A Closer Look to Your Business Network: Multitask Relation Extraction from Economic and Financial French Content

Hadjer Khaldi [1, 2],
Farah Benamara[2],
Camille Pradel [1],
Nathalie Aussenac Gilles[2],
[1] Geotrend, Toulouse, France
[2] IRIT- CNRS- Université de Toulouse, Toulouse, France
hadjer@geotrend.fr, farah.benamara@irit.fr, camille@geotrend.fr , nathalie.aussenac-gilles@irit.fr

Abstract

Online textual content constitutes a valuable source of information for market stakeholders, enabling them to unveil their business network's most important operations and interactions, and to gain insights about their customers, business partners, and competitors, in order to make well-informed strategic decisions. Due to the problem of information overload, manually extracting this information remains a laborious task for professionals, making the use of Information Extraction technologies a powerful asset. In this context, this paper concerns discovering business relations between companies (e.g. company-partner) from French content on the web. We present a new dataset for business relation extraction at the sentence level and develop a set of deep learning experiments to distinguish between business vs. non-business relations, as well as identify five types of business relations according to a predefined taxonomy. Our results are encouraging, showing that multitask architectures are the most productive beating several strong state of the art baselines.

Introduction

Web 2.0 has shaped the way information is shared in the world and accelerated the generation and spread of economic and financial information online making it accessible in different formats such as: companies announcements, industry research reports, online news articles, and policy statements. This strategic information is used by market stakeholders to take well-informed decisions (Oberlechner and Hocking 2004). For example, in order to maximize any gain while minimizing any possible losses, banks and investors need to analyze their clients and investees business relations and stock prices to evaluate the risks associated with giving them a loan or an investment. At the same time, companies should conversely have a real-time overview about their business network changes and competitors' strategies, to be able to detect threats or opportunities, therefore adapt their strategies in order to thrive and remain competitive in the market (Sewlal 2004).

In a rapidly changing business environment, economic and financial textual information is generated online in huge amounts and at a rapid pace, making its exploitation by market stakeholders a challenging task. Nowadays, professionals manually scrutinize tons of financial statements (Berns et al. 2021), companies announcements (Han, Hao, and Huang 2018), and online news articles (Liang et al. 2020) to identify knowledge about their competitors, customers, the market trends and movements while requiring a lot of time and effort. Hence, the availability of systems that automatically extract this information (e.g. named entities, relations, events) from textual content between market entities (e.g., *startups*, *companies*, *non-profit organizations*, etc.) can be an essential tool. For example, from the sentence (1) below taken from our dataset, one can infer that the company *Airbus* is a supplier for the company *Inmarsat*.

(1) Le groupe Airbus a signé avec Inmarsat un contrat de livraisons pour la réalisation de trois satellites géostationnaires reconfigurables en orbite.
(The Airbus group has signed a contract with Inmarsat for the delivery of three reconfigurable geostationary satellites in orbit).

Roughly, state of the art on information extraction from financial textual data concerns either event extraction or binary business relation extraction (henceforth BRE) involving two organizations. Event extraction aims at identifying event triggers, and their arguments (which can be companies and firms) (Lefever and Hoste 2016; Jacobs, Lefever, and Hoste 2018; Qian et al. 2019; Jacobs and Hoste 2021; Xingyue, Liping, and Zhiwei 2021; Wang et al. 2021) and has been used in multiple applications such as stock market prediction (Chen et al. 2019; Usmani and Shamsi 2021), perceiving market trends (Berns et al. 2021; Han, Hao, and Huang 2018), assisting investors' decisions and risk analysis (Liang et al. 2020; Hogenboom et al. 2015). BRE on the other hand aims at discovering either Inner-Organizational (Inner-ORG) relations linking a company and its components (e.g. company-employees, company-CEO), or Inter-Organizational (Inter-ORG) for relations involving different companies (e.g. company-customer, company-partner) (Zhao, Jin, and Liu 2010; Zuo et al. 2017). BRE has shown to be crucial to valuate companies (Zuo et al. 2017), analyze complex emerging business ecosystems (Braun et al. 2018), understand industries structures (Yamamoto et al. 2017), extract competitive intelligence (Zhao, Jin, and Liu 2010; Xu et al. 2011), and reduce credit risk for financial institutions by identifying links between customer groups (Yan et al. 2019).

