

BUSTER: a “BUSINESS Transaction Entity Recognition” dataset

Andrea Zugarini
expert.ai, Siena, Italy
azugarini@expert.ai

Marco Ernandes
expert.ai, Siena, Italy
mernandes@expert.ai

Andrew Zamai
expert.ai, Siena, Italy
azamai@expert.ai

Leonardo Rigutini
expert.ai, Siena, Italy
lrigutini@expert.ai

Abstract

Albeit Natural Language Processing has seen major breakthroughs in the last few years, transferring such advances into real-world business cases can be challenging. One of the reasons resides in the displacement between popular benchmarks and actual data. Lack of supervision, unbalanced classes, noisy data and long documents often affect real problems in vertical domains such as finance, law and health. To support industry-oriented research, we present *BUSTER*, a BUSINESS Transaction Entity Recognition dataset. The dataset consists of 3779 manually annotated documents on financial transactions. We establish several baselines exploiting both general-purpose and domain-specific language models. The best performing model is also used to automatically annotate 6196 documents, which we release as an additional silver corpus to *BUSTER*.

1 Introduction

Natural Language Processing (NLP) is a field potentially beneficial to a broad span of language-intensive domains, such as law and health. Whilst lots of Financial data are tabular, there is also crucial information stored in reports, news, transaction agreements, etc.

The abrupt developments in NLP (Vaswani et al., 2017) are favouring its adoption as assistance tools for human experts in many tasks, ranging from Document Classification (Chalkidis et al., 2019) to Information Extraction (Alvarado et al., 2015; Loukas et al., 2022) and even Text Summarization (Bhattacharya et al., 2019). However, transferring the emerging technologies into industry applications can be non-trivial. Adapting Large Language Models (LLMs) to vertical domains usually requires fine-tuning on domain-specific annotated data. Labeling is often a time-consuming, expensive process, especially when experts in the field are involved. Recently, several

benchmarks and datasets have been constructed for law (Chalkidis et al., 2022), health (Li et al., 2016) and finance (Loukas et al., 2022).

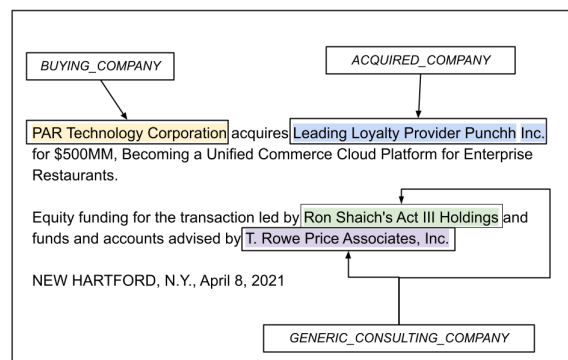


Figure 1: An annotated example extracted from *BUSTER*.

In this work, we support industry-oriented research community by presenting *BUSTER*: a BUSINESS Transaction Entity Recognition dataset. As the title suggests, *BUSTER* is an Entity Recognition (ER) benchmark that focus on the main actors involved in a business transaction. After collecting about ten thousands business transaction documents from EDGAR company acquisition reports, we constructed a dataset with 3779 manually annotated documents (the Gold corpus), from which we trained an LLM to automatically annotate the remaining 6196 documents (the Silver corpus). We analyze the properties of the proposed dataset and also evaluate the performance of some baselines. The dataset will be public and free to download as a benchmark for the NLP community.

The paper is organized as follows. First we review in Section 2 previous related works on Financial NER and document-level datasets. Then, we describe the data collection process and annotation methodologies in sections 3 and 4, respectively. A detailed description of *BUSTER* and its statistics follows in Section 5. In Section 6 we establish baselines with different LLMs. Finally, in Section 7 we

Tag Family	Tag Name	Description
Parties	<i>BUYING_COMPANY</i>	The company which is acquiring the target.
	<i>SELLING_COMPANY</i>	The company which is selling the target.
	<i>ACQUIRED_COMPANY</i>	The company target of the transaction.
Advisors	<i>LEGAL_CONSULTING_COMPANY</i>	A law firm providing advice on the transaction, such as: government regulation, litigation, anti-trust, structured finance, tax etc.
	<i>GENERIC_CONSULTING_COMPANY</i>	A general firm providing any other type of advice, such as: financial, accountability, due diligence, etc.
Generic_Info	<i>ANNUAL_REVENUES</i>	The past or present annual revenues of any company or asset involved in the transaction.

