds between anchor texts of hyperlinks pointing to the same Wikipedia article (see Figure 1), similar to the basic idea introduced in the Wikilinks dataset (Singh et al., 2012). Focusing on CD *event* coreference, we identify and target only Wikipedia articles denoting events. Anchor texts pointing to the same event article, along with some surrounding context, become candidate mentions for a corresponding event coreference cluster, undergoing extensive filtering. We apply our method to the English Wikipedia and extract **WEC-Eng**, our English version of a WEC dataset. The automatically-extracted data that we collected provides a training set of a very large scale compared to prior work, while our development and test sets underwent relatively fast manual validation.

Due to the large scale of the WEC-Eng training data, current state-of-the-art CD coreference models cannot be easily trained and evaluated on it, for scalability reasons. We therefore developed a new, more scalable, baseline model for the task, while adapting components of recent competitive



Figure 1: Example of two anchor texts (*100 metres*, *100 m final*) from different Wikipedia articles (Usain Bolt, Yohan Blake) pointing to the same event.

within-document coreference models (Lee et al., 2017; Kantor and Globerson, 2019; Joshi et al., 2019). In addition to setting baseline results for WEC-Eng, we assess our model’s competitiveness by presenting a new state-of-the-art on the commonly used ECB+ dataset. Finally, we propose that our automatic extraction and manual validation methods may be applied to generate additional annotated datasets, particularly for other languages. Overall, we suggest that future cross-document coreference models should be evaluated also on the WEC-Eng dataset, and address its complementary characteristics, while the WEC methodology may be efficiently applied to create additional datasets. To that end, our dataset and code¹² are released for open access.

2 Related Datasets

This section describes the main characteristics of notable datasets for CD event coreference (ECB+, MEANTIME, GVC). Table 1 presents statistics for all these datasets, as well as ours. We further refer to the Wikilinks dataset, which also leveraged Wikipedia links for CD coreference detection.

2.1 CD Event Corpora

ECB+ This dataset (Cybulska and Vossen, 2014), which is an extended version of the EventCoref-Bank (ECB) (Bejan and Harabagiu, 2010), is the most commonly used dataset for training and testing models for CD event coreference (Choubey and

¹WEC-<https://github.com/AlonEirew/extract-wec>

²Model-<https://github.com/AlonEirew/cross-doc-event-coref>

	Topics	Mentions	Clusters	Non-Singleton Clusters	Ambiguity	Diversity
WEC-Eng (train)	-	40,529	7,042	6,210	6.3	2.1
WEC-Eng (dev)	-	1,250	233	216	2.7	1.7
WEC-Eng (test)	-	1,893	322	306	2.6	1.9
ECB+ (train)	25	3,808	1,527	411	1.4	2.2
ECB+ (dev)	8	1,245	409	129	1.4	2.3
ECB+ (test)	10	1,780	805	182	1.4	2.2
GVC (all)	1	7,298	1,411	1,046	19.5	3.0
MEANTIME (all)	4	2,107	1,892	142	2.1	1.5

Table 1: Event Coreference Datasets Statistics (No train/dev/test split exists for MEANTIME and GVC). **Non-Singleton Clusters**: Number of clusters with more than a single mention. **Ambiguity**: Average number of different clusters in which a head lemma appears. **Diversity**: Average number of unique head lemmas within a cluster (excluding singletons for fair comparison).

Huang, 2017; Kenyon-Dean et al., 2018; Barhom et al., 2019). This corpus consists of documents partitioned into 43 clusters, each corresponding to a certain news topic. In order to introduce some ambiguity and to limit the use of lexical features, each topic is composed of documents describing two different events (called *sub-topics*) of the same event type (e.g. two different celebrities checking into rehab facilities). Nonetheless, as can be seen in Table 1, the ambiguity level obtained is still rather low. ECB+ is relatively small, where on average only 1.9 sentences per document were selected for annotation, yielding only 722 non-singleton coreference clusters in total.

MEANTIME Minard et al. (2016) proposed a dataset that is similar in some respects to ECB+, with documents partitioned into a set of topics. The different topics do not correspond to a specific news event but rather to a broad topic of interest (e.g. Apple, stock market). Consequently, different documents rarely share coreferring event mentions, resulting in only 11 event coreference clusters that include mentions from multiple documents, making this dataset less relevant for training CD coreference models.

Gun Violence Corpus (GVC) This dataset (Vossen et al., 2018) was triggered by the same motivation that drove us, of overcoming the huge complexity of direct manual annotation of CD event coreference from scratch. To create the dataset, the authors leveraged a structured database recording gun violence events, in which the record for an individual event points at documents describing that event. The annotators were then asked to examine the linked documents and mark in them

mentions of 5 gun-violence event classes (firing a gun, missing, hitting, injuring, death). Considering the recorded event as a pivot, all mentions found for a particular class were considered as coreferring. Using this process, they report an annotation rate of about 190 mentions per hour. As this corpus assumes a specific event structure scheme related to gun violence, it is more suitable for studying event coreference within a narrow domain rather than for investigating models for broad coverage event coreference.

2.2 Wikilinks

Wikilinks (Singh et al., 2012) is an automatically-collected large-scale cross-document coreference dataset, focused on entity coreference. It was constructed by crawling a large portion of the web and collecting as mentions hyperlinks pointing to Wikipedia articles. Since their method does not include mention distillation or validation, it was mostly used for training models for the Entity Linking task, particularly in noisy texts (Chisholm and Hachey, 2015; Eshel et al., 2017).