Our paper focuses on Inter-ORG BRE that may hold between two organizations expressed at the sentence level. Compared to domain-specific relation extraction (like the biomedical (Khachatrian et al. 2019; Li et al. 2020; Xing, Luo, and Song 2020) or food (Wiegand et al. 2012; Cenikj, Seljak, and Eftimov 2021) domains), BRE has received much less attention in the literature, despite its strategic importance for organizations' decision-makers. Most existing works have used semi-supervised approaches relying either on lexico-syntactic patterns generated from dependency trees (Braun et al. 2018), or lexical patterns based on a list of keywords which are specific to each predefined relation type (Lau and Zhang 2011). Yamamoto et al. (2017) used *DeepDive*, a machine learning system based on Markov Logic Network to extract business relations from web news articles. Neural networks have recently been used. For example, Yan et al. (2019) and Collovini et al. (2020) extract relations between Fintech companies from news text using Conditional Random Field and Bi-directional Gated Recurrent Units, while De Los Reyes et al. (2021) rely on BERT pre-trained language model (Devlin et al. 2019).

These approaches share four main limitations: **(1)** they rely on datasets that are either small (less than 3k) (De Los Reyes et al. 2021; Collovini et al. 2020) or not freely available to the research community (Yan et al. 2019; Yamamoto et al. 2017), **(2)** they generally consider only two relations (Yamamoto et al. 2017; Lau and Zhang 2011; De Los Reyes et al. 2021), **(3)** they target English data even if BRE has been explored in few other languages such as German (Braun et al. 2018), Chinese (Yan et al. 2019), and Portuguese (De Los Reyes et al. 2021; Collovini et al. 2020), and **(4)** they rarely account for the absence of business relations which makes BRE in real-world applications more challenging due to their irregular linguistic patterns and their high number of instances in the data (e.g., around 63% and 54% in the corpora used in (Khaldi et al. 2021) and (De Los Reyes et al. 2021) respectively).

We aim here to address these limitations through four main contributions:

1. *The first large French dataset*[1] *for BRE of about 10k instances* annotated according to a predefined taxonomy of five business relations.
2. *A set of deep learning monotask experiments to classify business relations* relying on different contextual representations of input sentences.
3. *A new multitask model* that jointly learns relation identification (a binary classification: Business Vs. Non-Business) as an auxiliary task learned jointly with relation classification (multiclass classifier) using multitask objectives in order to reduce the noise due to negative instances and help the model to learn more discriminative features about business relations.
4. *Qualitative analysis of our results* highlighting main causes of classification errors.

Our results show that the proposed multitask architecture achieves the best scores, beating several competitive state of the art baselines. The reminder of the paper is organized as follows: We first present our data in Section 2, then our models in Section 3. We describe the experiments carried out and the results in Section 4 and give the error analysis in Section 5.

BizRel Dataset

To build our dataset, we consider 5 business relations described below, following the characterization proposed in (Khaldi et al. 2021) where a relation linking two named entities of type *Organization* (henceforth *EO*) is expressed at the sentence level:

- INVESTMENT: an *EO* is a subsidiary of another *EO*, or *EO* holds (all or part) of the shares of another *EO*.

  (2) Le missilier européen MBDA (filiale commune d'Airbus, Leonardo et BAE) espère que l'accord signé à Helsinki lui donnera à terme accès à des financements pour développer de nouvelles versions de son missile antichar de moyenne portée (MMP). (The European missile MBDA (a joint subsidiary of Airbus, Leonardo, and BAE) hopes that the agreement signed in Helsinki will eventually give it access to financing to develop new versions of its medium-range anti-tank missile (MMP).)

- COMPETITION: a competition/rivalry between two *EO*s providing the same goods or services, or wanting to access the same relatively small market.

  (3) Boeing et l'avionneur brésilien Embraer, rival de Bombardier sur les avions régionaux, ont annoncé discuter sur un éventuel rapprochement de leurs activités. (Boeing and the brazilian aircraft manufacturer Embraer, Bombardier's regional aircraft rival, have announced discussions on a possible merger of their activities.)

- COOPERATION: a contractual cooperation between two *EO*s, or when two *EO*s work together on the same project.