Table 1: Description of the tag-set defined in *BUSTER*.

draw our conclusions and outline possible future research directions.

2 Related works

Several document datasets in the financial domain have been proposed in the literature, but few of them are dedicated to the Entity Recognition (ER) task. Furthermore, these few are mainly intended for the standard Named Entity Recognition (NER) task, such as (Alvarado et al., 2015; Francis et al., 2019; Hampton et al., 2016; Kumar et al., 2016).

In Alvarado et al. (2015) is presented a corpus (FIN) of eight documents from SEC which were manually annotated using the standard four NER data type: person, organization, location and miscellaneous. Unlike that dataset, in *BUSTER* we decided to focus on Entities that are involved in a financial transaction. FiNER-139 (Loukas et al., 2022) instead consists in a large corpus of SEC documents annotated via gold XBLR tags, that includes a label set of 139 numerical entities on about 1.1M sentences. The tag attribution mostly depends on context rather than the token itself, as it is in *BUSTER*. Beside the completely different tag set, the main difference between *BUSTER* and Finer-139 is the fact that we release a document-level benchmark. Indeed, the detection of roles like the buyer company can require scopes wider than a single sentence. Moreover, documents come from files with heterogeneous layouts, extensions and structure, which can sometimes hinder the segmentation of the document into single sentences.

Outside the financial domain, a variety of document-level datasets for NER have been proposed. DocRED (Yao et al., 2019) is a NER and Relation Extraction (RE) corpus built from Wikidata

and Wikipedia short text passages, while BioCreative (Li et al., 2016) is a dataset for NER/RE on health domain. In (Quirk and Poon, 2016), the authors propose a dataset for NER in medical area.

3 Data Collection

Our goal was to create a highly business-oriented dataset to recognize relevant entities involved in financial transactions. Unlike standard NER tasks, we focused on the problem of entity-role recognition, where the goal is to identify a set of entities but only where they appear with specific roles in a context, such as companies involved in an acquisition or consultants assisting in an operation.

Target documents

To collect such documents, we exploited the EDGAR (Electronic Data Gathering, Analysis, and Retrieval system) service of the U.S. Securities and Exchange Commission (SEC)¹. The SEC’s mission is to maintain fair, orderly, and efficient markets. In particular, the organization aims to give transparency to business activities and provide investors with more security on the companies in which they invest, facilitating capital formation. For this purpose, domestic and foreign companies conducting business in the US are required to provide regular reports to the SEC through EDGAR. Reports are filed based on a list of forms that correspond to certain filing types. The EDGAR service provides more than 150 different form types (*filing type*)² and of these, the *Form 8K* type deserves particular attention.

¹<https://www.sec.gov/edgar/>

²https://en.wikipedia.org/wiki/SEC_filing

		<i>JPA</i>	<i>CPA₁</i>	<i>CPA₂</i>	<i>Cov₁</i>	<i>Cov₂</i>	κ
<i>Parties</i>	<i>BUYING_COMPANY</i>	0.6514	0.7445	0.8389	0.8749	0.7764	0.6810
	<i>SELLING_COMPANY</i>	0.5026	0.6362	0.7053	0.7900	0.7126	0.6383
	<i>ACQUIRED_COMPANY</i>	0.5611	0.6658	0.7811	0.8427	0.7184	0.6119
<i>Advisors</i>	<i>LEGAL_CONSULTING_COMPANY</i>	0.8913	0.9011	0.9880	0.9891	0.9022	0.9405
	<i>GENERIC_CONSULTING_COMPANY</i>	0.6624	0.7273	0.8814	0.9108	0.7516	0.7862
<i>Generic_Info</i>	<i>ANNUAL_REVENUES</i>	0.5781	0.6894	0.7817	0.7590	0.7000	0.7246
	MICRO OVERALL	0.6100	0.7107	0.8115	0.8583	0.7517	0.7257
	MACRO OVERALL	0.6448	0.7504	0.8148	0.8566	0.7882	0.7402

Table 2: The quality assessment results of the output of the annotation process.