3 The WEC Methodology and Dataset

We now describe our methodology for gathering a CD event coreference dataset from Wikipedia, and the WEC-Eng dataset created by applying it to the English Wikipedia. We also denote how this methodology can be applied, with some language-specific adjustments, to other Wikipedia languages.

3.1 Dataset Structure

Our data is collected by clustering together anchor texts of (internal) Wikipedia links pointing to the same Wikipedia concept. This is generally justified

Cluster-1 (2010 Polish Air Force Tu-154 crash)	Cluster-2 (Lokomotiv Yaroslavl plane crash)
...following the death of President Lech Kaczyński in <i>a plane crash</i>On 7 September 2011, nearly the entire team perished in <i>an airplane crash</i> ...
...following the <i>Smolensk air disaster</i> which killed the incumbent Polish president...	...fourth season was overshadowed by the <i>Yaroslavl air disaster</i> on 7 September...
...died when the <i>presidential plane went down</i> about a half mile from the runway...	...Early in September, <i>tragedy</i> rocked the hockey world...

Table 2: Mentions from two generated WEC-Eng clusters, illustrating typical challenging ambiguity. To resolve, the model must consider the surrounding paragraph extracted along with the mention (not fully shown here).

since all these links refer to the same real world theme described by that article, as illustrated in Figure 1. Accordingly, our dataset consists of a set of mentions, each including the mention span corresponding to the link anchor text, the surrounding context, and the mention cluster ID. Since Wikipedia is not partitioned into predefined topics, mentions can corefer across the entire corpus (unlike most prior datasets).

Since mention annotation is not exhaustive, coreference resolution is performed over the gold mentions. Thus, our goal is to support the development of CD event coreference algorithms, rather than of mention extraction algorithms. Our dataset also includes metadata information, such as source and target URLs for the links, but these are not part of the data to be considered by algorithms, as our goal in this work is CD coreference development rather than Event Linking (Nothman et al., 2012).

3.2 Data Collection Process

In this paper, we focus on deriving from Wikipedia an *event* coreference dataset. The choice to focus on event coreference was motivated by two observations: (1) coreference resolution for Wikipedia anchor texts would be more challenging for event mentions than for entity mentions, since the former exhibits much higher degrees of both ambiguity and lexical diversity, and (2) event structures, with their arguments (such as participants, location and time) available in the surrounding context, would facilitate a more natural dataset for the corpus-wide CD coreference task, compared to Wikipedia entity mentions which are comprised mostly of named entities.

Accordingly, we seek to consider only Wikipedia pages denoting events, then collect hyperlinks pointing at these pages. All anchor texts pointing to the same event then become the mentions of a corresponding event coreference cluster, and are extracted along with their surrounding paragraph as context (see Table 2). The following paragraphs describe this process in detail, and how it was applied to generate the WEC-Eng dataset from English

Wikipedia.

Event Identification Many Wikipedia articles contain an infobox³ element. This element can be selected by a Wikipedia author from a pre-defined list of possible infobox types (e.g. “Civilian Attack”, “Game”, “Scientist”, etc.), each capturing typical information fields for that type of articles. For example, the “Scientist” infobox type consists of fields such as “birth date”, “awards”, “thesis” etc. We leverage the infobox element and its parameters in order to identify articles describing events (e.g. accident, disaster, conflict, ceremony, etc.) rather than entities (e.g. a person, organization, etc.).

To that end, we start by automatically compiling a list of all Wikipedia infobox types that are associated with at least dozens of Wikipedia articles. Of those, we manually identify all infobox types related to events (WEC-Eng examples include Awards, Meetings, Civilian Attack, Earthquake, Contest, Concert and more). We then (manually) exclude infobox types that are frequently linked from related but non-corefering mentions, such as sub-events or event characteristics, like location and time (see Appendix A.1 for further details). For WEC-Eng, we ended up with 28 English Wikipedia event infobox types (see Appendix A.2 for the full list).

Gathering Initial Dataset Once the infobox event list is determined, we apply a fully automatic pipeline to obtain an initial crude version of the dataset. This pipeline consists of: (1) Collecting all Wikipedia articles (event “pivot” pages) whose Infobox type is in our list. (2) Collecting all Wikipedia anchor texts (“mentions”) pointing to one of the pivot pages, along with their surrounding paragraph. (3) Filtering mentions that lack context or those belonging to Wikipedia metadata, such as tables, images, lists, etc., as well as mentions whose surrounding context contains obvious Wikipedia boilerplate code (i.e. HTML and JSON tags) (4) Finally, all collected mentions are clustered accord-

³<https://en.wikipedia.org/wiki/Help:Infobox>

Cluster	Mention link and context	Validation
1	..The Tar Heels made it to the final of the 2001 ACC Men's Basketball TournamentDuke finished the regular season with a 26–4 record entering the ACC Tournament ..	✓ ✗ Insufficient mention context
2	..was nominated for an Academy Award at the 1939 Oscars for Best.. .. <i>In the following year</i> , the Academy presented Disney an Honorary Academy Award..	✓ ✗ Link placed on event time
3	..was kidnapped during the Draper's Meadow massacre in July 1755.. ..with a terrible attack on Drapers Meadow by the Shawnee Indian Warriors..	✓ ✗ Link placed on event location
4	..War and was on board the USS "Forrestal" in 1967 when a fatal fire broke out.. ..the repairs to the supercarrier which had been extensively damaged off Vietnam..	✓ ✗ Subevent

Table 3: Examples of WEC-Eng manually validated mentions, showing validation results (validated/disqualified)

ing the pivot page at which they point.

