

Generating Harder Cross-document Event Coreference Resolution Datasets using Metaphoric Paraphrasing

Shafiuddin Rehan Ahmed¹ Zhiyong Eric Wang² George Arthur Baker¹

Kevin Stowe³ James H. Martin¹

Departments of ¹Computer Science & ²CLASIC, University of Colorado, Boulder, USA
{shah7567, zhwa3087}@colorado.edu

³Education Testing Service (ETS)

Abstract

The most widely used Cross-Document Event Coreference Resolution (CDEC) datasets fail to convey the true difficulty of the task, due to the lack of lexical diversity between corefering event triggers (words or phrases that refer to an event). Furthermore, there is a dearth of event datasets for figurative language, limiting a crucial avenue of research in event comprehension. We address these two issues by introducing ECB+META, a lexically rich variant of Event Coref Bank Plus (ECB+) for CDEC on figurative and metaphoric language. We use GPT-4 as a tool for the metaphoric transformation of sentences in the documents of ECB+, then tag the original event triggers in the transformed sentences in a semi-automated manner. In this way, we avoid the re-annotation of expensive coreference links. We present results that show existing methods that work well on ECB+ struggle with ECB+META, thereby paving the way for CDEC research on a much more challenging dataset.¹

1 Introduction

Cross-Document Event Coreference Resolution (CDEC) involves identifying mentions of the same event within and across documents. An issue with CDEC is that the widely used dataset, Event Coref Bank plus (ECB+; Cybulska and Vossen (2014)), is biased towards lexical similarities, both for triggers and associated event arguments, and therefore has a very strong baseline (Cybulska and Vossen, 2015; Kenyon-Dean et al., 2018; Ahmed et al., 2023a). To see this, consider the excerpts from ECB+ shown in Figure 1(a). This consists of three *killing* events selected from separate articles sharing a common trigger. An algorithm capable of matching the triggers and tokens within the sentences, such as "Vancouver" and "office," can readily discern that Event 2 is coreferent with Event 3, and not Event 1. This

leads to the question of whether the state-of-the-art methods using this corpus (Held et al., 2021) learn the semantics of event coreference, or are merely exploiting surface triggers.

Figurative language, encompassing metaphors, similes, idioms, and other non-literal expressions, is an effective tool for assessing comprehension across cognitive, linguistic, and social dimensions (Lakoff and Johnson, 1980; Winner, 1988; Gibbs, 1994; Palmer and Brooks, 2004; Palmer et al., 2006). Figurative language, by its nature, draws on a wide array of cultural, contextual, and imaginative resources to convey meanings in nuanced and often novel ways. Consequently, it employs a broader vocabulary and more unique word combinations than literal language (Stefanowitsch, 2006). Most recent work on metaphors has been focused on generation (Stowe et al., 2020, 2021b; Chakrabarty et al., 2021a), interpretation (Chakrabarty et al., 2022, 2023), and detection (Li et al., 2023; Joseph et al., 2023; Wachowiak and Gromann, 2023). Yet, there is a dearth of event datasets for figurative language which limits an important research direction of event comprehension.

In this paper, we address these two challenges by leveraging GPT-4 in *constrained metaphoric paraphrasing* of ECB+documents. We introduce a novel dataset named ECB+META, which we generate using a semi-automatic approach. This involves applying metaphoric transformations to the event triggers within ECB+ and then hand-correcting the tagged triggers in the new corpus. As depicted in Figure 1(b), the trigger word *killing* in Events 2 and 3 of ECB+ become *slaying* and *snuffing out the flame of life of* in ECB+META, respectively.

This approach preserves the coreference annotations from ECB+, thereby avoiding an expensive coreference re-annotation task. Thus, we create several versions of "tougher" CDEC benchmark datasets with enhanced lexical diversity with varying levels of metaphoricity. We present baseline

¹Code/data: github.com/ahmeshaf/llms_coref

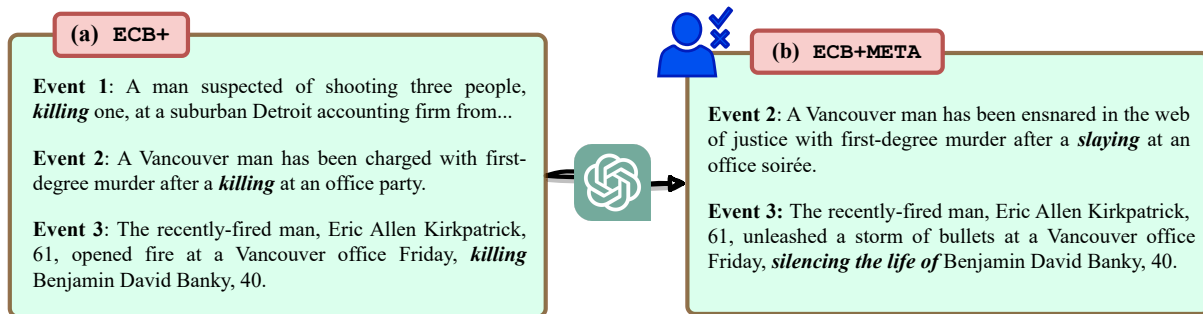


Figure 1: Using GPT-4 to Generate ECB+META from ECB+Corpus. **Event 2 & Event 3** are coreferent, while **Event 1** is not. ECB+META has metaphorically transformed triggers, e.g., *killing* -> *silencing the life*. The triggers are hand-corrected by an annotator. ECB+META challenges previous work—Held et al. (2021) & Ahmed et al. (2023a).

results using previous methods—Held et al. (2021) and Ahmed et al. (2023a) (described in §3.2), and show the limitation of these approaches on this dataset. Finally, we correlate lexical diversity and text complexity with CDEC and test the hypothesis that CDEC gets more difficult as the lexical diversity/complexity of the corpus increases.

