

IDIAPers @ Causal News Corpus 2022: Extracting Cause-Effect-Signal Triplets via Pre-trained Autoregressive Language Model

Martin Fajcik^{*, 1, 2}, Muskaan Singh¹, Juan Zuluaga-Gomez^{1, 3}, Esaú Villatoro-Tello^{1, 4}, Sergio Burdisso^{1, 5}, Petr Motlicek^{1, 2}, Pavel Smrz²

¹Idiap Research Institute, Martigny, Switzerland

²Brno University of Technology, Brno, Czech Republic

³Ecole Polytechnique Fédérale de Lausanne, Switzerland

⁴Universidad Autónoma Metropolitana Unidad Cuajimalpa, Mexico City, Mexico

⁵Universidad Nacional de San Luis (UNSL), San Luis, Argentina

*corresponding author: martin.fajcik@vut.cz

Abstract

In this paper, we describe our shared task submissions for Subtask 2 in CASE-2022, Event Causality Identification with Casual News Corpus. The challenge focused on the automatic detection of all cause-effect-signal spans present in the sentence from news-media. We detect cause-effect-signal spans in a sentence using T5 — a pre-trained autoregressive language model. We iteratively identify all cause-effect-signal span triplets, always conditioning the prediction of the next triplet on the previously predicted ones. To predict the triplet itself, we consider different causal relationships such as *cause*→*effect*→*signal*. Each triplet component is generated via a language model conditioned on the sentence, the previous parts of the current triplet, and previously predicted triplets. Despite training on an extremely small dataset of 160 samples, our approach achieved competitive performance, being placed second in the competition. Furthermore, we show that assuming either *cause*→*effect* or *effect*→*cause* order achieves similar results.¹

1 Introduction

Causality links the relationship between two arguments — cause and effect (Barik et al., 2016). Figure 1 shows examples extracted from the Causal News Corpus (CNC) (Tan et al., 2022b). *Cause* clauses appear in yellow, *Effect* in green, and *Signals* in pink; hereafter referred to as CES triplets. As shown in the example, “the bombing created panic among villagers”, illustrates that the event “bombing” caused the event “panic among villagers” termed as *effect*. The linkage among the cause and effect, i.e., the word “created”, is termed as *signal* and can be expressed explicitly or implicitly.

¹Code at <https://github.com/idiap/cncsharedtask>.

(A) Casual segment:	
The treating doctor said	Sangram lost around 5kgs due to
the hunger strike	.
The bombing	created panic among villagers
Dissatisfied with the package	workers staged an all-night sit-in
(B) Non Casual Segment: Thus we too joined the sloganerrring.	
The alliance claimed 4,000 took part last year.	

Figure 1: Examples from the Causal News Corpus, causes are in yellow, effects in green, and signals in pink. If a sentence has both — cause and effect — it is referred to as casual (A), otherwise, as non-casual (B).

Automatically detecting and extracting causality relations plays a vital role in many natural language processing (NLP) works to tackle inference and understanding (Dunietz et al., 2020; Fajcik et al., 2020; Jo et al., 2021; Feder et al., 2021a). It has applications in various down-streaming NLP tasks, namely, causal question-answering generation, explaining social media behavior, political phenomena, effective education, and gender bias in the research community (Tan et al., 2014; Wood-Doughty et al., 2018; Sridhar and Getoor, 2019; Veitch et al., 2020; Zhang et al., 2020; Feder et al., 2021b).

In this paper, we describe our methodology for CASE-2022 cause-effect-signal span detection shared task (Subtask 2). Overall, our main contributions are listed below:

1. We show that cause-effect-signal spans can be extracted by a simple pre-trained generative seq2seq model trained on just 160 instances.
2. We develop a method for extracting all causal triplets from the sentence in an iterative manner.
3. We investigate how language models deal with

the causal order of the cause and effect spans to answer the research question “*should cause be identified first, and only then effect, or vice-versa?*”.

4. We show that an efficient F1 best-substring matching algorithm, known for question answering, can be applied to deal with rare cases when a language model (LM) does not generate part of the input sequence.

