Eco EvE : Economic Event Extraction

Meriem Mkhinini, Mohamed Ali B.E.H Aissa, Aboubacar Sidiki Sidibe,
Paul Pike and **Aymen Khelifi**
Kaisens Data, 16 Pl. de l'Iris, 92400 Courbevoie
{mmkhinini, asidiki, ppike, aymen.khelifi}@kaisensdata.fr

Abstract

Every day, important events take place all over the world, but they are reported in various media outlets using various narrative tenses. Identifying real-world events have long been an important NLP problem. This paper, presents a comprehensive and up-to-date approach for economic events extraction in text-based context. We propose a zero-shot approach as an event extraction solution. The novelty of our approach rely in the use of separate glossaries to adapt to the domain application. It does require retraining to each specific type of events. The proposed approach, EcoEVE, is shown to be very effective when working with data from many platforms (Economic Calendar, Economic news. . .). Finally, we also present our ideas on future research directions.

1 Introduction

Event extraction is a challenging and long-researched task in information extraction (Grishman and Sundheim(1996)); (Riloff(1996)). The Automatic Content Extraction (ACE) program (ace()), defines an event as *something that happens. An Event can frequently be described as a change of state.* Based on the ACE program, we identify five main elements that form an event:

- **Event Type** : Thematic Event label describing the general nature of what's described in the sentence.
- **Trigger** : The word that most clearly expresses the event occurrence. In many cases, it is the main verb in the part of the sentence describing the event.
- **Agent** : The doer or instigator of the action denoted by the predicate.

Figure 1: An example of an event with extracted arguments.

- **Patient** : The undergoer of the action or event denoted by the predicate.
- **Time** : When the Event takes place
- **Place** : Where the Event takes place

Some event elements, such as the place or time maybe absent. However, an event my still occur. We give an example in Figure 1, which represents a 'Bankruptcy' event (the event type), triggered by "began" (the event trigger) and accompanied by its extracted arguments - text spans denoting entities that fulfill a set of (semantic) roles associated with the event type (e.g. `AGENT of the event, PATIENT or recipient of the event and TIME of the event`).

In this paper, we study event extraction based on context information, namely economical context. We address the following research question : Can context information improve the accuracy of event identification?

To address this question, we propose a three step model based on zero-shot classification. The later is a technique that allows the association of an appropriate label to a text. This association is irrespective of the text domain and the aspect.

In Section 2, we give a review of existing work. In section 3, we present the details of our proposed approach. In section 4, we present the data used to implement our model. Section 5, gives the results

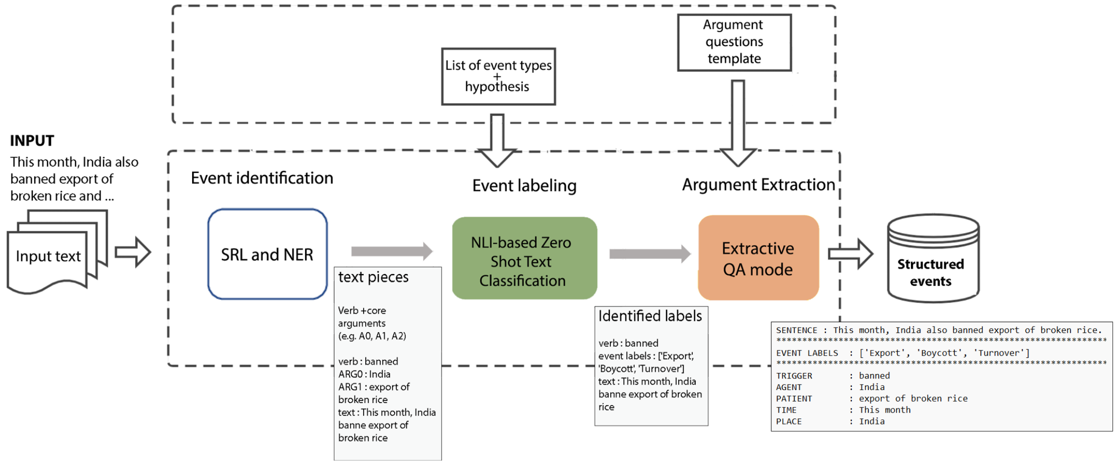

Figure 2: Our pipeline for event extraction.

obtained in our experiments. Finally, in section 6, we conclude the state of our work and present the future improvements.