2 Related Work

2.1 CDEC Datasets

ECB+² is the most widely used dataset for CDEC, yet it has limited utility in realistic applications because of how simple the dataset is. The Gun Violence Corpus (GVC; Vossen et al. (2018)), for instance, was introduced as a way of adding ambiguity to the task. Yet, both these datasets lack lexical diversity in terms of coreferent event triggers. Ravenscroft et al. (2021) is one such work that addresses the diversity question through cross-domain coreference, however, a dataset focusing CDEC on figurative language does not exist to our best knowledge.

Even with the use of modern annotation tools (Klie et al., 2018; Ahmed et al., 2023b), annotating CDEC datasets is expensive. Works such as Bugert and Gurevych (2021); Eirew et al. (2021) use Wikipedia as a way of bootstrapping ECR annotations automatically. In a similar vein, we bootstrap CDEC annotations for figurative language in a synthetic way using GPT-4.

2.2 Metaphoric Paraphrasing

The task of metaphoric paraphrasing has been explored through a variety of methods. A primary theme is sentential paraphrasing by replacing literal words with metaphors (Stowe et al., 2021a,b; Chakrabarty et al., 2021b). These approaches fine tune language models with control codes to indicate

metaphors, exploiting available metaphoric data to facilitate transformations from literal language to metaphoric. However, they rely on extensive data, and there is evidence that modern large language models excel at metaphor generation (Chakrabarty et al., 2023) and paraphrasing (Kojima et al., 2023; OpenAI, 2023). For this reason, we leverage GPT-4 via ChatGPT functionality for our experiments.

2.3 CDEC Methods

Non-filtering Methods: Previous works (Meged et al., 2020; Zeng et al., 2020; Cattan et al., 2021; Allaway et al., 2021; Caciularu et al., 2021; Yu et al., 2022) in CDEC have been successful using pairwise mention representation learning models, a method popularly known as cross-encoding. These methods use distributed and contextually-enriched “non-static” vector representations of mentions from Transformer-based (Vaswani et al., 2017) language models like various BERT-variants (Devlin et al., 2019; Beltagy et al., 2020) to calculate supervised pairwise scores for those event mentions. While these methods demonstrate SoTA performance, their applicability is hindered by their quadratic complexity at inference.

Filtering Methods: Keeping usability and tractability in mind, we experiment only with the recent work that adds a low-compute mention pair filtering step before crossencoding. These approaches aid in the removal of numerous irrelevant mention pairs, thereby directing focus toward the most pertinent pairs with resource-intensive models. For instance, in their work, Held et al. (2021) propose a retrieval, vector-based K-nearest neighbor method, that helps find and focus only on the hard negatives in the corpus. In contrast, Ahmed et al. (2023a) employ simplified lexical similarity metrics to filter out a substantial number of truly non-coreferent pairs in the corpus.

²Corpus detailed in §A

3 Methodology

We first synthetically create ECB+META by employing metaphoric paraphrasing of the original corpus. Then we tag the event triggers of the original corpus in ECB+META in a semi-automated manner. Finally, we adopt two existing CDEC methods to test this new dataset. We describe each of these steps:

3.1 Metaphoric Paraphrasing using GPT-4

We paraphrase ECB+'s sentences in a constrained manner in which we convert only the event triggers in a sentence into metaphors. We first extract the event mentions from each sentence of the documents in the corpus, then prompt GPT-4 to convert only the trigger words in the sentence to metaphors. We adopt a chain of thought prompting approach (Kojima et al., 2022), where we provide the steps that need to be followed in the conversion (see §B).

To enhance diversity and sample appropriate metaphors, we generate five metaphors for every trigger word in the sentence and then task GPT-4 to select the most coherent one from the list. We diversify metaphoricity levels by using both single-word and multi-word metaphors. As illustrated in Figure 3, the conversion of "killing" into a single-word metaphor is "slaying," while its transformation into a multi-word phrasal metaphor is "extinguishing the candle of life." We develop two versions of ECB+META, designated as ECB+META₁ for single-word transformations and ECB+META_m for multi-word transformations, respectively.

Using the generated conversions, we first automatically tag the original events in the transformed sentences. Then, we hand-correct cases where the conversion is ambiguous. In the end, we are left with two versions of the validation and the test sets of ECB+META preserving the original coreference annotations of ECB+.

3.2 CDEC Methods

Filtering Step for CDEC: The BiEncoder K-NN (KNN) approach, introduced by Held et al. (2021) involves a novel approach to mention pair retrieval before doing CDEC. This method focuses on selecting mentions that are most similar to a given target mention using their static vector representations and a Vector Store (like FAISS Johnson et al. (2019)). To achieve this, they fine-tune the RoBERTa-Large (Liu et al., 2019) pre-trained model using a contrastive Categorical Loss function, with categories corresponding to event clusters within the corpus. This fine-tuning process uti-

lizes token embeddings generated by the language model and trains on the centroid representations of gold standard event clusters. Due to computation constraints, we use RoBERTa-Base instead of RoBERTa-Large in this work. For the same reason, we use triplet-loss with mention pairs instead of the centroid of clusters.