  (4) Depuis le 25 novembre 2017, 32 associations et startups , 400.000 citoyen.nes, la Fondation Kering, Facebook et la Région Île-de-France ont travaillé ensemble avec Make.org pour élaborer le premier plan de actions de la société civile contre les violences faites aux femmes. (Since November 25th, 2017, 32 associations and startups, 400,000 citizens, the Kering Foundation, Facebook, and the Île-de-France region have worked together with Make.org to develop the first civil society action plan against violence against women.)

- LEGAL PROCEEDING: one *EO* launches a legal proceedings against another *EO*.

[1] Link to BizRel dataset

(5) Grégoire Triet a représenté Shionogi dans une action en contrefaçon de brevet portant sur un médicament contre le VIH, qui l'a opposé à Merck et ses filiales.
(Grégoire Triet represented Shionogi in a patent infringement action relating to an HIV drug, which brought him against Merck and its subsidiaries.)

- SALE-PURCHASE: one *EO* is a client of another, or supplies it with goods or services.

(6) Le capot d'un réacteur d'un Airbus A320 de la compagnie Frontier Airlines s'est rompu en plein décollage.
(The engine hood of a Frontier Airlines Airbus A320 broke during take-off.)

We add the relation OTHERS to account for any other non-business relation, or the absence of a relation. For example, in (7), one can infer two relations: *Investment (EO_1,EO_2)* and *Investment (EO_3, EO_2)*. However, since no business relation links EO_1 to EO_3, it is therefore annotated as *Others (EO_1,EO_3)* in our dataset.

(7) En 2016, [General Motors]$_1$ a acheté [Cruise Automation]$_2$ pour 1 milliard de dollars, rejoint début juin par le japonais [Softbank]$_3$, qui a annoncé y investir 2,25 milliards de dollars.
(In 2016, [General Motors]$_1$ bought [Cruise Automation]$_2$ for 1$ billion, joined in early June by Japan's [Softbank]$_3$, which announced it was investing 2.25$ billion.)

Note that this relation can have many other linguistic patterns which resemble the ones used to express one of our five business relations, making this relation irregular, then very hard to predict.

This taxonomy has been used to manually annotate sentences extracted from the web. We retrieved textual contents from various sources (e.g. companies websites, online news articles, industrial reports, excluding social media, e-commerce, and code versioning websites.) by requesting search engines API using a list of keywords related to various business activity fields such as *autonomous cars*, *3D printing*, etc. The sentences to annotate are selected according to two main criteria: **(i)** They must contain at least two entities of type ORG as predicted by both Spacy and StanfordNLP, two well-known named entity taggers; and **(ii)** Sentences whose words are at least 95% of type ORG are discarded. The collected sentences were manually annotated by six non-domain-expert French speakers via the collaborative annotation platform *Isahit*[2]. More details about the annotation process and guideline can be found in (Khaldi et al. 2021). The average Kappa (Cohen 1960) between the annotators and the experts is 0.685, which is an acceptable agreement given the complexity of the task (many relations are implicitly expressed and the large context within the sentence (41 words on average) makes the annotation hard).

Table 1 summarizes the distribution of instances per relation type. We can notice that OTHERS is the most frequent relation (near 68%), which is almost similar to the proportion observed in our English dataset (63%) (Khaldi et al. 2021). This is also in line with the proportion of negative classes reported in many binary relation extraction benchmarks (e.g 79% in TACRED (Zhang et al. 2017), 50% in BioRel (Xing, Luo, and Song 2020), and 85.3% in i2b2 2010 (Uzuner et al. 2011)).

#Total	Train.	Dev.	Test.	#Total
Invest.	220	48	47	**315**
Compet.	1,229	263	263	**1,825**
Cooperat.	598	128	129	**855**
Legal.	41	9	9	**58**
Sale.	188	40	40	**268**
Others	4,747	1,017	1,018	**6,782**
#Total	**7,023**	**1,505**	**1,505**	**10,033**

Table 1: Dataset statistics.

For the following experiments, the dataset has been split into a train (70%), a development (15%), and a test set (15%).

Proposed Models

We propose **MT-BizBERT**, a multitask model based on the pre-trained language model BERT (Devlin et al. 2019) fine-tuned on our business relations dataset.