2 Related Work

The problem of causality extraction from text is a challenging task as it requires semantic understanding and contextual knowledge. There were many attempts in the domain of linguistics for corpora creation for event extraction but with limited size such as CausalTimeBank (CTB) (Mirza et al., 2014) from news with 318 pairs, CaTeRS (Mostafazadeh et al., 2016) from short stories with 488 casual links, EventStoryLine (Caselli and Vossen, 2017) from online news articles with 1,770 casual event pairs, semantic relation corpora PDTB-3 (Webber et al., 2019) with over 7,000 causal relations and CNC corpus (Tan et al., 2022b,c) with 1,957 casual events with multiple event pairs. Compared to previous datasets, CNC differs by focusing on event sentences, accepting arguments which does not need to form a clause, and not limiting itself to pre-defined list of connectives, but instead including causal examples in more varied linguistic constructions. The previous work in this domain can be broadly classified into knowledge-based approaches, statistical ML, and deep-learning-based approaches. The knowledge-based approach uses linguistic patterns by predefining hand-crafted or keywords (Garcia et al., 1997; Khoo et al., 2000; Radinsky et al., 2012; Beamer et al., 2008; Girju et al., 2009; Ittoo and Bouma, 2013; Kang et al., 2014; Khoo et al., 1998; Bui et al., 2010).

Statistical techniques (Girju, 2003; Do et al., 2011) rely on building probabilistic models over features extracted via third-party NLP tools such as Wordnet (Miller, 1994). Deep-learning techniques map words and features into low-dimensional dense vectors, which may alleviate the feature sparsity problem. The most frequent used sequence to sequence models are feed-forward network (Ponti and Korhonen, 2017), long short-term memory networks (Kruengkrai et al., 2017; Dasgupta et al., 2018; Martínez-Cámara et al., 2017) convolutional neural networks (Jin et al., 2020; Kruengkrai et al.,

2017; Wang et al., 2016), recurrent neural networks (Yao et al., 2019), gated recurrent units (Chen et al., 2016) which embed semantic and syntactic information in local consecutive word sequences (Yao et al., 2019). Later unsupervised training model such as BERT (Devlin et al., 2018; Sun et al., 2019), RoBERTa (Becquin, 2020), graph convolution network (Zhang et al., 2018), graph attention networks and joint model for entity relation extraction (Li et al., 2017; Wang and Lu, 2020; Zhao et al., 2021; Bekoulis et al., 2018).

In this work, we base our model on T5 (Raffel et al., 2020), a sequence-to-sequence transformer model, pre-trained on a mixture of denoising objective and 25 supervised tasks such as machine translation, linguistic acceptability, abstractive summarization or question answering. The unsupervised denoising objective randomly replaces spans of the input with different mask tokens, and generates contents of these masked spans prefixed with these special mask tokens. Furthermore, our work shares similarities with pointer-network (Vinyals et al., 2015) based generative framework for various NER subtasks introduced by Yan et al. (2021). Contrastively, our work is more adapted to low-resource scenarios, as no extra parameters were added to our system, at the cost of errors, which can happen in the postprocessing matching step.

3 Problem Description

CASE-2022 shared task challenge (Tan et al., 2022a) aimed for event causality identification, and extraction in casual news corpus (Tan et al., 2022b). It comprised of two subtasks, namely casual event classification (Subtask 1) and cause-effect-signal span detection (Subtask 2)². Subtask 2 aims on extracting the spans corresponding to cause-effect-signal (CES) triplets, as shown in Figure 1. We trained a generative seq2seq model to address this challenge and extracted the CES triplets using an iterative procedure (see Section 4.1).

The dataset statistics are presented in Table 1. The number of total sentences is given by the column *#Sentences*, whereas a total number of CES triplets is in column *#Relations*. Column *#Signals* shows how many signal annotations were present in the total number of CES triplets.

²We participated in both subtasks, but report on Subtask 2 in this paper. For Subtask 1, we refer reader to our standalone publication (Burdizzo et al., 2022).

Split	#Sentences	#Relations	#Signals
Train	160	183	118 (64%)
Dev	15	18	10 (56%)
Test	89	119	98 (82%)

Table 1: Dataset statistics. See text for details.