The Lemma Heuristic (LH; Ahmed et al. (2023a)) leverages lexical features to pre-filter non-coreferent pairs before CDEC. This way, they eliminate the need for an additional fine-tuning step as required in the KNN approach. LH focuses on creating a balanced set of coreferent and non-coreferent pairs while minimizing the inadvertent exclusion of coreferent pairs (false negatives) by the heuristic. It accomplishes this by first generating a set of synonymous lemma pairs from the training corpus and then applying a sentence-level word overlap ratio to prune pairs that don't meet the threshold or lack synonymy. In this work, we use the LH method for filtering and also as a baseline lexical method following Ahmed et al. (2023a).

Cross-encoder³: The Cross-Encoder (CE) functions within CDEC as a pairwise classifier, leveraging joint representations of a mention pair (e_i, e_j) . First, it combines the two event mentions with their respective contexts into a single unified string to facilitate cross-attention. Next, it derives the token-level representations of each mention after encoding this unified string. Finally, the joint representation is the concatenation of the context-enhanced token representations (v_{e_i}, v_{e_j}) along with their element-wise product, as illustrated below:

$$v_{(e_i, e_j)} = [v_{e_i}, v_{e_j}, v_{e_i} \odot v_{e_j}] \quad (1)$$

The resulting vector $v_{(e_i, e_j)}$ is then refined through a binary cross-entropy loss function using logistic regression that learns coreference. In our work, we use the learned weights of the CE_{LH}⁴. For the KNN cross-encoder (CE_{KNN}), we trained the weights of RoBERTa-Base using the KNN to generate focused mention pairs. We carry out our experiments in a transfer learning format where we train the crossencoders only on the training set of ECB+ and use the test sets of ECB+META. This is motivated by the work of Ortony et al. (1978), which argues the human processes required for comprehension of figurative and literal uses of language are essentially similar.

³Described in more detail in §C

⁴Provided by the authors

GPT-4 as Pairwise Classifier: Yang et al. (2022) demonstrated the viability of a prompt-based binary coreference classifier using GPT-2, though the results were sub-par. Building on their work, we employ a similar prompting technique with GPT-4 to develop an enhanced classifier. This classifier determines whether a pair of events, identified by marked triggers in sentences, are coreferent by responding with “Yes” or “No”. Similar to CE, we vary this method by incorporating the two filtering techniques (GPT_{LH}, GPT_{KNN})

4 Results

4.1 Metaphor Quality Control

To assess the quality of the generated metaphors, an annotator familiar with the events in the ECB+ dataset manually examines the Dev_{small} sets. We chose a familiarized annotator because metaphors often abstract away many of the details that make coreference obvious, and we are interested in whether or not the generated paraphrases would (by any stretch of the imagination) reasonably be interpreted as referring to the original event.

The annotator examines each of the original event mentions alongside their paraphrased versions and makes a binary judgment as to whether the two can be reasonably interpreted as referring to the same event. We estimate based on the results that approximately 99% of ECB+META₁ and 95% of ECB+META_m could be reasonably interpreted by a human as being coreferent to the original event mentions from which they are derived.

4.2 Coreference & Lexical Diversity

We use B³ (Bagga and Baldwin, 1998) and CoNLL (Denis and Baldridge, 2009; Pradhan et al., 2012) clustering metrics, in which we use the B_R³ for estimating recall, CoNLL as the overall metric (evaluated using CoVal (Moosavi et al., 2019)). For the methods that use LH as the filtering step, we follow Ahmed et al. (2023a)’s clustering with connected components. For KNN as the filtering step, we use Held et al. (2021)’s greedy agglomeration.

Filtering Scores: Following previous work, we first assess the B_R³ score on oracle results. This tests how well the filtering methods perform in minimizing false negatives (coreferent pairs that are eliminated inadvertently). From Table 1 we observe a substantial difference in the recall measures of ECB+ and ECB+META versions. The LH approach particularly takes a toll because

	Method	Dev	Dev _{small}	Test
ECB+	LH	76.3	87.9	81.5
	KNN	95.7	95.3	94.9
ECB+META ₁	LH	45.8	64.6	58.2
	KNN	91.8	93.7	91.4
ECB+META _m	LH	38.4	59.4	51.3
	KNN	84.4	86.5	85.6

Table 1: B_R³ Oracle Results on Dev, Dev_{small} and Test sets of ECB+, ECB+META₁, and ECB+META_m.

it relies on synonymous lemma pairs from the train set. Interestingly, KNN does well on the ECB+META versions, with only a minor drop in recall for ECB+META₁ and about 10% drop for ECB+META_m. Between ECB+META₁ and ECB+META_m, as expected, the recall drops more in ECB+META_m as more complex metaphors are used here.

Method	ECB+	ECB+META ₁	ECB+META _m
LH	74.1	49.8	54.0
CE _{LH}	78.1	60.9	50.6
CE _{KNN}	78	71.4	54.8
GPT _{LH}	78.23	62.5	55.6
GPT _{KNN}	67.73	60.15	55.5

Table 2: CoNLL F1 Baseline and Cross-encoder results on ECB+, ECB+META₁ and ECB+META_m Test sets.