Multitask architectures have shown to be effective for generic and specific relation extraction by learning additional implicit information from either generic auxiliary tasks (Wang and Hu 2020; Zhou et al. 2019) (e.g., dependency parsing, recognizing textual entailment task from GLUE Benchmark), or from multiple domain related tasks (Yadav et al. 2020) (e.g., Protein-Protein Interaction, Drug-Drug interaction). Recently, Lyu et al. (2022) showed that adding binary classification to discriminate between positive and negative instances from the same dataset, as an auxiliary task, could improve the overall performance of relation extraction on SemEval 2010 Task 8 data. Our contribution consists of combining different auxiliary tasks from the same dataset that can help the model learn more discriminative features between business vs. non-business instances, and between business instances only.

The model we propose is shown in Figure 1. Our objective is to assign to a relation instance noted $i = (S, e_1, e_2)$, where S is the sentence, e_1 and e_2 are target entities, one relation type r from a set of predefined relations R. We consider three values for R: bin = $\{$ *business*, *Others*$\}$, biz = $\{$ *Invest., Compet., Cooperat., Legal., Sale.*$\}$, and all = $\{$ *Invest., Compet., Cooperat., Legal., Sale., Others*$\}$.

Let T_R be a relation classification task, where R is the set of pre-defined relations to consider in this task. Our main task is T_{all} performing business relation classification while accounting for the negative relation OTHERS given its importance in end-user systems. We consider two different auxiliary tasks to be learnt jointly with the main task.

- T_{bin} a binary relation classification task (business Vs. non-business) that helps the main task to learn more

[2] https://isahit.com/en/

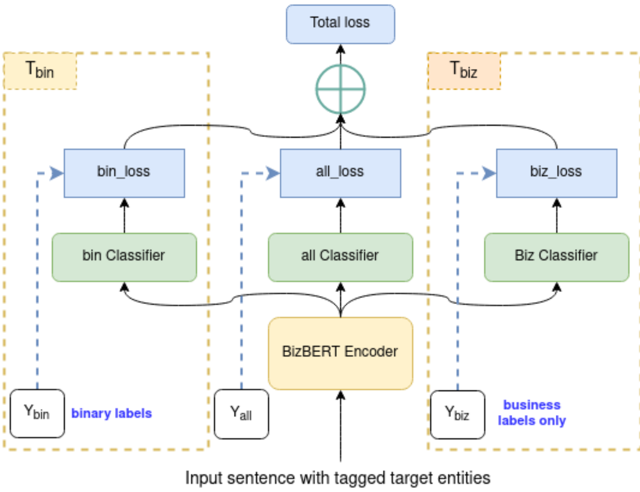

Figure 1: Multitask BRE model architecture.

generic features about business relations and discriminates them from non-business ones (OTHERS),

- T_{biz} a multi-class business relation classification task that learns more specific features about business relations while discarding the noisy negative relation OTHERS which has irregular patterns.

All these three tasks have a shared sentence encoder, named **BizBERT**, initialized using mBERT[3], a pre-trained BERT multilingual language model fine-tuned on our French dataset. At the input level, the beginning and end of each target entity are marked using typed entity marker (i.e. *SUB-ORG* or *OBJ-ORG*) as suggested by (Zhou and Chen 2021). The input sentence is then fed to the sentence encoder to generate a contextualized sentence representation. We use the concatenation of CLS token representation and entities representations as an input for the classifier layer. Each task has its own classifier accounting for the number of relations to consider: 2 for T_{bin}, 5 for T_{biz}, and 6 for T_{all}.

We experiment with three models aiming to improve the main classification task by considering three different combinations of auxiliary tasks: **MT-BizBERT**$^{all+bin}$, **MT-BizBERT**$^{all+biz}$ and **MT-BizBERT**$^{all+biz+bin}$. The models have been trained using cross-entropy loss, where one loss is calculated per task in each model.