Mention-level Filtering An event coreference dataset mined this way may still require some refinement in order to further clean the dataset at the individual mention level. Indeed, we observed that many Wikipedia editors have a tendency to position event hyperlinks on an event argument, such as a Named Entity (NE) related to the event date or location (as in the case of the disqualified mention for cluster 3 in Table 3). To automatically filter out many of the cases where the hyperlink is placed on an event argument instead of on the event mention itself, we use a Named Entity tagger and filter out mentions identified by one of the following labels: PERSON, GPE, LOC, DATE and NORP (for WEC-Eng we used the SpaCy Named Entity tagger (Honnibal and Montani, 2017)).

Controlling Lexical Diversity So far, we addressed the need to avoid having invalid mentions in a cluster, which do not actually refer to the linked pivot event.

Next, we would like to ensure a reasonably balanced lexical distribution of the mentions within each cluster. Ideally, it would be desired to preserve the “natural” data distributions as much as possible. However, we observed that in many Wikipedia hyperlinks, the anchor texts used in an event mention may be lexically unbalanced. Indeed, Wikipedia authors seem to have a strong bias to use the pivot article title when creating hyperlinks pointing at that article, while additional ways by which the event can be referred are less frequently hyperlinked. Consequently, preserving the original distribution of hyperlink terms would create a too low level of lexical diversity. As a result, training a model on such a dataset might overfit to identifying only the most common mention phrases, leaving little room for identifying the less frequent ones.

To avoid this, we applied a simple filter⁴ that allows a maximum of 4 mentions having identical strings in a given cluster. This hyperparameter was tuned by making the lexical repetition level in our clusters more similar to that of ECB+, in which lexical diversity was not controlled (resulting with an average of 1.9 same-string mentions per cluster in WEC-Eng train set compared to 1.3 in ECB+).

Using this process we automatically generated a large dataset. We designated the majority of the automatically generated data to serve as the WEC-Eng training set. The remaining data was left for the development and test sets, which underwent a manual validation phase, as described next.

3.3 Manual Validation

Inevitably, some noise would still exist in the automatically derived dataset just described. While partially noisy training data can be effective, as we show later, and is legitimate to use, the development set, and particularly the test set, should be of high quality to allow for proper evaluation. To that end, we manually validated the mentions in the development and test sets.

For CD coreference evaluation, we expect to include a mention as part of a coreferring cluster only if it is clear, at least from reading the given surrounding context, that this mention indeed refers to the linked pivot event. Otherwise, we cannot expect a system to properly detect coreference for that mention. Such cases occasionally occur in Wikipedia, where identifying context is missing while relying on the provided hyperlink (see cluster-1 in Table 3, where the tournament year is not mentioned). Such mentions are filtered out by the manual validation. Additionally, misplaced mention boundaries that do not include the correct event trigger (Table 3 cluster-2,3), as well as

⁴We also release the non controlled version of the dataset.

mentions of subevents of the linked event (Table 3 cluster-4), are filtered.

Summing up, to filter out these cases, we used a strict and easy-to-judge manual validation criterion, where a mention is considered valid only if: (1) the mention boundaries contain the event trigger phrase; (2) the mention’s surrounding paragraph suffices to verify that this mention refers to the pivot page and thus belongs to its coreference cluster; and (3) the mention does not represent a subevent of the referenced event. Table 3 shows examples of validated vs. disqualified mentions judged for the WEC-Eng development set.

For the WEC-Eng validation, we randomly selected 588 clusters and validated them, yielding 1,250 and 1,893 mentions for the development and test sets, respectively. Table 1 presents further statistics for WEC-Eng. The validation was performed by a competent native English speaker, to whom we explained the guidelines, after making a practice session over 150 mentions. Finally, all training mentions that appeared in the same (source) article with a validated mention were discarded from the training set.

Our manual validation method is much faster and cheaper compared to a full manual coreference annotation process, where annotators would need to identify and compare all mentions across all documents. In practice, the average annotation rate for WEC-Eng yielded 350 valid mentions per hour, with the entire process taking only 9 hours to complete. In addition, since our validation approach is quite simple and does not require linguistic expertise, the eventual data quality is likely to be high. To assess the validation quality, one of the authors validated 50 coreference clusters (311 mentions), randomly selected from the development and test sets, and then carefully consolidated these annotations with the original validation judgements by the annotator. Relative to this reliable consolidated annotation, the original annotations scored at 0.95 Precision and 0.96 Recall, indicating the high quality of our validated dataset (the Cohen’s Kappa (Cohen, 1960) between the original and consolidated annotations was 0.75, considered substantial agreement).

In all, 83% of the candidate mentions were positively validated in the development and test sets, indicating a rough estimation of the noise level in the training set. That being said, we note that a majority of these noisy mentions were not totally

wrong mentions but rather were filtered out due to the absence of substantial surrounding context or the misplacement of mention boundaries (see examples in Table 3).