CDEC Scores: We present the overall CoNLL F1 scores in Table 2 for the baseline (LH), the two fine-tuned cross-encoders (CE_{LH}, CE_{KNN}), and the methods that use GPT-4 (GPT_{LH}, GPT_{KNN}). From the table, it is evident that LH is no longer a strong baseline for ECB+META versions with a drop in 20% score. Both CE_{LH} and CE_{KNN} show a pattern of reducing score from ECB+META₁ to ECB+META_m, with CE_{LH} performing considerably worse. Interestingly, the drop in scores for CE_{KNN} is not substantial for ECB+META₁ but there is a dramatic drop of 20% for ECB+META_m. GPT_{LH} achieves the highest scores on ECB+ and ECB+META_m, demonstrating that GPT-4’s performance aligns with the state-of-the-art, unlike its predecessor GPT-2. However, the financial implications of using GPT_{LH} and GPT_{KNN} are noteworthy; running CDEC with these methods incurred approximately \$75 in API costs to OpenAI.

From these results, we can conclude three things: a) ECB+ is an easy dataset, b) datasets with complex metaphors are harder benchmarks, and c) GPT-4 is only as good as the CE methods with a significant amount of added costs.

Lexical Diversity: We estimate the lexical diversity (MLTD; McCarthy and Jarvis (2010)) of the mention triggers of event clusters. We first eliminate singleton clusters. Then we calculate a weighted average (by cluster size) of the MLTD score for each cluster. The scores we achieved for the test sets of each version of ECB+ are as follows: ECB+: 7.33. ECB+META1: 11.92, ECB+META_m: 26.48. From the lower CDEC scores from Table 2 and the increasing diversity scores of the more complex corpus, we can establish a negative correlation between CDEC scores and MLTD.

Overall, the results confirm our hypothesis that when a dataset a) moves away from strong lexical overlap and b) has figurative language usage, the CDEC scores drop.

5 Analysis

5.1 Coreference Resolution Difficulty

We evaluate whether the paraphrased versions are more difficult for humans to determine as coreferent. On the Dev_{small} splits of ECB+META₁, ECB+META_m, and ECB+, a human annotator reaches the same coreference verdict regardless of the degree of figurative language approximately 98% of the time. Cases in which the human annotator did not reach the same verdict generally involved convergent metaphorical language, for example:

Event a: The Indian navy unfurled the words that it had ensnared 23 pirates *in the law’s net* who cast ominous shadows over a merchant vessel in the Gulf of Aden on Saturday, the latest in a series of recent violent ballets with Somali pirates.

Event b: Indian Naval Ship *throws a net over* three pirate vessels in a single orchestrated symphony .

were incorrectly identified as coreferent; in actuality the former refers to the arrest of the pirates but the latter refers to the interception of their ships. This analysis supports the findings of Ortony et al. (1978): that, for humans, figurative language use and literal language do not substantially affect comprehension.

5.2 Qualitative Error Analysis

We examined the coreference predictions of CE_{KNN} on 142 common mention pairs between ECB+, ECB+META₁, and ECB+META_m, as CE_{KNN} achieved the best overall performance. For mention pairs that CE_{KNN} correctly predicted as coreferent across all versions, we noticed a pattern: the same event trigger was shared in each (see Figure 4).

In cases where CE_{KNN} got the prediction right on ECB+ but wrong on the META versions, the event triggers in ECB+ were changed to different ones in the META versions (see Figure 5). When CE_{KNN} incorrectly predicted coreference on ECB+ but correctly predicted it in the META versions, it was because the same triggers in ECB+ were altered to different ones (see Figure 6). This further affirms that the model heavily relies on surface triggers for making coreference decisions.

6 Future Work

Future research could explore applying more recent CDEC techniques on ECB+META. These techniques could include symbolic grounding, as discussed in Ahmed et al. (2024b,a), and event type categorical cross-encoding, as proposed by Otmazgin et al. (2023). Another outcome of this research is to use CDEC as a text complexity metric (Hale, 2016) of a corpus. We argue that a corpus is more complex if a CDEC algorithm is not able to identify that different explanations of the same event are the same. An interesting line of future work would be to automatically generate an optimally complex CDEC corpus, i.e., a corpus that yields the lowest coreference score.

In this work, we rely on the GPT-4’s metaphor list and substitution choice. The only control we have is to make a coherent choice, however, we find ourselves subjected to the unpredictable outputs, colloquially referred to as “hallucinations”, generated by GPT-4. In the future, we aim to integrate human feedback into the process of metaphor selection and to employ annotated metaphor databases from studies such as Joseph et al. (2023).

7 Conclusion

In this paper, we introduced ECB+META a lexically rich variant of ECB+ using constrained metaphoric paraphrasing of the original corpus. We provide hand-corrected event trigger annotations of two versions of ECB+META differing in the kind of metaphoric transformation using either single words or phrases. We finally provide baseline results using existing SoTA methods on this dataset and show their limitations when there is substantial lexical diversity in the corpus. Through the provided data and methodology, we lay a path forward for future research in Cross-Document Event Coreference Resolution on more challenging datasets.

Limitations

The study faced several limitations, including its focus on a single language-English. Some experiments were conducted within a small sample space, especially for Dev_{small} , potentially leading to biased results and limiting the generalizability of the findings. Finally, while the study utilized variations within a single dataset, the reliance on this sole dataset could introduce inherent biases, affecting the broader applicability of the research outcomes.

Reproducibility Concern: All the coreferencing experiments are reproducible, but the generation of ECB+META is not. So we may have vastly different results if a new version of ECB+META is created with the methodology. However, we released all the generated text that came out of our work and the code to run the experiments.