Experiments & Results

Baselines

We compare our models to monotask baselines that have shown to be quite effective for generic binary relation extraction (e.g. *Component-Whole*, *Cause-Effect*) on the SemEval 2010 Task 8 dataset (Hendrickx et al. 2010). Note that more competitive models exist such as LUKE (Yamada et al. 2020) and K-Adapter (Wang et al. 2020)). However the adaptation of these models to French is costly since it requires either a re-training from scratch of the language models or the use of language-dependent linguistic resources which are not always available for the French language. The baselines are as follows:

– CNN (Zeng et al. 2014). It is based on a convolutional neural network that uses FastText (Mikolov et al. 2018) pre-trained word embedding vectors of 300-dimension, three 1D convolutional layers of different window sizes (3, 4, and 5 respectively). Each layer is followed by a max-pooling layer. The output layer is composed of a fully connected layer followed by a softmax classifier. The results reported here were obtained by using a dropout of 50% and optimized using the Adam optimizer with a learning rate of 10^{-3}.

– Attention-BiLSTM (Zhou et al. 2016). As for CNN, FastText embeddings are used to initialize the embedding layer. The LSTM layer uses two LSTM hidden layers to extract high-level features while taking into account left and right contexts. The attention layer merges the word-level features extracted by LSTM layers into a sentence level vector by multiplying them by a weight vector calculated from the outputs of the LSTM layers. The final output layer is used for relation classification. During experiments, best results have been obtained by using 100 hidden units, an embedding dropout rate of 70%, a final layer dropout rate of 70% and the Adam optimizer with a learning rate of 1.

– R-mBERT (Wu and He 2019). This is an adaptation of pre-trained BERT multilingual language models for RE as suggested by (Wu and He 2019) in order to take into account entities representation. The models are fine-tuned on our dataset for 5 *epochs* to provide contextualized sentence and entity representations that are concatenated and fed into a relation classifier. We use the same hyperparameters of the original paper.

In addition to these baselines, we newly propose the two models below to deal with the specificities of French data:

–R-FlauBERT and **R-CamemBERT**. Following the adaptation of BERT for relation classification (Wu and He 2019), we adapt the French pre-trained language models CamemBERT (Martin et al. 2019) and FlauBERT(Le et al. 2019) by adding a classification layer on top of the pre-trained language model, to which the contextualized representations of the input sentence and target entities are given to predict the relation type. The models are fine-tuned on our dataset for 5 *epochs*.

All baseline models perform the main task of multi-class business relation classification including OTHERS (i.e., T_{all}).

Results

Both the baselines and the multitask models were evaluated on the BizRel test set. The results are reported in Table 3 in terms of macro precision, recall, and F1-score.

We can observe that the baseline models based on pre-trained language models are more effective than the ones based on static embeddings (CNN and Bi-LSTM), with the multilingual model R-mBERT giving the best F1 score result among them. The multitask models achieve good results outperforming the best baseline R-mBERT (+ 1.3% in F1). From the obtained results, two interesting findings can

[3]Link to mBert

Model	P	R	F1
CNN (Zeng et al. 2014)	66.8	51.3	57.0
Bi-LSTM (Zhou et al. 2016)	56.6	55.0	53.3
R-mBERT (Wu and He 2019)	71.2	64.1	67.1
R-CamemBERT	74.6	53.8	59.5
R-FlauBERT	**77.2**	59.7	66.3
MT-BizBERT$^{all+bin}$	66.2	**67.1**	66.5
MT-BizBERT$^{all+biz}$	72.9	**64.8**	**68.0**
MT-BizBERT$^{all+bin+biz}$	71.4	**66.1**	**68.4**

Table 2: Our results. Best scores are underlined while bold ones are those that outperform the best baseline.

Relation	P	R	F
Invest.	65.8	53.2	58.8
Compet.	67.7	71.9	69.7
Cooperat.	63.8	68.2	65.9
Legal.	83.3	62.5	71.4
Sale.	61.1	55.0	57.9
Others	86.9	86.1	86.5

Table 3: Best performing model results per relation type.

be drawn. First, considering T_{bin} as an auxiliary task in the multitask model could improve the model recall, scoring the highest value (+ 3% in R over baseline), yet low precision. Second, when the auxiliary task is T_{biz}, the model achieves a better precision compared to T_{bin}, however, the recall is still low. Combining the two auxiliary tasks $T_{bin} + T_{biz}$ with the main task offers a good compromise between precision and recall, achieving therefore the best F1-score (+ 1.3% over the baselines).

A closer look into results per class for our best model (cf. Table 3) shows that the relation types with the best F1-score are the ones with more training data (*Competition*, *Cooperation*). The relation type *Legal Proceeding* scores a high F1-score, which can be due to the similarity and few variations of relation instance patterns because of the few examples we have.