4 Methodology

4.1 Language Model Training

We utilize T5 (Raffel et al., 2020), a pre-trained autoregressive transformer-based language model trained on a mixture of unsupervised and supervised tasks that require language understanding. The model is conditioned $n \times 3$ times for each example, as there can be n CES triplets in one sentence (up to $n = 4$ triplets in training data). Each time, we condition the language model 3 times for every example and its corresponding CES triplet, generating a different triplet component (cause, effect, and signal) to learn to generate the entire CES triplet. As these triplets are unordered, we uniformly sample a random path among them (e.g., 2-3-1-4, for sample with four triplets) during training. We only train with as many triplets, as available in the training data. We now describe the input format, further illustrated in Appendix B.

Firstly, the model’s encoder is conditioned with sentence tokens `<sentence>` followed by the history of already generated CES triplets for this example (empty if there was none) as

```
<sentence> _history : <history>.
```

The history is always prepended with `_history:` tokens. The content of the history are the already generated triplets. Each part of the triplet is prepended with its corresponding `_cause:`, or `_effect:`, or `_signal:` sequence. Concurrently, model’s decoder is prefixed with `_cause:` sequence. In this case, the probability of cause sequence is maximized.

Secondly, the model is conditioned with sentence tokens `<sentence>` and cause tokens `<cause>`, prepended with `_cause:` token as

```
<sentence> _cause : <cause>
_history : <history>.
```

This time, the decoder is prompted with `_effect:` prefix, and the probability of effect sequence is maximized.

Thirdly, the model is conditioned with sentence tokens `<sentence>`, cause tokens `<cause>`, and effect tokens `<effect>` with `_effect:` token prepended as

```
<sentence> _cause : <cause> _effect
: <effect> _history : <history>.
```

Analogically, decoder is prompted with `_signal:` prefix and probability of signal sequence is maximized. As the signal might not always be part of the CES triplet, we let the model generate `_empty` token in these cases.

4.2 Experimental Details

We use cross-entropy (CE) loss to train the T5. We firstly average CE loss over tokens, then over inputs per example (for all CES triplets), and then across mini-batch. We use greedy search to generate the sequences. In inference time, we always generate 4 CES triplets for each sentence, as that is the maximum we observed in the training data.

As we don’t constrain the decoding, the generated sequence does not have to match certain sub-string in the input. However, the extractive task requires inserting tags around a cause, effect, or signal span inside the input sentence. Therefore we map the generated sequences back to the input sentence via F1 matching. In particular, for each generated sequence, we find the most similar substring in the input, where the similarity is measured via token-level F1 score. We utilize an efficient F1 matching technique, which prunes out a significant part of the search space, presented in the Appendix C.1 of Fajcik et al. (2021)³. We base our implementation on PyTorch (Paszke et al., 2019), Transformers (Wolf et al., 2020) libraries and use AdamW (Loshchilov and Hutter, 2017) for optimization. We tune hyperparameters via HyperOpt (Bergstra et al., 2015) and report the exact hyperparameters in Appendix A.

4.3 Evaluation Metrics

In this section, we describe the metrics we used to evaluate the system.

F1: F1 score was the official main evaluation metric in the challenge. It is computed over B, and I tags in sequence following the BIO tagging scheme for every example and every CES triplet component separately, using `seqeval`⁴. The F1 is then

³Implemented at <https://shorturl.at/kxEVW>.

⁴<https://github.com/chakki-works/seqeval>.

System	CE	Cause	Effect	Signal	Overall
Baseline	-	-	-	-	2.2
T5-NoHistory	.181	-	-	-	67.7±2
T5-ECS	.168	75.9±5	71.3±4	76.1±5	73.5±2
T5-CES	.183	81.0±4	67.8±2	66.7±5	73.0±2
T5-CES _{LARGE}	.159	73.5±8	74.1±4	77.2±7	74.8±2

Table 2: Main results, in terms of Cross-Entropy (CE) and F1, with \pm standard deviations on dev data.

averaged firstly across dataset examples, obtaining F1 for each component (*Cause F1*, *Effect F1*, *Signal F1*). *Overall F1* is computed as a weighted average of component examples by their frequency.