LLMs on ECB+. Contamination Concern The GPT-4 has likely been contaminated by the test sets of ECB+, i.e., GPT-4 has been pretrained on this benchmark. With the recent work involving GPT and ECB+ (Yang et al., 2022; Ravi et al., 2023a,b), it seems likely the test set is also been used in the instruction fine-tuning of GPT-4. But we stress the synthesizing of datasets to battle contamination as we do in our work.

Ethics Statement

AI-generated text should always be thoroughly scrutinized before being used for any application. In our work, we provide methods to synthesize new versions of the same real articles. This can have unintentional usage in the propagation of disinformation. This work is only intended to be applied to research in broadening the field of event comprehension. Our work carries with it the inherent biases in news articles of ECB+ corpus and has the potential of exaggerating it with the use of GPT-4, which in itself has its own set of risks and biases.

Acknowledgements

We thank the anonymous reviewers for their helpful suggestions that improved our paper. We are also grateful to Susan Brown, Alexis Palmer, and Martha Palmer from the BoulderNLP group for their valuable feedback before submission. Thanks also to William Held and Vilém Zouhar for their insightful comments. We gratefully acknowledge the support of DARPA FA8750-18-2-0016-AIDA – RAMFIS: Representations of vectors and Abstract Meanings for Information Synthesis and a

sub-award from RPI on DARPA KAIROS Program No. FA8750-19-2-1004. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA or the U.S. government.

References

- Shafiuddin Rehan Ahmed, George Arthur Baker, Evi Judge, Michael Reagan, Kristin Wright-Bettner, Martha Palmer, and James H. Martin. 2024a. [Linear cross-document event coreference resolution with X-AMR](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 10517–10529, Torino, Italia. ELRA and ICCL.
- Shafiuddin Rehan Ahmed, Jon Cai, Martha Palmer, and James H. Martin. 2024b. [X-AMR annotation tool](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 177–186, St. Julians, Malta. Association for Computational Linguistics.
- Shafiuddin Rehan Ahmed, Abhijnan Nath, James H. Martin, and Nikhil Krishnaswamy. 2023a. [2 * n is better than n²: Decomposing event coreference resolution into two tractable problems](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1569–1583, Toronto, Canada. Association for Computational Linguistics.
- Shafiuddin Rehan Ahmed, Abhijnan Nath, Michael Reagan, Adam Pollins, Nikhil Krishnaswamy, and James H. Martin. 2023b. [How good is the model in model-in-the-loop event coreference resolution annotation?](#) In *Proceedings of the 17th Linguistic Annotation Workshop (LAW-XVII)*, pages 136–145, Toronto, Canada. Association for Computational Linguistics.
- Emily Allaway, Shuai Wang, and Miguel Ballesteros. 2021. [Sequential cross-document coreference resolution](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4659–4671, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Amit Bagga and Breck Baldwin. 1998. Algorithms for scoring coreference chains. In *The first international conference on language resources and evaluation workshop on linguistics coreference*, volume 1, pages 563–566. Citeseer.
- 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.

- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv:2004.05150*.
- Michael Bugert and Iryna Gurevych. 2021. [Event coreference data \(almost\) for free: Mining hyperlinks from online news](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 471–491, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Avi Caciularu, Arman Cohan, Iz Beltagy, Matthew Peters, Arie Cattan, and Ido Dagan. 2021. [CDLM: Cross-document language modeling](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2648–2662, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Arie Cattan, Alon Eirew, Gabriel Stanovsky, Mandar Joshi, and Ido Dagan. 2021. [Cross-document coreference resolution over predicted mentions](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 5100–5107, Online. Association for Computational Linguistics.
- Tuhin Chakrabarty, Arkadiy Saakyan, Debanjan Ghosh, and Smaranda Muresan. 2022. Flute: Figurative language understanding through textual explanations. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7139–7159.
- Tuhin Chakrabarty, Arkadiy Saakyan, Olivia Winn, Artemis Panagopoulou, Yue Yang, Marianna Apidianaki, and Smaranda Muresan. 2023. [I spy a metaphor: Large language models and diffusion models co-create visual metaphors](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 7370–7388, Toronto, Canada. Association for Computational Linguistics.
- Tuhin Chakrabarty, Xurui Zhang, Smaranda Muresan, and Nanyun Peng. 2021a. Mermaid: Metaphor generation with symbolism and discriminative decoding. *arXiv preprint arXiv:2103.06779*.
- Tuhin Chakrabarty, Xurui Zhang, Smaranda Muresan, and Nanyun Peng. 2021b. [MERMAID: Metaphor generation with symbolism and discriminative decoding](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4250–4261, Online. Association for Computational Linguistics.
- 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 Language Resources Association (ELRA).
- Agata Cybulska and Piek Vossen. 2015. [Translating granularity of event slots into features for event coreference resolution](#). In *Proceedings of the The 3rd Workshop on EVENTS: Definition, Detection, Coreference, and Representation*, pages 1–10, Denver, Colorado. Association for Computational Linguistics.
- Pascal Denis and Jason Baldridge. 2009. Global joint models for coreference resolution and named entity classification. *Procesamiento del lenguaje natural*, 42.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alon Eirew, Arie Cattan, and Ido Dagan. 2021. [WEC: Deriving a large-scale cross-document event coreference dataset from Wikipedia](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2498–2510, Online. Association for Computational Linguistics.
- Raymond W. Gibbs. 1994. *The Poetics of Mind: Figurative Thought, Language, and Understanding*. Cambridge University Press.
- John Hale. 2016. [Information-theoretical complexity metrics](#). *Language and Linguistics Compass*, 10(9):397–412.
- William Held, Dan Iter, and Dan Jurafsky. 2021. [Focus on what matters: Applying discourse coherence theory to cross document coreference](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1406–1417, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.
- Rohan Joseph, Timothy Liu, Aik Beng Ng, Simon See, and Sunny Rai. 2023. [NewsMet : A ‘do it all’ dataset of contemporary metaphors in news headlines](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10090–10104, Toronto, Canada. 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.