2 Related Work

Recent successful approaches to event extraction usually require supervision (e.g., (Lin et al.(2020)Lin, Ji, Huang, and Wu)). Particularly, methods relying on expert systems, to define rules based on the occurrence of events in text. Such approaches can be labor-intensive and ignore the semantic meaning of event type labels.Economic events, however, can be formulated in many ways depending on specific domain they report (Arendarenko and Kakkonen(2012)).defining a domain ontology and rules for every application can be time consuming and difficult. Furthermore, defining a set of strict rules often results in low recall scores, since these rules usually cover only a portion of the many various ways in which certain information can be worded.

Zero-shot learning (ZSL) most often refers to a fairly specific type of task: learning a classifier on one set of labels, and then evaluating it on another set of labels that the classifier has never seen before. It has been used much more broadly to make a model do something for which it has not been explicitly trained. Evaluate a language model on downstream tasks (Radford et al.(2019)Radford, Wu, Child, Luan, Amodei, and Sutskever) without refining it directly on those tasks. In their pioneering work on more general zero-shot models, (Yin et al.(2019)Yin, Hay, and Roth) propose to formulate text classification tasks as a textual entailment problem (Dagan et al.(2006)Dagan, Glickman, and Magnini). This correspondence allows for the use of a trained model on natural language inference (NLI) to be used as a zero-shot text classifier for a wide variety of unseen downstream tasks.

Recent work ((Liu et al.(2020)Liu, Chen, Liu, Bi, and Liu); (Du and Cardie(2020))) have emphasized the link between question answering (QA) and EA in supervised system development. Similarly, several efforts have explored unsupervised methods. Using similarity-based methods (Peng et al.(2016)Peng, Song, and Roth) attempted to extract event triggers with minimal supervision. (Huang et al.(2018)Huang, Ji, Cho, Dagan, Riedel, and Voss) and (Lai et al.(2020)Lai, Nguyen, and Dernoncourt) explored trigger and argument extraction in a slightly different setting: training on certain event types and testing on unseen event types. Recently, (Liu et al.(2020)Liu, Chen, Liu, Bi, and Liu) proposed a QA-based argument extraction method that does not handle triggers. To the extent of our knowledge no method has been proposed to extract both event triggers and arguments without any event extraction training data. In this paper, we investigate the possibility of a paradigm for the event extraction task - formulating it as a Zeroshot/question answering (QA) task (Zhang et al.(2021)Zhang, Wang, and Roth).

3 Methodology

We propose a three step based approach for event extraction as illustrated in Figure 2. The first step is event detection: Given input text, we apply Semantic role labeling (SRL) (Collobert and Weston(2008)) to understand the role of each word in a sentence. Then we use Named Entity Recognition(NER) (Schmitt et al.(2019)Schmitt, Kubler, Robert, Papadakis, and LeTraon), to identify key elements in text like names of people, places, monetary values,etc. Once the pre-processing step accomplished, we use zero-shot classification to detect potential events in the Event labeling step. The third and final step is argument extraction. We use a pre-trained QA models to extract specific arguments concerning the event. All pre-trained models we use are based on BERT and BART, including a Zero Shot, a bart-large (Lewis et al.(2019)Lewis, Liu, Goyal, Ghazvininejad, Mohamed, Levy, Stoyanov, and Zettlemoyer) model trained on the MultiNLI (MNLI) data-set (Williams et al.(2018)Williams, Nangia, and Bowman) , and extractive QA model (Lyu et al.(2021)Lyu, Zhang, Sulem, and Roth) trained on QAMR. We illustrate each step of our approach in details in the following sections.