CE: is an average token cross-entropy, computed as described in Section 4.2.

ES Acc: is an empty-signal accuracy, i.e., an accuracy of the model predicting no signal span in the CES triplet when given golden cause and effect.

4.4 Baseline Model

As a baseline model, we used the CASE-2022 organizers’ provided model for Subtask 2: a random generator that uniformly samples a cause, effect, and signal spans⁵ from the sentence. This baseline guarantees the cause and the effect do not overlap.

5 Results & Discussion

We now report the results obtained from averaging at least ten measured performances from 10 checkpoints trained with different seeds⁶. We studied 4 different variants of our system. System T5-CES is our vanilla model described in 4.1, based on T5-base. System T5-CES_{LARGE} is the same model based on T5-large. Unlike T5-CES, system T5-ECS reverses the generation order by generating the first effect and cause, followed by the signal (assuming causal order *effect*→*cause*→*signal*, hence the suffix ECS). Lastly, we studied the effect of conditioning the model on the history of already generated triplets. We remove the history from the input at all times in training and predict the four identical CES triplets for each example in test time. Our ablated results are available in Table 2.

Firstly, the model with no history at input performs significantly worse, validating our hypothesis that the model can learn to decrease the probability of the triplets already contained within the

⁵Available at <https://shorturl.at/msY04>.

⁶Dev set predictions from our best t5-base model are available at <https://shorturl.at/bjVZ9>.

System	Dev F1	Dev ₁ F1	Dev ₂ F1	Dev ES Acc	Test F1
T5-ECS	77.7	80.9	71.1	82	43.4
T5-CES _{LARGE}	78.3	77.4	80.0	70	43.7
T5-CES	77.5	79.6	73.3	70	48.8

Table 3: Top checkpoints submitted to the leaderboard.

input, even from just 160 samples. Secondly, we observed a general trend that *in the Cause F1 T5-CES outperforms T5-ECS* and *in Effect F1, T5-ECS outperforms T5-CES*. This leads to the hypothesis that whichever part of the triplet, cause or effect, is generated first, the language model performs better in its case. Thirdly, we observed that the large model achieved the best results on average. It also achieved our best single-checkpoint performance on the dev set (78.3 Overall F1). However, given the sample size of the dev set, the differences between T5-CES, T5-ECS, and T5-CES_{LARGE} can hardly be deemed significant.

Next we present our results on the test set in Table 3. We submitted checkpoints with the best overall F1 score on the dev set (Dev F1) to the leaderboard while varying the model types. We observed a significant drop in performance on the test data. As the annotation on the test data is not released at the time of writing, the causes of this performance drop remain unknown. We hypothesize it could have been caused by a covariate shift in the test data, as supported by #Signals statistics in Table 1.

Additionally, we include extra statistics (Dev₀ F1, Dev₁ F1, Dev ES Acc) for our best checkpoints. We expected the performance on the dev subset with two triplets (Dev₂ F1) per example to be worse than on the dev subset with one triplet per sentence (Dev₁ F1). Performance-wise this does not always seem to be the case. Upon manual analysis, we found that the model often failed in the second round of triplet extraction. We found 2 LM hallucinations out of 18 dev samples in the second generation round.

6 Inference Speed

Measuring the inference speed on test set, we used Intel i5-based 2080Ti GPU workstation. The inference of 4 CES triplets without postprocessing per 1 sentence example took 1.46 seconds on average. The postprocessing runtime was negligible, taking 0.025 seconds per sentence example on average.

7 Conclusion

In this work, we have analyzed our CASE-2022 2nd place submissions on Subtask 2. We showed that a generative model could extract cause-effect-signal triplets at the competitive level using just 160 annotated samples. We investigated causal assumptions about the generation order of cause and effect to answer the research question “*should cause be identified first, and only then effect, or vice-versa?*” and found that while the Overall F1 won’t change significantly, whichever component was generated first achieved better performance on average (Cause first achieved better Cause-F1, and Effect first Effect-F1 respectively). Finally, we showed the F1 difference between the dev subset with 1 or 2 causal triplets per sentence is negligible.