An 8-K provides investors with timely notification of significant changes at listed companies such as acquisitions, bankruptcy, the resignation of directors, or changes in the fiscal year³. Optionally, but very frequently, the *Form 8K* includes a document called *Exhibit 99.1* (often abbreviated on *EX-99.1*). It consists of a disclosure document which summarizes all the details of the operation announced in the form and it is designed to provide investors with a complete and detailed view on the operation.

Crawling, filtering and processing

To collect the *EX-99.1* disclosure documents from EDGAR reporting company acquisitions, ownership changes and share purchase, we make use of the full index tool of the EDGAR site. Limiting to 2021, we downloaded about 120,000 *EX-99.1* disclosure documents in HTML format. After parsing, cleaning and removing any empty or too short documents, we selected the relevant documents using transaction-related keywords (acquisition, acquire, ownership, etc.) obtaining a final raw dataset of about 10,000 text files.

4 Annotation

For data labeling, we used a double-blind manual procedure. Specifically, we utilized two annotators (*ann₁* and *ann₂*), who were trained on the financial transactions topic and who were provided with a tag-set and specific guidelines to follow in the entity tagging procedure.

Tag-set

In designing the tag-set, we identified three families of tags: (a) *Parties* which groups tags used to identify the entities directly involved in the transaction;

(b) *Advisors* which groups tags identifying any external facilitator and advisor of the transaction and (c) *Generic_Info* which identifies tags reporting any information about the transaction. For each family, we defined a set of related tags. The tag-set is reported in Table 1.

Guidelines and General instructions

In order to improve annotation coherency, the schema definitions outlined in Table 1 were prepared as guidelines to the annotators. Moreover, the following general instructions were provided:

- **Annotate linguistically apparent instances only** – Tag only instances of entities where the class is linguistically evident. Do not tag a string just because you know that it is an instance of an entity: the context must make it obvious that it is an instance of such class.
- **Evaluate sentence context only** – Tag only instances of entities in which there is evidence within a sentence that the instance is of that entity. Each sentence should be evaluated for entities in isolation from the rest of the document context.

Annotation Procedure

To monitor the annotation procedure, the data set was divided into “sprints” which have been provided sequentially to the annotators. Each sprint consists of a pair of document batches that have been submitted independently to the two annotators. Additionally, we designed each sprint so that its two batches shared a certain percentage of documents. In this way, in each sprint, a portion of documents will be tagged by both annotators. Although this choice reduces the number of documents processed over time, it allows subsequent estimation of the annotation quality in each sprint.

³<https://www.sec.gov/investor/pubs/readan8k.pdf>

		Gold					Silver
		<i>fold</i> ₁	<i>fold</i> ₂	<i>fold</i> ₃	<i>fold</i> ₄	<i>fold</i> ₅	Total
	N. Docs	753	759	758	755	754	3779
	N. Tokens	685K	680K	687K	697K	688K	3437K
	N. Annotations	4119	4267	4100	4103	4163	20752
<i>Parties</i>	<i>BUYING_COMPANY</i>	1734	1800	1721	1707	1717	8679
	<i>SELLING_COMPANY</i>	460	447	456	426	439	2228
	<i>ACQUIRED_COMPANY</i>	1399	1473	1362	1430	1447	7111
	Total	3593	3720	3539	3563	3603	18018
<i>Advisors</i>	<i>LEGAL_CONSULTING_COMPANY</i>	142	132	152	146	153	721
	<i>GENERIC_CONSULTING_COMPANY</i>	256	267	261	248	256	1279
	Total	398	399	413	394	409	2013
<i>Generic_Info</i>	<i>ANNUAL_REVENUES</i>	128	148	148	146	151	721
	Total	128	148	148	146	151	696

Table 3: The statistics of the 5 folds Gold and Silver data.