Error Analysis

We performed a detailed error analysis on the best performing model in order to gain insights into the main shortcomings of the current approach. We can notice two main sources of errors.

The first one concerns sentences containing more than one relation between different entity pairs, as in (8). In this example, only the relation linking the two EO underlined has to be identified. Our best model predicts COOPERATION (EO_2,EO_3), whereas the ground-truth annotation is OTHERS (EO_2,EO_3). Note that a COOPERATION relation actually exists between EO_1 and EO_2 and between EO_1 and EO_3 .

(8) À ce jour, les partenaires TOP du [CIO]$_1$ qui seront associés aux JO de Paris 2024 sont Alibaba, Bridgestone, [Intel]$_2$, Omega, [Panasonic]$_3$, Toyota et Visa.
(To date, the TOP partners of the [IOC]$_1$ who will be associated with the Paris 2024 Olympic Games are Alibaba, Bridgestone, [Intel]$_2$, Omega, [Panasonic]$_3$, Toyota and Visa)

The second source of error arises from the use of generic lexical clues to express certain business relations, as in (9). The lexical clue "de" (of) is generally used to express the relation type *Investment* referring to a subsidiary link between two organizations in French language. However, in this example, it does not. Our model misclassified this sentence as *Investment*(EO_1,EO_2) whereas the ground-truth annotation is OTHERS (EO_1,EO_2).

(9) Si [Google]$_1$ est sorti de [Stanford]$_2$, il y a aussi des startups françaises connues qui sont nées au sein d'incubateurs des écoles.
(If [Google]$_1$ came out of [Stanford]$_2$, there are also well-known French startups that were born within school incubators)

Conclusion

We introduce the first French dataset for business relation extraction manually annotated of 5 business relations and one negative relation *Others* of an importance for end-user systems. We also propose a multitask model trained on this dataset, accounting for generic and specific features of this type of relations. The obtained results are very encouraging and can help advance this line of research.

References

Berns, J.; Bick, P.; Flugum, R.; and Houston, R. 2021. Do changes in MD&A section tone predict investment behavior? *Financial Review*.

Braun, D.; Faber, A.; Hernandez-Mendez, A.; and Matthes, F. 2018. Automatic Relation Extraction for Building Smart City Ecosystems using Dependency Parsing. In Basile, P.; Basile, V.; Croce, D.; Dell'Orletta, F.; and Guerini, M., eds., *Proceedings of the 2nd Workshop on Natural Language for Artificial Intelligence (NL4AI 2018) co-located with 17th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2018), Trento, Italy, November 22nd to 23rd, 2018*, volume 2244 of *CEUR Workshop Proceedings*, 29–39. CEUR-WS.org.

Cenikj, G.; Seljak, B. K.; and Eftimov, T. 2021. FoodChem: A food-chemical relation extraction model. *arXiv preprint arXiv:2110.02019*.

Chen, D.; Zou, Y.; Harimoto, K.; Bao, R.; Ren, X.; and Sun, X. 2019. Incorporating Fine-grained Events in Stock Movement Prediction. In *Proceedings of the Second Workshop on Economics and Natural Language Processing*, 31–40.

Cohen, J. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1): 37–46.

Collovini, S.; Gonçalves, P. N.; Cavalheiro, G.; Santos, J.; and Vieira, R. 2020. Relation Extraction for Competitive Intelligence. In *International Conference on Computational Processing of the Portuguese Language*, 249–258. Springer.

De Los Reyes, D.; Barcelos, A.; Vieira, R.; and Manssour, I. 2021. Related Named Entities Classification in the Economic-Financial Context.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.

Han, S.; Hao, X.; and Huang, H. 2018. An event-extraction approach for business analysis from online Chinese news. *Electronic Commerce Research and Applications*, 28: 244–260.

Hendrickx, I.; Kim, S. N.; Kozareva, Z.; Nakov, P.; Séaghdha, D. Ó.; Padó, S.; Pennacchiotti, M.; Romano, L.; and Szpakowicz, S. 2010. SemEval-2010 Task 8: Multi-Way Classification of Semantic Relations between Pairs of Nominals. In *Proceedings of the 5th International Workshop on Semantic Evaluation*, 33–38.