Acknowledgements

This work was supported by CRITERIA, an EU project, funded under the Horizon 2020 programme, grant agreement no. 101021866 and the Ministry of Education, Youth and Sports of the Czech Republic through the e-INFRA CZ (ID:90140). Esaú Villatoro-Tello, was supported partially by Idiap, SNI CONACyT, and UAM-Cuajimalpa Mexico.

References

Biswanath Barik, Erwin Marsi, and Pinar Øzturk. 2016. Event causality extraction from natural science literature.

Brandon Beamer, Alla Rozovskaya, and Roxana Girju. 2008. Automatic semantic relation extraction with multiple boundary generation. In *AAAI*, pages 824–829.

Guillaume Becquin. 2020. Gbe at fincausal 2020, task 2: span-based causality extraction for financial documents. In *Proceedings of the 1st Joint Workshop on Financial Narrative Processing and MultiLing Financial Summarisation*, pages 40–44.

Giannis Bekoulis, Johannes Deleu, Thomas Demeester, and Chris Develder. 2018. Adversarial training for multi-context joint entity and relation extraction. *arXiv preprint arXiv:1808.06876*.

James Bergstra, Brent Komer, Chris Eliasmith, Dan Yamins, and David D Cox. 2015. Hyperopt: a python library for model selection and hyperparameter optimization. *Computational Science & Discovery*, 8(1):014008.

Quoc-Chinh Bui, Breannán Ó Nualláin, Charles A Boucher, and Peter Sloot. 2010. Extracting causal

relations on hiv drug resistance from literature. *BMC bioinformatics*, 11(1):1–11.

Sergio Burdisso, Juan Zuluaga-Gomez, Martin Fajcik, Esaú Villatoro-Tello, Muskaan Singh, Petr Motlicek, and Pavel Smrz. 2022. IDIAPers @ causal news corpus 2022: Efficient causal relation identification through a prompt-based few-shot approach. In *The 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE @ EMNLP 2022)*. Association for Computational Linguistics.

Tommaso Caselli and Piek Vossen. 2017. The event storyline corpus: A new benchmark for causal and temporal relation extraction. In *Proceedings of the Events and Stories in the News Workshop*, pages 77–86.

Jifan Chen, Qi Zhang, Pengfei Liu, Xipeng Qiu, and Xuan-Jing Huang. 2016. Implicit discourse relation detection via a deep architecture with gated relevance network. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1726–1735.

Tirthankar Dasgupta, Rupsa Saha, Lipika Dey, and Abir Naskar. 2018. Automatic extraction of causal relations from text using linguistically informed deep neural networks. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 306–316.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Quang Do, Yee Seng Chan, and Dan Roth. 2011. Minimally supervised event causality identification. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 294–303.

Jesse Dunietz, Gregory Burnham, Akash Bharadwaj, Owen Rambow, Jennifer Chu-Carroll, and David Ferrucci. 2020. To test machine comprehension, start by defining comprehension. *arXiv preprint arXiv:2005.01525*.

Martin Fajcik, Martin Docekal, Karel Ondrej, and Pavel Smrz. 2021. [R2-D2: A modular baseline for open-domain question answering](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 854–870, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Martin Fajcik, Josef Jon, Martin Docekal, and Pavel Smrz. 2020. [BUT-FIT at SemEval-2020 task 5: Automatic detection of counterfactual statements with deep pre-trained language representation models](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 437–444, Barcelona (online). International Committee for Computational Linguistics.