3.1 Pipeline Overview

Our pipeline for event extraction relies on The Zero Shot model (green box in Figure 2). In the pre-processing stage, we segment the text into sentences and apply data cleaning techniques based on the Spacy Python library (Honnibal and Montani(2017)). Then, for each sentence, given a set of event types , it creates hypothesis template of "this example is . . . " for each type to predict the type of the premise. If the inference is entailment, it means that the premise belongs to that type. Finally, the extractive QA model (orange box in Figure 2) takes as input a context and a Wh-question. For each sentence, it extract the event trigger and type. Thus iteratively identifying candidate event arguments (spans of text) in the input sentence.

3.2 Trigger Extraction

We formalize Trigger Extraction as a Zero Shot Classification task. To get potential event triggers from a sentence, we first perform semantic role labeling (SRL) (Gardner et al.(2017)Gardner, Grus, Neumann, Tafjord, Dasigi, Liu, Peters, Schmitz, and Zettlemoyer) as a pre-processing step. We use the BERT-based Verb SRL model. The sentence is then split into "text fragments", each containing its SRL predicate and its core arguments (A0, A1, A2, etc.). Then, we pass each text fragment as a premise to the zero-shot model as well as a list of event types. The model returns the list of event types sorted with scores (most likely to be linked to the text fragment) with the first being the highest entailment probability. Then, we pass the highest three labels combined with hypotheses of the form "this text is about . . . " or "is related to . . . " inspired by (Yin et al.(2019)Yin, Hay, and Roth). For every hypothesis, the model returns the probability that it is entailed by the premise. If the very best entailment probability throughout all occasion sorts surpasses a threshold, we output the corresponding SRL predicate as an event trigger of this type.

3.3 Argument Extraction

We formalize the task of Argument Extraction as a sequence of QA interactions with the pre-trained extractive QA model. Given an input sentence and the extracted trigger, we ask a set of questions, and retrieve the QA model's answers as argument predictions. We design two templates with annotation guideline based questions as shown in table 1.

We describe the agent as the noun phrase or pronoun that identifies the person or thing which initiates or performs an action in a sentence. The patient being the person or thing that receives an action in a sentence. Time and place are straightforward. The place (Locative in linguistics) is the specification of the place where the action or event denoted by the predicate is situated. The time or date in the other hand is the period when the action or event took place. Then, the cause being the reason why the action happens and the aim is the reason for doing the action. Finally the variation and old/new value being the values that changed or are talked about in the sentence.

For each question, the model returns the probability for the answer. If the highest probability across two question templates surpasses a threshold, we output the corresponding argument. Since many argument types in the event template do not occur in every sentence. For example in the sentence : `The imports of their struggling economy drastically outweigh the exports.`, there is only an AGENT and PATIENT argument.

4 Data Description

One of the main difficulties we faced building our model, is finding open source event data sets. In this section, we describe the EcoEVE economic event labels annotation dictionary. The goal of

Argument	Template(1)	Template(2)
AGENT	Who is responsible for the {trigger} ?	what is responsible for the {trigger} ?
PATIENT	who is {triggered} ?	what is {triggered} ?
TIME	when the {trigger} happen ?	in what time the {trigger} happen ?
PLACE	where the {trigger} happen ?	in what place the {trigger} happen ?
AIM	why the {trigger} happen ?	for what reason the {trigger} happen ?
OLD_VALUE	what is the old value before the {trigger} ?	from what value it have {triggered} ?
NEW_VALUE	what is the new value after the {trigger} ?	to what value it have {triggered} ?
VARIATION	what is the variation of the {trigger} ?	
CAUSE	how the {trigger} happen ?	what is cause of the {trigger} ?