Hogenboom, F.; de Winter, M.; Frasincar, F.; and Kaymak, U. 2015. A news event-driven approach for the historical value at risk method. *Expert Systems with Applications*, 42(10): 4667–4675.

Jacobs, G.; and Hoste, V. 2021. SENTiVENT: enabling supervised information extraction of company-specific events in economic and financial news. *LANGUAGE RESOURCES AND EVALUATION*.

Jacobs, G.; Lefever, E.; and Hoste, V. 2018. Economic event detection in company-specific news text. In *Proceedings of the First Workshop on Economics and Natural Language Processing*, 1–10.

Khachatrian, H.; Nersisyan, L.; Hambardzumyan, K.; Galstyan, T.; Hakobyan, A.; Arakelyan, A.; Rzhetsky, A.; and Galstyan, A. 2019. BioRelEx 1.0: Biological Relation Extraction Benchmark. In *Proceedings of the 18th BioNLP Workshop and Shared Task*, 176–190.

Khaldi, H.; Benamara, F.; Abdaoui, A.; Aussenac-Gilles, N.; and Kang, E. 2021. Multilevel Entity-Informed Business Relation Extraction. In *International Conference on Applications of Natural Language to Information Systems*, 105–118. Springer.

Lau, R.; and Zhang, W. 2011. Semi-supervised statistical inference for business entities extraction and business relations discovery. In *SIGIR 2011 workshop, Beijing, China*, 41–46.

Le, H.; Vial, L.; Frej, J.; Se that gonne, V.; Coavoux, M.; Lecouteux, B.; Allauzen, A.; Crabbé, B.; Besacier, L.; and Schwab, D. 2019. FlauBERT: Unsupervised Language Model Pre-training for French.

Lefever, E.; and Hoste, V. 2016. A Classification-based Approach to Economic Event Detection in Dutch News Text. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, 330–335. Portorož, Slovenia: European Language Resources Association (ELRA).

Li, Z.; Lian, Y.; Ma, X.; Zhang, X.; and Li, C. 2020. Bio-semantic relation extraction with attention-based external knowledge reinforcement. *BMC Bioinformatics*, 21: 1–18.

Liang, X.; Cheng, D.; Yang, F.; Luo, Y.; Qian, W.; and Zhou, A. 2020. F-HMTC: Detecting Financial Events for Investment Decisions Based on Neural Hierarchical Multi-Label Text Classification. In *IJCAI*, 4490–4496.

Lyu, S.; Cheng, J.; Wu, X.; Cui, L.; Chen, H.; and Miao, C. 2022. Auxiliary Learning for Relation Extraction. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 6(1): 182–191.

Martin, L.; Muller, B.; Suárez, P. J. O.; Dupont, Y.; Romary, L.; de la Clergerie, É. V.; Seddah, D.; and Sagot, B. 2019. CamemBERT: a Tasty French Language Model.

Mikolov, T.; Grave, É.; Bojanowski, P.; Puhrsch, C.; and Joulin, A. 2018. Advances in Pre-Training Distributed Word Representations. In *Proceedings of LREC*.

Oberlechner, T.; and Hocking, S. 2004. Information sources, news, and rumors in financial markets: Insights into the foreign exchange market. *Journal of economic psychology*, 25(3): 407–424.

Qian, Y.; Deng, X.; Ye, Q.; Ma, B.; and Yuan, H. 2019. On detecting business event from the headlines and leads of massive online news articles. *Information Processing & Management*, 56(6): 102086.

Sewlal, R. 2004. Effectiveness of the Web as a competitive intelligence tool. *South African Journal of Information Management*, 6(1).

Usmani, S.; and Shamsi, J. A. 2021. News sensitive stock market prediction: literature review and suggestions. *PeerJ Computer Science*, 7: e490.

Uzuner, Ö.; South, B. R.; Shen, S.; and DuVall, S. L. 2011. 2010 i2b2/VA challenge on concepts, assertions, and relations in clinical text. *Journal of the American Medical Informatics Association*, 18(5): 552–556.

Wang, H.; Zhu, T.; Wang, M.; Zhang, G.; and Chen, W. 2021. A Prior Information Enhanced Extraction Framework for Document-level Financial Event Extraction. *Data Intelligence*, 1–12.