3.4 Dataset Content

This section describes the WEC-Eng dataset content and some of its characteristics. The final WEC-Eng dataset statistics are presented in Table 1. Notably, the training set includes 40,529 mentions distributed into 7,042 coreference clusters, facilitating the training of deep learning models.

The relatively high level of lexical ambiguity shown in the table⁵ is an inherent characteristic caused by many events (coreference clusters) sharing the same event type, and thus sharing the same terms, as illustrated in Table 2. Identifying that identical or lexically similar mentions refer to different events is one of the major challenges for CD coreference resolution, particularly in the corpus-wide setting, where different documents might refer to similar yet different events.

With respect to the distinction between *descriptive* and *referential* event mentions, proposed in Section 1, WEC-Eng mentions are predominantly referential. This stems from the fact that its mentions correspond to hyperlinks that point at a different Wikipedia article, describing the event, while the mention’s article is describing a different topic. On the other hand, ECB+, being a news dataset, is expected to include predominantly descriptive mentions. Indeed, manually analyzing a sample of 30 mentions from each dataset, in WEC-Eng 26 were referential while in ECB+ 28 were descriptive. This difference also imposes different lexical distributions for mentions in the two datasets, as sampled in Appendix A.3. When describing an event, verbs are more frequently used as event mentions, but nominal mentions are abundant as well. This is apparent for the predominately descriptive ECB+, where 62% of the mentions in our sample are verbal vs. 38% nominal. On the other hand, when a previously known event is only referenced, it is mostly referred by a nominal mention. Indeed, in the predominantly referential WEC-Eng, a vast majority of the mentions are nominal (93% in our sample).

⁵The higher ambiguity level of the training sets stems from its larger size — including many more coreference clusters in total, and accordingly more clusters under individual event types.

3.5 Potential Language Adaptation

While our process was applied to the English Wikipedia, it can be adapted with relatively few adjustments and resources to other languages for which a large-scale Wikipedia exists. Here we summarize the steps needed to apply our dataset creation methodology to other Wikipedia languages. To generate a dataset, the first step consists of manually deciding on the list of suitable infobox types corresponding to (non-noisy) event types. Then, the automatic corpus creation process can be applied for this list, which takes only a few hours to run on a single CPU. After the initial dataset was created, a language specific named-entity tagger should be used to filter mentions of certain types, like time and location (see Mention-level Filtering (3.2)). Next, the criterion for ensuring balanced lexical diversity in a cluster (see Controlling Lexical Diversity (3.2)), which was based on a simple same-string test for English, may need to be adjusted for languages requiring a morphological analyzer. Finally, as we perform manual validation of the development and test sets, this process should be performed for any new dataset (see Section 3.3). Supporting this step, our validation guidelines are brief and simple, and are not language specific. They only require identifying subevents and misplaced mention boundaries, as well as validating the sufficiency of the mention context.

4 Baseline and Analysis

4.1 Model

The current state-of-the-art CD event coreference system (Barhom et al., 2019) cannot be effectively trained on WEC-Eng for two main reasons: (1) computational complexity and (2) reliance on verbal SRL features. With respect to computation time, the training phase of this model simulates the clustering operations done at inference time, while recalculating new mention representations and pairwise scores after each cluster merging step. Consequently, training this model on our large scale training data, which is further not segmented to topics, is computationally infeasible. In addition, the model of Barhom et al. (2019) uses an SRL system to encode the context surrounding *verbal* event mentions, while WEC-Eng is mostly composed of *nominal* event mentions (Section 3.4).

We therefore develop our own, more scalable, model for CD event coreference resolution, estab-

lishing baseline results for WEC-Eng. As common in CD coreference resolution, we train a pairwise scorer $s(i, j)$ indicating the likelihood that two mentions i and j in the dataset are corefering, and then apply agglomerative clustering over these scores to find the coreference clusters. Following the commonly used average-link method (Choubey and Huang, 2017; Kenyon-Dean et al., 2018; Barhom et al., 2019), the merging score for two clusters is defined as the average mention pair score $s(i, j)$ over all mention pairs (i, j) across the two candidate clusters to be merged.

For the pairwise model, we replicate the architecture of mention representation and pairwise scorer from the end-to-end within document coreference model in (Lee et al., 2017), while including the recent incorporation of transformer-based encoders (Joshi et al., 2019; Kantor and Globerson, 2019). Concretely, we first apply a pre-trained RoBERTa (Liu et al., 2019) language model (without fine-tuning), separately for each mention. Given a mention span i , we include as context, T (set to 250) tokens to the left of i and T tokens to the right of i . Applying RoBERTa to this window, we represent each mention by a vector g_i , which is the concatenation of three vectors: the contextualized representations of the mention span boundaries (first and last) and the weighted sum of the mention token vectors according to the head-finding attention mechanism in (Lee et al., 2017). The two mention representations g_i and g_j , and the element-wise multiplication of these vectors are then concatenated and fed into a simple MLP, which outputs a score $s(i, j)$, indicating the likelihood that mentions i and j belong to the same cluster. The head-attention layer and the MLP are trained to optimize the standard binary cross-entropy loss over all pairs of mentions, where the label is 1 if they belong to the same coreference cluster and 0 otherwise.⁶

4.2 Experiments

We first train and evaluate our model on the commonly used dataset ECB+, to assess its relevance as an effective baseline model, and then evaluate it on WEC-Eng, setting baseline results for our dataset. We also present the performance of the challenging *same-head-lemma* baseline, which clusters mentions sharing the same syntactic-head lemma. For

⁶We note that this optimization is different than the (linear) antecedent ranking in the model of Lee et al. (2017), since in the CD setting there is no linear order between mentions from different documents.