Table 1: Arguments and corresponding questions from templates.

	title	event	sentence
0	Consumer Inflation Expectations	Inflation	Inflation expectations in Australia increased ...
1	Consumer Inflation Expectations	Inflation	The November inflation expectation figure, bas...
2	Consumer Inflation Expectations	Inflation	The report further noted that uncertainty abou...
3	Consumer Inflation Expectations	Inflation	Wage expectations continue to be weak, with be...
4	Unemployment Rate	Employment	The number of unemployed grew by 120 thousand ...

Figure 3: test data frame.

the EcoEVE labels is to enable unsupervised data-driven event extraction in economic news. To do so, we use a lexicon of English event labels. We scrapped articles from the news site The Financial Times, Wikipedia articles describing companies and major economic events, Economic calendars(containing indicators in real-time as economic events are announced and the immediate global market impact) and The Economist articles (Authoritative global news coverage of world politics, economics and business). In total, we collected over 500 news articles. Combined with Glossary of economics, containing 473 economic term and definition. We identified 70 event labels. These events and activities relate to specific instances of events mentioned in the articles. For example, in some economic calendars, events are divided into categories that describe the event like Interest Rate, Inflation, GDP Growth, Foreign Trade,etc.

5 Experimental Setup and Results

To evaluate our approach, we built the tool EcoEVE. Event extraction has two tasks: Trigger/argument identification and event labeling, with trigger/argument having three sub tasks (trigger identification, argument identification, and argument classification). We test each task separately.

5.1 Trigger/Argument Identification

To evaluate trigger and argument extraction, we use an existing Data Collection used by (Liu et al.(2019)Liu, Huang, and Zhang) including 574 news groups, 2433 news reports, 5830 sentences. This data set gives us the arguments and the trigger verb for each sentence. However, since the results of this dataset only take into account the root verb of a sentence, we made some adjustments to our model. Thus, for this part, we used the Spacy Linguistic Features model to only work with the ROOT verbs as the syntactic dependency, i.e. the relationship between tokens. Our approach gives us results of more than 80% of triggers and about 50% of arguments with semantically correct types were successfully mapped.

5.2 Event Labeling

For a lack of official open source test set, we collected data from Trading Economics [1]. The site provides accurate information for 196 countries, including historical data and forecasts for more than 20 million economic indicators, exchange rates,

[1] https://tradingeconomics.com/

stock indices, government bond yields and commodity prices. We suppose that the articles titles contain the main event discussed in the article text. Then, for each text, we segment it into sentences. We manually selected the sentences that relate to the event proposed by the title. Our final data frame contains 436 sentences. Figure 3 shows the first few rows of our data set, including titles, events, and sentences. We identified 18 event types manually.

We tested our approach without changing the original event type lexicon. In doing so, we obtain real results as we would on a potential unknown use case. Since our model can predict more than one event label for a sentence, we suppose that if one of them is the same as the manually identified event type, the event label is correct. The tool successfully mapped 89% of the event labels in our test set.

6 Conclusion and Perspective

In this paper, we present a novel approach for event extraction based on zero-shot and QA event extraction system. We study the performance of QA/zero-shot models on event extraction data sets and how these strategies affect the performance of our pipeline. Our approach have shown positive result and performance. However, we also identified several key challenges of the current approach. For instance, a more generic formulation in the event labeling stage can lead to better performance and flexibility. For future work, we are working on incorporating a broader context, a paragraph/document level context, into our methods to improve prediction accuracy. We could also further refine the QA/zero-shot models to improve their performance for the event extraction task.

References

ACE (Automatic Content Extraction) English Annotation Guidelines for Event.

Ernest Arendarenko and Tuomo Kakkonen. 2012. Ontology-based information and event extraction for business intelligence. In *AIMSA*.

Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th international conference on Machine learning*, pages 160–167.