- Amir Feder, Katherine A Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E Roberts, et al. 2021a. Causal inference in natural language processing: Estimation, prediction, interpretation and beyond. *arXiv preprint arXiv:2109.00725*.
- Amir Feder, Nadav Oved, Uri Shalit, and Roi Reichart. 2021b. Causalm: Causal model explanation through counterfactual language models. *Computational Linguistics*, 47(2):333–386.
- Daniela Garcia et al. 1997. Coatis, an nlp system to locate expressions of actions connected by causality links. In *International Conference on Knowledge Engineering and Knowledge Management*, pages 347–352. Springer.
- Roxana Girju. 2003. Automatic detection of causal relations for question answering. In *Proceedings of the ACL 2003 workshop on Multilingual summarization and question answering*, pages 76–83.
- Roxana Girju, Preslav Nakov, Vivi Nastase, Stan Szpakowicz, Peter Turney, and Deniz Yuret. 2009. Classification of semantic relations between nominals. *Language Resources and Evaluation*, 43(2):105–121.
- Ashwin Ittoo and Gosse Bouma. 2013. Minimally-supervised learning of domain-specific causal relations using an open-domain corpus as knowledge base. *Data & Knowledge Engineering*, 88:142–163.
- Xianxian Jin, Xinzhi Wang, Xiangfeng Luo, Subin Huang, and Shengwei Gu. 2020. Inter-sentence and implicit causality extraction from chinese corpus. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 739–751. Springer.
- Yohan Jo, Seojin Bang, Chris Reed, and Eduard Hovy. 2021. Classifying argumentative relations using logical mechanisms and argumentation schemes. *Transactions of the Association for Computational Linguistics*, 9:721–739.
- Ning Kang, Bharat Singh, Chinh Bui, Zubair Afzal, Erik M van Mulligen, and Jan A Kors. 2014. Knowledge-based extraction of adverse drug events from biomedical text. *BMC bioinformatics*, 15(1):1–8.
- Christopher SG Khoo, Syin Chan, and Yun Niu. 2000. Extracting causal knowledge from a medical database using graphical patterns. In *Proceedings of the 38th annual meeting of the association for computational linguistics*, pages 336–343.
- Christopher SG Khoo, Jaklin Kornfilt, Robert N Oddy, and Sung Hyon Myaeng. 1998. Automatic extraction of cause-effect information from newspaper text without knowledge-based inferencing. *Literary and Linguistic Computing*, 13(4):177–186.
- Canasai Kruengkrai, Kentaro Torisawa, Chikara Hashimoto, Julien Kloetzer, Jong-Hoon Oh, and Masahiro Tanaka. 2017. Improving event causality recognition with multiple background knowledge sources using multi-column convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Fei Li, Meishan Zhang, Guohong Fu, and Donghong Ji. 2017. A neural joint model for entity and relation extraction from biomedical text. *BMC bioinformatics*, 18(1):1–11.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Eugenio Martínez-Cámara, Vered Shwartz, Iryna Gurevych, and Ido Dagan. 2017. Neural disambiguation of causal lexical markers based on context. In *IWCS 2017—12th International Conference on Computational Semantics—Short papers*.
- George A. Miller. 1994. *WordNet: A lexical database for English*. In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Paramita Mirza, Rachele Sprugnoli, Sara Tonelli, and Manuela Speranza. 2014. Annotating causality in the tempeval-3 corpus. In *Proceedings of the EACL 2014 Workshop on Computational Approaches to Causality in Language (CAtoCL)*, pages 10–19.
- Nasrin Mostafazadeh, Alyson Grealish, Nathanael Chambers, James Allen, and Lucy Vanderwende. 2016. Caters: Causal and temporal relation scheme for semantic annotation of event structures. In *Proceedings of the Fourth Workshop on Events*, pages 51–61.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. *Pytorch: An imperative style, high-performance deep learning library*. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Edoardo Ponti and Anna-Leena Korhonen. 2017. Event-related features in feedforward neural networks contribute to identifying implicit causal relations in discourse.
- Kira Radinsky, Sagie Davidovich, and Shaul Markovitch. 2012. Learning causality for news events prediction. In *Proceedings of the 21st international conference on World Wide Web*, pages 909–918.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits

- of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Dhanya Sridhar and Lise Getoor. 2019. Estimating causal effects of tone in online debates. *arXiv preprint arXiv:1906.04177*.
- Cong Sun, Zhihao Yang, Ling Luo, Lei Wang, Yin Zhang, Hongfei Lin, and Jian Wang. 2019. A deep learning approach with deep contextualized word representations for chemical–protein interaction extraction from biomedical literature. *IEEE Access*, 7:151034–151046.
- Chenhao Tan, Lillian Lee, and Bo Pang. 2014. The effect of wording on message propagation: Topic and author-controlled natural experiments on twitter. *arXiv preprint arXiv:1405.1438*.
- Fiona Anting Tan, Ali Hürriyetoğlu, Tommaso Caselli, Nelleke Oostdijk, Hansi Hettiarachchi, Tadashi Nomoto, Onur Uca, and Farhana Ferdousi Liza. 2022a. Event causality identification with causal news corpus - shared task 3, CASE 2022. In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2022)*, Online. Association for Computational Linguistics.
- Fiona Anting Tan, Ali Hürriyetoğlu, Tommaso Caselli, Nelleke Oostdijk, Tadashi Nomoto, Hansi Hettiarachchi, Iqra Ameer, Onur Uca, Farhana Ferdousi Liza, and Tiancheng Hu. 2022b. [The causal news corpus: Annotating causal relations in event sentences from news](#). In *Proceedings of the Language Resources and Evaluation Conference*, pages 2298–2310, Marseille, France. European Language Resources Association.
- Fiona Anting Tan, Xinyu Zuo, and See-Kiong Ng. 2022c. [Unicausal: Unified benchmark and model for causal text mining](#).
- Victor Veitch, Dhanya Sridhar, and David Blei. 2020. Adapting text embeddings for causal inference. In *Conference on Uncertainty in Artificial Intelligence*, pages 919–928. PMLR.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. *Advances in neural information processing systems*, 28.
- Jue Wang and Wei Lu. 2020. Two are better than one: Joint entity and relation extraction with table-sequence encoders. *arXiv preprint arXiv:2010.03851*.
- Linlin Wang, Zhu Cao, Gerard De Melo, and Zhiyuan Liu. 2016. Relation classification via multi-level attention cnns. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1298–1307.
- Bonnie Webber, Rashmi Prasad, Alan Lee, and Aravind Joshi. 2019. The penn discourse treebank 3.0 annotation manual. *Philadelphia, University of Pennsylvania*, 35:108.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Zach Wood-Doughty, Ilya Shpitser, and Mark Dredze. 2018. Challenges of using text classifiers for causal inference. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, volume 2018, page 4586. NIH Public Access.
- Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. 2021. A unified generative framework for various ner subtasks. 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 5808–5822.
- Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Graph convolutional networks for text classification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 7370–7377.
- Justine Zhang, Sendhil Mullainathan, and Cristian Danescu-Niculescu-Mizil. 2020. Quantifying the causal effects of conversational tendencies. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2):1–24.
- Yuhao Zhang, Peng Qi, and Christopher D Manning. 2018. Graph convolution over pruned dependency trees improves relation extraction. *arXiv preprint arXiv:1809.10185*.
- Shan Zhao, Minghao Hu, Zhiping Cai, and Fang Liu. 2021. Modeling dense cross-modal interactions for joint entity-relation extraction. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 4032–4038.

Hyperparameter	Value
learning rate	.0002
hidden dropout	.1436
attention dropout	.4719
weight decay	.0214
minibatch size	8
warmup proportion	.1570
scheduler	constant (no lr decrease)
max steps	10,000
max gradient norm	1

Table 4: Hyperparameter setting used in this work.

A Hyperparameters

In Table 4, we report the exact hyperparameters used when fine-tuning T5. Warmup proportion, weight decay, and dropouts are in the (0,1) range (for instance, .4719 means 47.19%).

B Example of Inputs

The input format and label format for a single training example, a sentence with 2 CES triplets, are illustrated in Figure. 2.

ENCODER INPUT: " _I _think _independent _film _producers _have _the _responsibility _to _document _what _mainstream _media _failed _to _report _on . _" _But _on _the _e ve _of _the _protest s _' _second _anniversary _ , _Chan _claims _all _of _Hong _Kong _' _s _major _cinema s _are _refus ing _to _show _his _film _ , _the _result _ , _he _suspect s _ , _of _creep ing _self - censor ship _as _businesses _shy _away _from _off ending _Beijing _ . **_history :**