	MUC			B^3			CEAF			CoNLL
	R	P	F_1	R	P	F_1	R	P	F_1	F_1
Lemma Baseline	76.5	79.9	78.1	71.7	85	77.8	75.5	71.7	73.6	76.5
DISJOINT (Barhom et al., 2019)	75.5	83.6	79.4	75.4	86	80.4	80.3	71.9	75.9	78.5
JOINT (Barhom et al., 2019)	77.6	84.5	80.9	76.1	85.1	80.3	81	73.8	77.3	79.5
Our model	84.2	81.8	83	80.8	81.7	81.3	76.7	79.5	78	80.8

Table 4: Event coreference results on the ECB+ test set

	MUC			B^3			CEAF			CoNLL
	R	P	F_1	R	P	F_1	R	P	F_1	F_1
Lemma Baseline	85.5	79.9	82.6	74.5	32.8	45.5	25.9	39.4	31.2	53.1
Our model	78	83.6	80.7	66.1	55.3	60.2	53.4	40.3	45.9	62.3

Table 5: Event coreference results on the WEC-Eng test set

the experiment on ECB+, we follow the recent evaluation setting (Kenyon-Dean et al., 2018; Barhom et al., 2019), clustering gold mentions and concatenating all test documents into one meta-document, as proposed by Upadhyay et al. (2016). For a fair comparison, we use the output of pre-processing document clustering obtained by Barhom et al. (2019) and apply our coreference model separately on each predicted document cluster.

For both datasets, the positive examples for training consist of all the mention pairs in the dataset that belong to the same coreference cluster. For the ECB+ model, we consider only negative examples that belong to the same subtopic, while for WEC-Eng we sample k (tuned to 10) negative examples for each positive one. Results are reported by precision, recall, and F1 for the standard coreference metrics MUC, B^3 , CEAF-e, and the average F1 of the three metrics, using the official CoNLL scorer (Pradhan et al., 2012).⁷

4.3 Results

Table 4 presents the results on ECB+. Our model outperforms state-of-the-art results for both the JOINT model and the DISJOINT event model of Barhom et al. (2019), with a gain of 1.3 CoNLL F_1 points and 2.3 CoNLL F_1 points respectively. The JOINT model jointly clusters event and entity mentions, leveraging information across the two subtasks, while the DISJOINT event model considers only event mentions, taking the same input as our model. These results assess our model as a suitable baseline for WEC-Eng.

⁷<https://github.com/conll/reference-coreference-scorers>

Table 5 presents the results on WEC-Eng. First, we observe that despite the certain level of noise in the automatically gathered training data, our model outperforms the same-head-lemma baseline by 9.2 CoNLL F_1 points. In fact, it achieves similar error reduction rates relative to the lemma baseline as obtained over ECB+, where training is performed on a clean but smaller training data (18.3% error reduction in ECB+ and 19.6% in WEC). Furthermore, the performance of both the same-head-lemma baseline and our model are substantially lower on WEC-Eng (Table 5) than on ECB+ (Table 4). This indicates the more challenging nature of WEC-Eng, possibly due to its corpus wide nature and higher degree of ambiguity (Tables 1).

Further examining the different nature of the two datasets, we applied cross-domain evaluation, applying the ECB+ trained model on WEC-Eng test data and vice versa. The results suggest that due to their different characteristics, with respect to mention type (descriptive vs. referential) and structure (topic-based vs. corpus wide), a model trained on one dataset is less effective (by 8-12 points) when applied to the other (further details are presented in Appendix B.1).

4.4 Qualitative Analysis

To obtain some qualitative insight about the learned models for both ECB+ and WEC-Eng, we manually examined their most certain predictions, looking at the top 5% instances with highest predicted probability and at the bottom 5%, of lowest predictions. Some typical examples are given in Appendix B.2. Generally, both models tend to assign the highest probabilities to mention pairs that

share some lemma, and occasionally to pairs with different lemmas with similar meanings, with the WEC-Eng model making such lexical generalizations somewhat more frequently. Oftentimes in these cases, the models fail to distinguish between (gold) positive and negative cases, despite quite clear distinguishing evidence in the context, such as different times or locations. This suggests that the RoBERTa-based modeling of context may not be sufficient, and that more sophisticated models, injecting argument structure more extensively, may be needed.

In both models, the lowest predictions (correctly) correspond mostly to negative mention pairs, and occasionally to positive pairs for which the semantic correspondence is less obvious (e.g. *offered* vs. *candidacy*). In addition, we observe that longer spans common in WEC-Eng challenge the span representation model of Lee et al. (2017). This model emphasizes mention boundaries, but these often vary across lexically-similar coreferring mentions with different word order.