We set the size of each sprint to 500 documents, 100 of which were shared between the two annotators (20%). The two annotators processed 8 sprints, thus obtaining 4000 annotated documents, 800 of which were labeled by both annotators. Finally, after removing documents without any labels, the resulting dataset was composed of 3779 labeled documents.

Validation

To evaluate the quality output of the annotation process, we exploited the shared set of documents that had been tagged by both annotators. In particular, indicating with L_1 and L_2 the two sets of annotations⁴ inserted respectively by the two annotators ann_1 and ann_2 in the shared documents, we calculated several standard indexes⁵:

- Joint Probability of Agreement, which measures the chance of having a match between the two annotators: $JPA = \frac{\#(L_1 \cap L_2)}{\#(L_1 \cup L_2)}$.
- Conditional Probability of Agreement of ann_k , which measures the naive probability that annotations inserted by an annotator k have a match with annotations entered by the other: $CPA_k = \frac{\#(L_1 \cap L_2)}{\#(L_k)}$, $k \in \{1, 2\}$.
- Coverage of ann_k , which measures the probability that a randomly selected annotation was entered by the annotator k : $Cov_k = \frac{\#(L_k)}{\#(L_1 \cup L_2)}$, $k \in \{1, 2\}$.
- Cohen’s kappa (κ), which extends the Joint Probability of Agreement taking into account

that agreement may occur by chance (Cohen, 1960): $\kappa = \frac{p_o - p_e}{1 - p_e}$ where $p_o = JPA$ is the observed agreement, $p_e = \frac{\#(L_1) \times \#(L_2)}{N^2}$ estimates the probability of a random agreement and $N = \#(L_1 \cup L_2)$ is the total number of inserted annotations.

The results are reported in the Table 2 and the values of Cohen’s kappa (κ) show a substantial agreement between the two evaluators (Landis and Koch, 1977).

Managing annotations in shared documents

In creating the final dataset, it was required to manage shared sets annotated by both annotators. Firstly, we accepted all non-overlapping annotations from both annotators. Secondly, we fixed overlapping, incoherent, annotations by involving a third annotator who manually assigned the correct label. Moreover, for pairs of overlapping annotations with boundaries $l_1 = [s_1, e_1]$ and $l_2 = [s_2, e_2]$, we merged them into a new annotation such that $l = [s, e] = [\min(s_1, s_2), \max(e_1, e_2)]$.

5 The BUSTER dataset

The final BUSTER dataset is composed of 3779 labeled documents. In Figure 1, we show an example of an annotated text passage inside a document. As explained, those documents were manually annotated and represent the “gold” BUSTER corpus. We randomly split the data into 5 folds to yield a statistically robust benchmark. Indeed, such division allows the use of a standard k-fold cross-validation approach.

The data set has been used as benchmark for 4 state-of-the art ER models (as described in Sec-

⁴Each ‘annotation’ refers to an entire annotated phrase.

⁵https://en.wikipedia.org/wiki/Inter-rater_reliability

Model	μ -Precision	μ -Recall	μ -F1	M-Precision	M-Recall	M-F1
BERT	61.16 \pm 1.65	67.42 \pm 2.72	64.06 \pm 0.90	55.12 \pm 1.75	66.60 \pm 2.79	59.80 \pm 1.23
SEC-BERT	66.76 \pm 0.74	74.18 \pm 1.99	70.28 \pm 0.90	70.30 \pm 0.96	78.10 \pm 1.82	73.98 \pm 1.14
RoBERTa	69.84 \pm 1.41	75.08 \pm 1.42	72.34 \pm 0.39	72.38 \pm 0.64	79.34 \pm 1.17	75.58 \pm 0.66
Longformer	69.28 \pm 2.71	73.40 \pm 1.31	71.24 \pm 1.34	70.02 \pm 3.27	77.34 \pm 1.49	73.30 \pm 2.25

Table 4: Micro (μ -) and macro (M-) scores of the four baseline models evaluated using 5-Fold Cross Validation.

tion 6) and the best performing model has been used to automatically annotate the remaining 6196 documents. The resulting annotated data was released as a “silver” extra corpus in *BUSTER* benchmark. The details of the 5 folds and of the silver extra corpus are reported in Table 3.