- Jan-Christoph Klie, Michael Bugert, Beto Boulosa, Richard Eckart de Castilho, and Iryna Gurevych. 2018. [The INCEpTION platform: Machine-assisted and knowledge-oriented interactive annotation](#). In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 5–9, Santa Fe, New Mexico. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2023. [Large language models are zero-shot reasoners](#).
- George Lakoff and Mark Johnson. 1980. *Metaphors We Live By*. University of Chicago Press.
- Yucheng Li, Shun Wang, Chenghua Lin, Frank Guerin, and Loic Barrault. 2023. [FrameBERT: Conceptual metaphor detection with frame embedding learning](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1558–1563, Dubrovnik, Croatia. Association for Computational Linguistics.
- 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](#).
- Philip M McCarthy and Scott Jarvis. 2010. Mtd, vocd-d, and hd-d: A validation study of sophisticated approaches to lexical diversity assessment. *Behavior research methods*, 42(2):381–392.
- Yehudit Meged, Avi Caciularu, Vered Shwartz, and Ido Dagan. 2020. [Paraphrasing vs coreferring: Two sides of the same coin](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4897–4907, Online. Association for Computational Linguistics.
- Nafise Sadat Moosavi, Leo Born, Massimo Poesio, and Michael Strube. 2019. [Using automatically extracted minimum spans to disentangle coreference evaluation from boundary detection](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4168–4178, Florence, Italy. Association for Computational Linguistics.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Andrew Ortony, Diane L. Schallert, Ralph E. Reynolds, and Stephen J. Antos. 1978. [Interpreting metaphors and idioms: Some effects of context on comprehension](#). *Journal of Verbal Learning and Verbal Behavior*, 17(4):465–477.
- Shon Otmazgin, Arie Cattan, and Yoav Goldberg. 2023. [LingMess: Linguistically informed multi expert scorers for coreference resolution](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2752–2760, Dubrovnik, Croatia. Association for Computational Linguistics.
- Barbara C Palmer and Mary Alice Brooks. 2004. Reading until the cows come home: Figurative language and reading comprehension. *Journal of Adolescent & Adult Literacy*, 47(5):370–379.
- Barbara C Palmer, Vikki S Shackelford, Sharmane C Miller, and Judith T Leclere. 2006. Bridging two worlds: Reading comprehension, figurative language instruction, and the english-language learner. *Journal of Adolescent & Adult Literacy*, 50(4):258–267.
- 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.
- James Ravenscroft, Amanda Clare, Arie Cattan, Ido Dagan, and Maria Liakata. 2021. [CD²CR: Coreference resolution across documents and domains](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 270–280, Online. Association for Computational Linguistics.
- Sahithya Ravi, Raymond Ng, and Vered Shwartz. 2023a. [Comet-m: Reasoning about multiple events in complex sentences](#). *ArXiv*, abs/2305.14617.
- Sahithya Ravi, Chris Tanner, Raymond Ng, and Vered Shwartz. 2023b. [What happens before and after: Multi-event commonsense in event coreference resolution](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1708–1724, Dubrovnik, Croatia. Association for Computational Linguistics.
- Anatol Stefanowitsch. 2006. Words and their metaphors: A corpus-based approach. *Trends in Linguistics Studies and Monographs*, 171:63.
- Kevin Stowe, Nils Beck, and Iryna Gurevych. 2021a. [Exploring metaphoric paraphrase generation](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 323–336, Online. Association for Computational Linguistics.
- Kevin Stowe, Tuhin Chakrabarty, Nanyun Peng, Smaranda Muresan, and Iryna Gurevych. 2021b. [Metaphor generation with conceptual mappings](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6724–6736, Online. Association for Computational Linguistics.

Kevin Stowe, Leonardo Ribeiro, and Iryna Gurevych. 2020. Metaphoric paraphrase generation. *arXiv preprint arXiv:2002.12854*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. *Attention is all you need*. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

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).

Lennart Wachowiak and Dagmar Gromann. 2023. Does gpt-3 grasp metaphors? identifying metaphor mappings with generative language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1018–1032.

Ellen Winner. 1988. *The Point of Words: Children's Understanding of Metaphor and Irony*. Harvard University Press.

Xiaohan Yang, Eduardo Peynetti, Vasco Meerman, and Chris Tanner. 2022. What gpt knows about who is who. *arXiv preprint arXiv:2205.07407*.

Xiaodong Yu, Wenpeng Yin, and Dan Roth. 2022. *Pairwise representation learning for event coreference*. In *Proceedings of the 11th Joint Conference on Lexical and Computational Semantics*, pages 69–78, Seattle, Washington. Association for Computational Linguistics.

Yutao Zeng, Xiaolong Jin, Saiping Guan, Jiafeng Guo, and Xueqi Cheng. 2020. *Event coreference resolution with their paraphrases and argument-aware embeddings*. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3084–3094, Barcelona, Spain (Online). International Committee on Computational Linguistics.