DECODER PREFIX: **_cause :**

DECODER TARGET: **_cause :** _businesses _shy _away _from _off ending _Beijing

ENCODER INPUT: " _I _think _independent _film _producers _have _the _responsibility _to _document _what _mainstream _media _failed _to _report _on . _" _But _on _the _e ve _of _the _protest s _' _second _anniversary _ , _Chan _claims _all _of _Hong _Kong _' _s _major _cinema s _are _refus ing _to _show _his _film _ , _the _result _ , _he _suspect s _ , _of _creep ing _self - censor ship _as _businesses _shy _away _from _off ending _Beijing _ . **_cause :** business e s _shy _away _from _off ending _Beijing **_history :**

DECODER PREFIX: **_effect :**

DECODER TARGET: **_effect :** _creep ing _self - censor ship

ENCODER INPUT: " _I _think _independent _film _producers _have _the _responsibility _to _document _what _mainstream _media _failed _to _report _on . _" _But _on _the _e ve _of _the _protest s _' _second _anniversary _ , _Chan _claims _all _of _Hong _Kong _' _s _major _cinema s _are _refus ing _to _show _his _film _ , _the _result _ , _he _suspect s _ , _of _creep ing _self - censor ship _as _businesses _shy _away _from _off ending _Beijing _ . **_cause :** business e s _shy _away _from _off ending _Beijing **_effect :** cre e ping _self - censor ship **_history :**

DECODER PREFIX: **_signal :**

DECODER TARGET: **_signal :** _as

ENCODER INPUT: " _I _think _independent _film _producers _have _the _responsibility _to _document _what _mainstream _media _failed _to _report _on . _" _But _on _the _e ve _of _the _protest s _' _second _anniversary _ , _Chan _claims _all _of _Hong _Kong _' _s _major _cinema s _are _refus ing _to _show _his _film _ , _the _result _ , _he _suspect s _ , _of _creep ing _self - censor ship _as _businesses _shy _away _from _off ending _Beijing _ . **_history :** **_cause :** _businesses _shy _away _from _off ending _Beijing **_effect :** _creep ing _self - censor ship **_signal :** _as

DECODER PREFIX: **_cause :**

DECODER TARGET: **_cause :** _creep ing _self - censor ship

ENCODER INPUT: " _I _think _independent _film _producers _have _the _responsibility _to _document _what _mainstream _media _failed _to _report _on . _" _But _on _the _e ve _of _the _protest s _' _second _anniversary _ , _Chan _claims _all _of _Hong _Kong _' _s _major _cinema s _are _refus ing _to _show _his _film _ , _the _result _ , _he _suspect s _ , _of _creep ing _self - censor ship _as _businesses _shy _away _from _off ending _Beijing _ . **_cause :** cre e ping _self - censor ship **_history :** **_cause :** _businesses _shy _away _from _off ending _Beijing **_effect :** _creep ing _self - censor ship **_signal :** _as

DECODER PREFIX: **_effect :**

DECODER TARGET: **_effect :** _all _of _Hong _Kong _' _s _major _cinema s _are _refus ing _to _show _his _film

ENCODER INPUT: " _I _think _independent _film _producers _have _the _responsibility _to _document _what _mainstream _media _failed _to _report _on . _" _But _on _the _e ve _of _the _protest s _' _second _anniversary _ , _Chan _claims _all _of _Hong _Kong _' _s _major _cinema s _are _refus ing _to _show _his _film _ , _the _result _ , _he _suspect s _ , _of _creep ing _self - censor ship _as _businesses _shy _away _from _off ending _Beijing _ . **_cause :** cre e ping _self - censor ship **_effect :** all _of _Hong _Kong _' _s _major _cinema s _are _refus ing _to _show _his _film **_history :** **_cause :** _businesses _shy _away _from _off ending _Beijing **_effect :** _creep ing _self - censor ship **_signal :** _as

DECODER PREFIX: **_signal :**

DECODER TARGET: **_signal :** _the _result _ , _he _suspect s _ , _of

Figure 2: Example of tokenized inputs for a sentence with two annotated CES triplets. Phrases "ENCODER INPUT", "DECODER PREFIX" and "DECODER TARGET" are not parts of the input, and are included for illustrative purposes only. Special sequences (**_cause:**, **_effect:**, **_signal:**, **_history:**) used between concatenated parts of the input are in bold.