Statistics

Figure 2 shows the distribution of document lengths. The documents appear to have an aver-

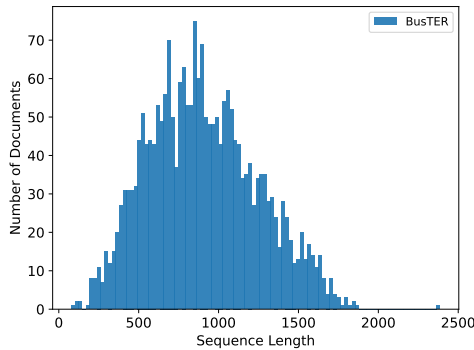


Figure 2: Sequence length distribution of *BUSTER* documents in terms of words.

age length of around 700 words and most of them fall into the 500-1000 range. Also, documents with more than 2000 words are extremely rare.

In figure 3, we report the distribution of the three tags families based on their position within the documents. We can observe how the tags belonging to the *Parties* family (in orange) are centered in the initial parts of the documents, while the remaining are distributed more uniformly and, in any case, located towards the second part of documents. However, no tags occurs beyond the 1500th word.

6 Experiments

To establish baselines, we performed several experiments using both generic and domain-specific language models.

Experimental Setup

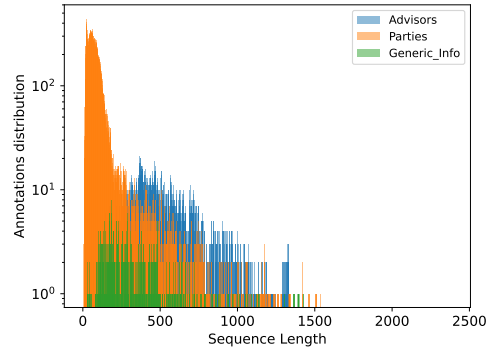


Figure 3: Distribution of tags families inside the documents.

In the experiments, we followed a 5-folds cross validation approach using the folds in Table 3 .

Metrics. We adopt traditional NER metrics for evaluation, i.e. micro and macro F1 scores, referred as μ -F1 and M-F1, respectively. True positives are counted in a strict sense, i.e. an entity is considered correctly predicted if and only if all of its constituent tokens are well identified, and no additional tokens belong to the entity.

Dealing with long documents. As shown in Figure 2, the vast majority of documents in *BUSTER* has more than 500 words, which typically exceeds the maximum sequence length that LLMs (e.g. BERT (Devlin et al., 2018)) can take in input. Truncation would cause a major drop of most of the document and a significant loss of information. Therefore, we split documents into contiguous chunks of text. Chunking is done such that no token is truncated at all and we fill each chunk sequence as much as possible. All the baselines are trained and tested on chunks with the exception of Longformers, since they are capable of processing longer sequences up to 4096 tokens.

Baseline Models

We considered several transformer-based models that report state-of-the-art performance in NLP. In particular, we have selected the following 4 models.

		Precision	Recall	F1
<i>Parties</i>	<i>BUYING_COMPANY</i>	74.06 ± 2.06	78.38 ± 1.47	76.12 ± 0.85
	<i>SELLING_COMPANY</i>	65.34 ± 2.35	75.04 ± 3.15	69.82 ± 0.77
	<i>ACQUIRED_COMPANY</i>	64.42 ± 1.11	70.38 ± 0.63	67.26 ± 0.38
<i>Advisors</i>	<i>LEGAL_CONSULTING_COMPANY</i>	84.86 ± 3.33	90.90 ± 2.33	87.72 ± 1.46
	<i>GENERIC_CONSULTING_COMPANY</i>	73.98 ± 1.97	77.98 ± 3.27	75.90 ± 2.04
<i>Generic_Info</i>	<i>ANNUAL_REVENUES</i>	61.88 ± 5.95	79.36 ± 4.66	69.30 ± 4.24

Table 5: Tag-wise precision, recall and F1-score values obtained with the RoBERTa baseline using 5-Fold Cross Validation.