A ECB+ Corpus

	Train	Dev*	Dev _{small} *	Test
Topics	25	8	8	10
Documents	594	156	40	206
Mentions	3808	968	277	1780

Table 3: Corpus statistics for event mentions in ECB+

The ECB+ corpus (Cybulska and Vossen, 2014) is a popular English corpus used to train and evaluate systems for event coreference resolution. It extends the Event Coref Bank corpus (ECB; Bejan and Harabagiu (2010)), with annotations from

Metaphoric Paraphrasing

You are a metaphor expert. Your task is to transform specific words in a given sentence into metaphors. These metaphors can only be **single-word/multi-word** replacements. Here are the detailed steps you need to follow:

Read the Sentence Provided: Focus on understanding the context and meaning of the sentence.

Review the Word List: This list contains the words you need to transform into metaphors.

Generate Metaphors: Create **5 distinct** single-word/multi-word metaphors for each word in the list.

Compose a New Sentence: Replace the original words with your chosen metaphors randomly. Ensure the new sentence maintains logical and grammatical coherence.

Sentence to Transform:

""{{sentence}}""

Word List to Convert into Metaphors:

""{{trigger_list}}""

Output Requirements: Provide your final output in JSON format, including:

- The "Original Sentence".
- The "Original Word List".
- The "Metaphoric Word List" (with your chosen metaphors).
- The "Metaphoric Sentence" (the sentence with metaphors incorporated).

Remember, the goal is to use metaphors to convey the original sentence's meaning in a more nuanced or impactful way without altering the core information.

Figure 2: Metaphoric Paraphrasing Prompt following Chain of Thought Reasoning. We provide the steps in this prompt to follow.

around 500 additional documents. The corpus includes annotations of text spans that represent events, as well as information about how those events are related through coreference. We divide the documents from topics 1 to 35 into the training and validation sets⁵, and those from 36 to 45 into the test set, following the approach of Cybulska and Vossen (2015). We further break the documents of the validation set into two subsets: Dev and Dev_{small} for our error analysis. Full corpus statistics can be found in Table 3.

⁵Validation set includes documents from the topics 2, 5, 12, 18, 21, 34, and 35

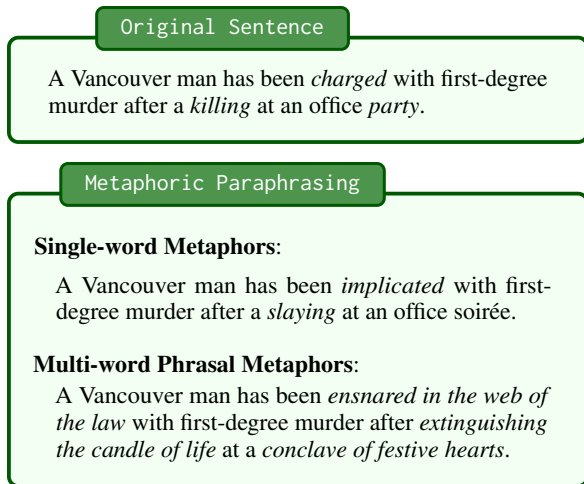


Figure 3: Metaphoric Paraphrasing: Transforming a Sentence with Figurative Language. Event triggers, indicated in italics, undergo modification in paraphrased versions, annotated by GPT-4 with two variations.

B Metaphoric Paraphrase Prompt

We present the prompt used with GPT-4 in Figure 2 for generating the Metaphoric Paraphrasing of ECB+ documents. We use two separate prompts for generating single-word metaphors and multi-word metaphors. We ran this prompt on the validation and test sets of ECB+ using GPT-4 as the LLM and a temperature value of 0.7. We force GPT-4 to produce JSON-style output to avoid parsing issues. It costs about \$16 to generate ECB+META₁ and \$18 to generate ECB+META_m with GPT-4 API calls. In the future, we plan to provide this conversion of the training set of ECB+ as well.

C Experiment Setup

LH details: we set the sentence-level word overlap ratio threshold at 0.005. We employ spaCy 3.7.4 as the lemmatizer to extract the root forms of words.

KNN details: we adopt the RoBERTa-Base model, enhanced with a triplet loss function calculated by `F.triplet_margin_loss` with a 10 margin, L2 norm ($p = 2$), and $\epsilon = 1e - 6$ for stability, without swapping and mean reduction. Our optimization uses AdamW, targeting bi-encoder parameters with a 1×10^{-5} learning rate across 20 iterations and batches of 4.

CE_{LH} details: We utilize the RoBERTa-Base model with the AdamW optimizer. Learning rates are set to 1×10^{-5} for BERT class parameters and 1×10^{-4} for the classifier. The model is trained over 20 epochs, using the sentences in which the

Split	Method	B _R ³	B _P ³	B _{F1} ³	CoNLL
Dev	LH	51.8	64.5	57.4	56.3
	CE _{LH}	47.2	77.3	58.6	55.3
	CE _{KNN}	42.4	86.2	56.8	49.2
Dev _{small}	LH	68.4	78.3	73.1	62.0
	CE _{LH}	64.8	84.7	73.4	59.0
	CE _{KNN}	62.4	91.6	74.2	55.5

Table 4: Baseline and Cross-encoder results on ECB+META_m Dev and Dev_{small} sets.

two mentions occur as context, and mention pairs generated by LH.

CE_{KNN} details: It mirrors the CE_{LH} configuration but it is trained on mention pairs from KNN exclusively.