5 Conclusion and Future Work

In this paper, we presented a generic low-cost methodology and supporting tools for extracting cross-document event coreference datasets from Wikipedia. The methodology was applied to create the larger-scale WEC-Eng corpus, and may be easily applied to additional languages with relatively few adjustments. Most importantly, our dataset complements existing resources for the task by addressing a different appealing realistic setup: the targeted data is collected across a full corpus rather than within topical document clusters, and, accordingly, mentions are mostly referential rather than descriptive. Hence, we suggest that future research should be evaluated also on WEC-Eng, while future datasets, particularly for other languages, can be created using the WEC methodology and tool suite, all made publicly available. Our released model provides a suitable baseline for such future work.

Acknowledgments

We would like to thank Valentina Pyatkin, Daniela Stepanov and Oren Pereg for their valuable assistance in the data validation process. The work described herein was supported in part by grants from Intel Labs, Facebook, the Israel Science Foundation grant 1951/17, the Israeli Ministry of Sci-

ence and Technology and the German Research Foundation through the German-Israeli Project Cooperation (DIP, grant DA 1600/1-1).

References

- Jun Araki, Zhengzhong Liu, Eduard Hovy, and Teruko Mitamura. 2014. [Detecting subevent structure for event coreference resolution](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Shany Barhom, Vered Shwartz, Alon Eirew, Michael Bugert, Nils Reimers, and Ido Dagan. 2019. [Revisiting joint modeling of cross-document entity and event coreference resolution](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4179–4189, Florence, Italy. Association for Computational Linguistics.
- Cosmin Bejan and Sanda Harabagiu. 2010. [Unsupervised event coreference resolution with rich linguistic features](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1412–1422, Uppsala, Sweden. Association for Computational Linguistics.
- Andrew Chisholm and Ben Hachey. 2015. [Entity disambiguation with web links](#). *Transactions of the Association for Computational Linguistics*, 3:145–156.
- Prafulla Kumar Choubey and Ruihong Huang. 2017. [Event coreference resolution by iteratively unfolding inter-dependencies among events](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2124–2133, Copenhagen, Denmark. Association for Computational Linguistics.
- Jacob Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and Psychological Measurement*, 20(1):37–46.
- Agata Cybulska and Piek Vossen. 2014. Using a sledgehammer to crack a nut? Lexical diversity and event coreference resolution. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 4545–4552, Reykjavik, Iceland. European Languages Resources Association (ELRA).
- Bhuwan Dhingra, Qiao Jin, Zhilin Yang, William Cohen, and Ruslan Salakhutdinov. 2018. [Neural models for reasoning over multiple mentions using coreference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 42–48, New Orleans, Louisiana. Association for Computational Linguistics.

- Yotam Eshel, Noam Cohen, Kira Radinsky, Shaul Markovitch, Ikuya Yamada, and Omer Levy. 2017. [Named entity disambiguation for noisy text](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 58–68, Vancouver, Canada. Association for Computational Linguistics.
- Tobias Falke, Christian M. Meyer, and Iryna Gurevych. 2017. [Concept-map-based multi-document summarization using concept coreference resolution and global importance optimization](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 801–811, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- Eduard Hovy, Teruko Mitamura, Felisa Verdejo, Jun Araki, and Andrew Philpot. 2013. Events are not simple: Identity, non-identity, and quasi-identity. In *Workshop on events: Definition, detection, coreference, and representation*, pages 21–28.
- Mandar Joshi, Omer Levy, Luke Zettlemoyer, and Daniel Weld. 2019. [BERT for coreference resolution: Baselines and analysis](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5802–5807, Hong Kong, China. Association for Computational Linguistics.
- Ben Kantor and Amir Globerson. 2019. [Coreference resolution with entity equalization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 673–677, Florence, Italy. Association for Computational Linguistics.
- Kian Kenyon-Dean, Jackie Chi Kit Cheung, and Doina Precup. 2018. [Resolving event coreference with supervised representation learning and clustering-oriented regularization](#). In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 1–10, New Orleans, Louisiana. Association for Computational Linguistics.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. [End-to-end neural coreference resolution](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197, Copenhagen, Denmark. Association for Computational Linguistics.
- Kexin Liao, Logan Lebanoff, and Fei Liu. 2018. [Abstract Meaning Representation for multi-document summarization](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1178–1190, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Xueling Lin, Haoyang Li, Hao Xin, Zijian Li, and Lei Chen. 2020. KBPearl: A Knowledge Base Population System Supported by Joint Entity and Relation Linking. *Proceedings of the VLDB Endowment*, 13(7):1035–1049.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Anne-Lyse Minard, Manuela Speranza, Ruben Urizar, Begoña Altuna, Marieke van Erp, Anneleen Schoen, and Chantal van Son. 2016. MEANTIME, the NewsReader multilingual event and time corpus. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 4417–4422, Portorož, Slovenia. European Language Resources Association (ELRA).
- Joel Nothman, Matthew Honnibal, Ben Hachey, and James R Curran. 2012. Event linking: Grounding event reference in a news archive. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 228–232.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. [CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes](#). In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea. Association for Computational Linguistics.
- Sameer Singh, Amarnag Subramanya, Fernando Pereira, and Andrew McCallum. 2012. Wikilinks: A large-scale cross-document coreference corpus labeled via links to Wikipedia. Technical Report UM-CS-2012-015.
- Shyam Upadhyay, Nitish Gupta, Christos Christodoulopoulos, and Dan Roth. 2016. [Revisiting the evaluation for cross document event coreference](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1949–1958, Osaka, Japan. The COLING 2016 Organizing Committee.
- Piek Vossen, Filip Ilievski, Marten Postma, and Roxane Segers. 2018. [Don’t annotate, but validate: a data-to-text method for capturing event data](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Haoyu Wang, Mo Yu, Xiaoxiao Guo, Rajarshi Das, Wenhan Xiong, and Tian Gao. 2019. [Do multi-hop readers dream of reasoning chains?](#) In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 91–97, Hong Kong, China. Association for Computational Linguistics.