Ido Dagan, Oren Glickman, and Bernardo Magnini. 2006. The pascal recognising textual entailment challenge. In *Machine Learning Challenges. Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Tectual Entailment*, pages 177–190, Berlin, Heidelberg. Springer Berlin Heidelberg.

Xinya Du and Claire Cardie. 2020. https://doi.org/10.18653/v1/2020.emnlp-main.49 Event extraction by answering (almost) natural questions. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 671–683, Online. Association for Computational Linguistics.

Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke S. Zettlemoyer. 2017. http://arxiv.org/abs/arXiv:1803.07640 Allennlp: A deep semantic natural language processing platform.

Ralph Grishman and Beth Sundheim. 1996. https://aclanthology.org/C96-1079 Message Understanding Conference- 6: A brief history. In *COLING 1996 Volume 1: The 16th International Conference on Computational Linguistics*.

Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.

Lifu Huang, Heng Ji, Kyunghyun Cho, Ido Dagan, Sebastian Riedel, and Clare Voss. 2018. https://doi.org/10.18653/v1/P18-1201 Zero-shot transfer learning for event extraction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2160–2170, Melbourne, Australia. Association for Computational Linguistics.

Viet Dac Lai, Thien Huu Nguyen, and Franck Dernoncourt. 2020. https://doi.org/10.18653/v1/2020.nuse-1.5 Extensively matching for few-shot learning event detection. In *Proceedings of the First Joint Workshop on Narrative Understanding, Storylines, and Events*, pages 38–45, Online. Association for Computational Linguistics.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2019. http://arxiv.org/abs/1910.13461 BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *CoRR*, abs/1910.13461.

Ying Lin, Heng Ji, Fei Huang, and Lingfei Wu. 2020. https://doi.org/10.18653/v1/2020.acl-main.713 A joint neural model for information extraction with global features. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7999–8009, Online. Association for Computational Linguistics.

Jian Liu, Yubo Chen, Kang Liu, Wei Bi, and Xiaojiang Liu. 2020. https://doi.org/10.18653/v1/2020.emnlp-main.128 Event extraction as machine reading comprehension. In *Proceedings of the 2020 Conference* *on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1641–1651, Online. Association for Computational Linguistics.

Xiao Liu, Heyan Huang, and Yue Zhang. 2019. https://doi.org/10.18653/v1/P19-1276 Open domain event extraction using neural latent variable models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2860–2871, Florence, Italy. Association for Computational Linguistics.

Qing Lyu, Hongming Zhang, Elior Sulem, and Dan Roth. 2021. https://doi.org/10.18653/v1/2021.acl-short.42 Zero-shot event extraction via transfer learning: Challenges and insights. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 322–332, Online. Association for Computational Linguistics.

Haoruo Peng, Yangqiu Song, and Dan Roth. 2016. https://doi.org/10.18653/v1/D16-1038 Event detection and co-reference with minimal supervision. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 392–402, Austin, Texas. Association for Computational Linguistics.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.

Ellen Riloff. 1996. *Automatically generating extraction patterns from untagged text.* In Proceedings of the national conference on artificial intelligence.

Xavier Schmitt, Sylvain Kubler, Jérémy Robert, Mike Papadakis, and Yves LeTraon. 2019. A replicable comparison study of ner software: Stanfordnlp, nltk, opennlp, spacy, gate. In *2019 Sixth International Conference on Social Networks Analysis, Management and Security (SNAMS)*, pages 338–343. IEEE.

Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. http://aclweb.org/anthology/N18-1101 A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.

Wenpeng Yin, Jamaal Hay, and Dan Roth. 2019. https://doi.org/10.18653/v1/D19-1404 Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. 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 3914–3923, Hong Kong, China. Association for Computational Linguistics.

Hongming Zhang, Haoyu Wang, and Dan Roth. 2021. https://doi.org/10.18653/v1/2021.findings-acl.114 Zero-shot Label-aware Event Trigger and Argument Classification. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1331–1340, Online. Association for Computational Linguistics.