Wang, R.; Tang, D.; Duan, N.; Wei, Z.; Huang, X.; Ji, J.; Cao, G.; Jiang, D.; and Zhou, M. 2020. K-Adapter: Infusing Knowledge into Pre-Trained Models with Adapters. *CoRR*, abs/2002.01808.

Wang, W.; and Hu, W. 2020. Improving relation extraction by multi-task learning. In *Proceedings of the 2020 4th High Performance Computing and Cluster Technologies Conference & 2020 3rd International Conference on Big Data and Artificial Intelligence*, 152–157.

Wiegand, M.; Roth, B.; Lasarcyk, E.; Köser, S.; and Klakow, D. 2012. A Gold Standard for Relation Extraction in the Food Domain. In *Proceedings of LREC*.

Wu, S.; and He, Y. 2019. Enriching Pre-Trained Language Model with Entity Information for Relation Classification. In *Proceedings of ACM CIKM'19*, 2361–2364.

Xing, R.; Luo, J.; and Song, T. 2020. BioRel: towards large-scale biomedical relation extraction. *BMC bioinformatics*, 21(16): 1–13.

Xingyue, C.; Liping, N.; and Zhiwei, N. 2021. Financial Event Extraction Based on ELECTRA and Part-of-Speech. *Data Analysis and Knowledge Discovery*, 1.

Xu, K.; Liao, S. S.; Li, J.; and Song, Y. 2011. Mining comparative opinions from customer reviews for Competitive Intelligence. *Decision Support Systems*, 50(4): 743–754. Enterprise Risk and Security Management: Data, Text and Web Mining.

Yadav, S.; Ramesh, S.; Saha, S.; and Ekbal, A. 2020. Relation extraction from biomedical and clinical text: Unified multitask learning framework. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*.

Yamada, I.; Asai, A.; Shindo, H.; Takeda, H.; and Matsumoto, Y. 2020. Luke: deep contextualized entity representations with entity-aware self-attention. *arXiv preprint arXiv:2010.01057*.

Yamamoto, A.; Miyamura, Y.; Nakata, K.; and Okamoto, M. 2017. Company relation extraction from web news articles for analyzing industry structure. In *2017 IEEE 11th International Conference on Semantic Computing (ICSC)*, 89–92. IEEE.

Yan, C.; Fu, X.; Wu, W.; Lu, S.; and Wu, J. 2019. Neural Network Based Relation Extraction of Enterprises in Credit Risk Management. In *2019 IEEE International Conference on Big Data and Smart Computing (BigComp)*, 1–6. IEEE.

Zeng, D.; Liu, K.; Lai, S.; Zhou, G.; and Zhao, J. 2014. Relation Classification via Convolutional Deep Neural Network. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, 2335–2344. Dublin, Ireland: Dublin City University and Association for Computational Linguistics.

Zhang, Y.; Zhong, V.; Chen, D.; Angeli, G.; and Manning, C. D. 2017. Position-aware Attention and Supervised Data Improve Slot Filling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 35–45. Copenhagen, Denmark: Association for Computational Linguistics.

Zhao, J.; Jin, P.; and Liu, Y. 2010. Business relations in the web: Semantics and a case study. *Journal of Software*, 5(8): 826–833.

Zhou, K.; Luo, X.; Wang, H.; and Xu, R. 2019. Multi-task Learning for Relation Extraction. In *31st International Conference on Tools with Artificial Intelligence (ICTAI)*, 1480–1487. IEEE.

Zhou, P.; Shi, W.; Tian, J.; Qi, Z.; Li, B.; Hao, H.; and Xu, B. 2016. Attention-Based Bidirectional Long Short-Term Memory Networks for Relation Classification. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 207–212. Berlin, Germany: Association for Computational Linguistics.

Zhou, W.; and Chen, M. 2021. An Improved Baseline for Sentence-level Relation Extraction. *arXiv preprint arXiv:2102.01373*.

Zuo, Z.; Loster, M.; Krestel, R.; and Naumann, F. 2017. Uncovering Business Relationships: Context-sensitive Relationship Extraction for Difficult Relationship Types. In Leyer, M., ed., *Lernen, Wissen, Daten, Analysen (LWDA) Conference Proceedings, Rostock, Germany, September 11-13, 2017*, volume 1917 of *CEUR Workshop Proceedings*, 271. CEUR-WS.org.