A WEC Dataset

A.1 Infobox Distillation

Excluding infobox types of broad general events, consisting of many sub-events is necessary as often Wikipedia authors link to a broad event from anchor texts that refer to a subevent (and should be regarded as non-coreferring by standard definitions of the event coreference task (Hovy et al., 2013; Araki et al., 2014)). For example, in English Wikipedia many event articles containing the infobox “*Election*” tend to be pointed at from anchor texts that describe subevents, such as *2016 primaries* linking to *2016 United States presidential election*. Additionally, for some infobox types, pointing mentions often correspond to related, but not coreferring, named entities; hence, we discard the infobox types to avoid such noisy mentions. For example, articles of the infobox type “Race” are often linked from mentions denoting the name of the country in which the race took place. Table 6 presents excluded infobox types due to the above reasons.

A.2 WEC-Eng Infobox Types

As mentioned in the paper (Section 3), we manually explored the various Wikipedia infobox types and selected only those denoting an event. Table 7 shows the number of coreference clusters (event pages) and mentions for each selected infobox type. In the table, infobox types falling under the same general category are grouped together.

A.3 Lexical Distribution of ECB+ and WEC-Eng Mentions

Event mentions can appear in various lexical forms in a document, such as verbs (e.g. *exploded*), nominalization (e.g. *crash*), common nouns (e.g. *party*, *accident*) and proper nouns (e.g. *Cannes Festival 2016*). In order to have a rough estimation of the distribution of these different forms, we manually analyze 100 sampled mentions, from each of WEC-Eng and ECB+, and present the statistics in Table 8.

B Further Analysis

B.1 Cross-Domain Experiment and Results

To further assess the difference between the ECB+ and WEC-Eng datasets, we evaluate the cross-domain ability of the model trained on one dataset, and tested on the other one. We use the same pairwise model as in Tables 4 and 5, while only tun-

ing the stop criterion for the agglomerative clustering, in the corresponding development set. Table 9 presents the cross-domain performance of our model (1) trained on ECB+ and evaluated on WEC-Eng and (2) trained on WEC-Eng and evaluated on ECB+, together with the corresponding in-domain performances (from Tables 4 and 5) for comparison. The performance drop for both cross-domain experiments indicates that these two datasets address different challenges of cross-document event coreference resolution.

B.2 Qualitative Analysis Examples

Illustrating Section 4.4, Tables 10 and 11 present examples sampled from the predictions of the ECB+ and WEC-Eng models on their respective development sets.

Infobox	Mention link and context	Non-coreferring relation
Race	..He crashed in <i>Italy</i> , and finished two laps off the lead in Portugal..	✗ Location
	..after several unrewarded drives the year before, namely in <i>Italy</i> ..	✗ Location
	..did not appear at the <i>British</i> or German Grands Prix..	✗ Location
Sport	..in the <i>Gold Medal final</i> straight sets by an identical score of 16 – 21..	✗ Subevent
	..2012 FINA World Swimming Championships (<i>25 m</i>) and in the 50 m freestyle..	✗ Subevent
Wars	..throughout most of World War II, in France <i>Holland</i> and Belgium..	✗ Location
	..During the <i>ensuing campaign</i> the Netherlands were defeated and occupied ..	✗ Subevent
	..and made possible further <i>advance into Crimea</i> and industrially developed Eastern..	✗ Subevent

Table 6: Sample of disqualified general infoboxs types, usually linked from many erroneous mentions

Infobox Type	Event Pages	Mentions
Awards (2 Infoboxes)	1,848	8,910
Meetings (4 Infoboxes)	1,212	5,942
Civilian Attack	1,178	9,490
Airliner Accident (6 Infoboxes)	596	2,640
Festival	541	2,770
Beauty Pageant	431	1,616
Earthquake	387	3,526
Contest	268	2,029
Concert	262	1,082
News Event	241	1,663
Terrorist Attack	230	1,296
Wildfire	135	720
Flood	121	990
Weapons Test (2 Infoboxes)	65	345
Eruption	31	343
Solar Eclipse	31	121
Oilspill	18	184
Rail Accident	2	5
Total (28)	7,597	43,672

Table 7: Final 28 WEC-Eng extracted infobox types list

	WEC-Eng	ECB+
# Verbs	7	62
# Proper nouns	38	2
# Nominalizations	24	23
# Other common nouns	31	13

Table 8: Distribution of lexical types for 100 randomly sampled mentions from ECB+ and WEC-Eng