All Non-GPT experiments are conducted on a single NVIDIA RTX 3090 with 24GB of VRAM. For generating the META datasets, we utilized GPT-4 (model version: gpt4-0613), setting the temperature parameter to 0.7.

D ECB+META_m Complete Results

We provide the baseline results for validation sets of ECB+META_m. As shown in Table 4, the results are consistent even for the development sets, where we see significantly low coreference scores with the used methods. Interestingly, LH performs better than the cross-encoder methods on these splits.

E Error Analysis

For more examples, please checkout the provided excel file in data repository.

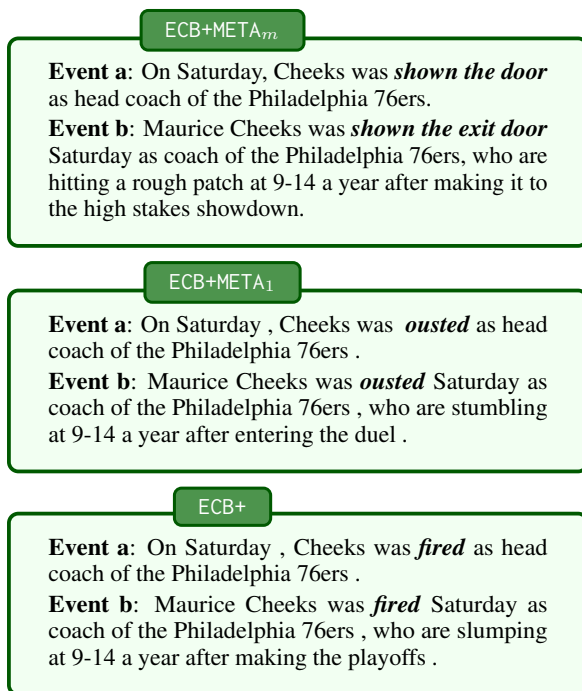


Figure 4: Correct prediction of coreferent mention pair across all datasets with CE_{KNN} . Pairs have the same event trigger in each case.

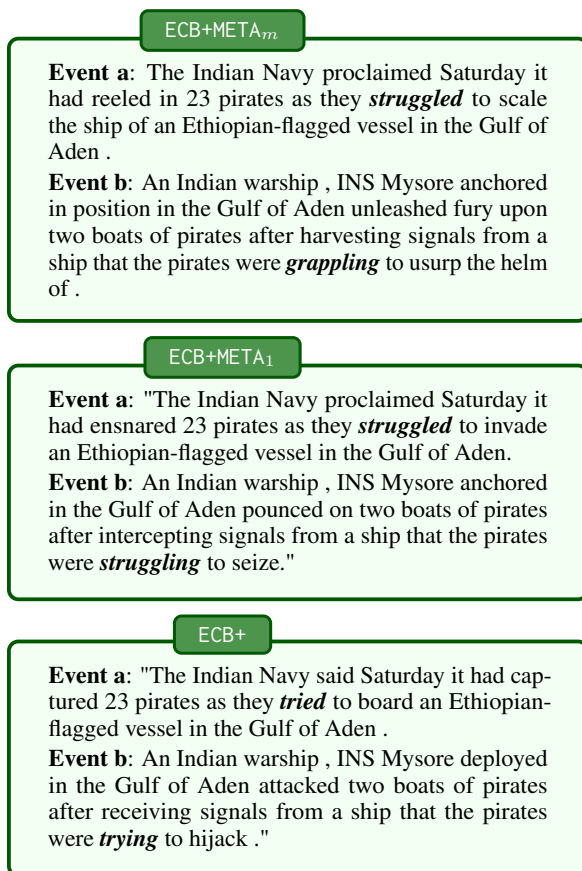


Figure 5: Correct coreference prediction in ECB+ but not in the META versions, simply because the triggers got changed.

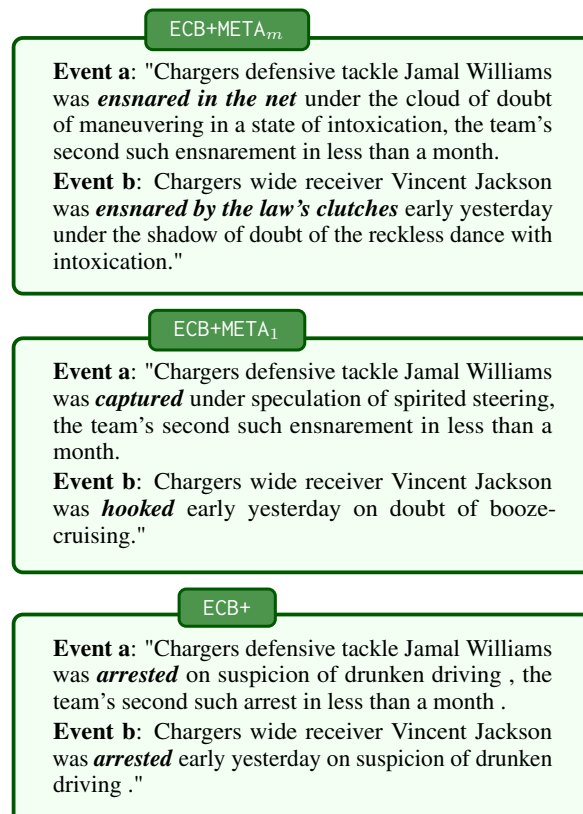


Figure 6: Correct non-coreference prediction in ECB+META but not in ECB+, simply because the META versions' event triggers were changed.