	MUC			B^3			CEAF			CoNLL
	R	P	F_1	R	P	F_1	R	P	F_1	F_1
WEC-Eng to WEC-Eng	78	83.6	80.7	66.1	55.3	60.2	53.4	40.3	45.9	62.3
ECB+ to WEC-Eng	70.1	79.3	74.4	53.8	54.6	54.2	42.9	27.4	33.5	54
ECB+ to ECB+	84.2	81.8	83	80.8	81.7	81.3	76.7	79.5	78	80.8
WEC-Eng to ECB+	77.1	68.5	72.6	74.2	73.1	73.6	56.4	66.5	61	69.1

Table 9: Cross-domain results of model trained on ECB+ and evaluated on WEC-Eng test set, and vice-versa

Gold	Mention Link and Context
	Highest Predicted Probabilities
+	..Nov 12 2011, 01:19 hrs Warship INS Sukanya on Thursday foiled a <i>piracy</i> attempt in the Gulf of Aden.. ..November 11, 2011 18:50 IST Indian Naval ship, INS Sukanya, thwarted a <i>piracy</i> attack in the Gulf of Aden and..
+	..We are thrilled to have Ellen DeGeneres <i>host</i> the Oscars, "producers Craig Zadan and Neil Meron said in a.. ..the Oscars Published August 02 , 2013 Ellen DeGeneres will <i>host</i> the Oscars for a second time next year..
+	..this morning that Ellen DeGeneres would return to the <i>Academy Awards</i> as emcee , marking her second hosting stint.. ..DeGeneres just tweeted she will be hosting the <i>Oscars</i> this year..
-	..A British man has been killed after <i>falling</i> around 2,000ft.. ..A 36-year-old Australian climber is dead after <i>falling</i> about 150 metres..
-	..Comedian Ellen DeGeneres picked to host 2014 <i>Oscars</i> Fri Aug 2, 2013 6:50pm EDT.. ..DeGeneres , the star of her own daytime talk show "Ellen," first hosted the <i>Oscars</i> in 2007, becoming only..
-	..play in Sunday 's AFC divisional <i>playoff</i> game at Pittsburgh.. ..drink-driving five days before a key NFL <i>playoff</i> game..
	Lowest Predicted Probabilities
-	..I'm hosting the Oscars! "DeGeneres <i>tweeted</i> Friday.. ..The comedian and talk show host made the <i>announcement</i> via her Twitter feed..
-	..and chief medical correspondent for CNN, has been <i>offered</i> the post of Surgeon General.. ..Since CNN learned of Gupta's <i>candidacy</i> , the cable network..
-	..Cheeks returns as Sixers' fixer Jim O'Brien was <i>shown the door</i> after a single tumultuous season.. ..The Sixers made a power move this offseason, <i>firing</i> head coach Jim O'Brien..
-	..the California Highway Patrol <i>said</i> . Williams, 32, was going west in his 2008 Bentley.. .. <i>According to</i> the California Highway Patrol, defensive tackle Jamal Williams..

Table 10: Examples of ECB+ highest predicted probabilities and the lowest predicted probabilities, along with their gold standard label.

Gold	Mention Link and Context
	Highest Predicted Probabilities
+	..Bangladesh Army He was the chief organiser of the <i>assassination</i> of Sheikh Mujibur Rahman.. ..When Sheikh Mujibur Rahman was <i>killed</i> on 15 August 1975 by members of the Bangladesh Army..
+	..On 6 April 1968 , the Hall was the host venue for the <i>Eurovision Song Contest</i> which was.. ..that at the <i>1968 Contest</i> , the voting had been rigged by Spanish dictator Francisco Franco and..
+	..Dagestani separatist movement, combined with a series of <i>apartment bombings</i> in Russia.. ..rebellious Chechnya republic as a retaliation for <i>terrorist bombings</i> in Moscow and other cities..
-	..towns and villages were destroyed by the <i>earthquake</i> and alluvium on May 31, 1970.. ..passing through the area struck by an <i>earthquake</i> in May 2012..
-	..sent troops to the rebellious Chechnya republic as a retaliation for <i>terrorist bombings</i> in Moscow and.. ..Alaoui died from injuries suffered in a <i>terrorist attack</i> in Ouagadougou, Burkina Fas..
-	..in monasteries in Austria-Hungary, after the recent <i>massacres</i> and expulsions of Albanians by their Slavic neighbours.. ..gained international fame in 2014 through the <i>genocide</i> of the Islamic State on the Yazidis..
-	..and was a gold Democrat delegate to the <i>1896 Democratic National Convention</i>while watching Pat Buchanan's televised speech at the <i>1996 Republican National Convention</i> ..
	Lowest Predicted Probabilities
-	..As the official representative of her country to the 2011 <i>Miss Universe</i> pageant.. ..as the final was held too late to send the winner to the <i>2011 edition</i> ..
-	..the ground the funeral party came under a hand-grenade <i>attack</i> from a lone Loyalist paramilitary.. ..the Hill bar shooting, the <i>Milltown massacre</i> , the Sean Graham's and James Murray's bookmakers' shootings..
-	..Christian-Jaque's 1946 film "A Lover's Return" was entered into the <i>1946 Cannes Film Festival</i>This film won two prizes at the first <i>Cannes Film Festival in 1946</i> ..
-	..The epicenter of the <i>2006 Kiholo Bay earthquake</i> was some offshore of the village.. ..However, the unit was damaged in an <i>earthquake in October 2006</i> ..

Table 11: Examples of WEC-Eng highest predicted Probabilities and the lowest predicted probabilities, along with their gold standard label.