BERT. BERT (Devlin et al., 2018) constitutes a standard baseline since it is one of the most popular LLMs nowadays.

RoBERTa. Similarly to BERT, RoBERTa (Liu et al., 2019) is a widely-used Language Model in the NLP community. The model is an optimized version of BERT and generally outperforms it.

SEC-BERT. We also consider a domain-specific LLM. We consider SEC-BERT (Loukas et al., 2022), a model pre-trained from scratch on EDGAR-CORPUS, a large collection of financial documents (Loukas et al., 2021).

Longformer. Longformer (Beltagy et al., 2020) is a transformer architecture equipped with self-attention mechanisms that scales linearly with the sequence length. Longformers were specifically designed to deal with long documents, hence they are a natural good candidate for processing *BUSTER*.

Results

The baselines’ performance are presented in Table 4. RoBERTa turned out to be the best performing model, with Longformer achieving similar levels of accuracy. BERT base, instead, underperformed with respect to the other baselines. However, when fine-tuning BERT on the financial domain (SEC-BERT) there is a clear F1 improvement.

Inspecting the scores of single entity tags obtained by the best model, i.e. RoBERTa (Table 5), we can observe that the *Advisors* family is generally well captured by the model. For *Parties* and *Generic_Info* families instead, the results are different. The model performs very well on *BUYING_COMPANY*, while *ACQUIRED_COMPANY*, *SELLING_COMPANY* and *ANNUAL_REVENUES* appear more complex to discriminate, especially in terms of precision. In our analysis, this depends on some structural characteristics of these entities. The first two tags (*ACQUIRED_COMPANY* and

SELLING_COMPANY) are strongly related to each other and often they are not easy to disambiguate even for human annotators, as confirmed by the quality assessment outlined in Table 2. The definition of *ANNUAL_REVENUES* instead, is very specific and detailed (Section 4) and this makes it hard to distinguish it from occurrences of other economic data present in the text, e.g. EBITDA. Finally, the inherent complexity inevitably increases the noise in the gold annotations, thus affecting the training of the model itself.

7 Conclusions and future works

In this work, we presented *BUSTER*, an Entity Recognition (ER) benchmark for business transaction-related entities. It consists of a corpus of 3779 manually annotated documents on financial transaction (the Gold data) which has been randomly divided into 5 folds, plus an additional set of 6196 automatically annotated documents (the Silver data) that were created from the fine-tuned RoBERTa model. We plan to make it publicly available and free-to-download and we believe it could become a reference benchmark in the field of Entity Recognition, in particular for the financial domain.

In the future, we intend to work in two directions. On one side, we plan to increase the amount of manually labeled data and to extend the labels set with more transaction-related tags. On the other hand, we aim to introduce some specific types of relations between entities in order to extend the dataset to Relational Extraction.

Acknowledgements

A huge thank you to Bianca Vallarano and Stefano Genua who participated as annotators. Thanks to Daniela Baiamonte who supported us in the production of the guidelines and in the validation of the annotation process. Thanks to Paolo Lombardi who prepared the scripts to download and process

the documents from EDGAR.

This work was supported by the IBRIDAI project, a project financed by the Regional Operational Program “FESR 2014-2020” of Emilia Romagna (Italy), resolution of the Regional Council n. 863/2021.

References

- Julio Cesar Salinas Alvarado, Karin Verspoor, and Timothy Baldwin. 2015. Domain adaption of named entity recognition to support credit risk assessment. In *Proceedings of the Australasian Language Technology Association Workshop 2015*, pages 84–90.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Paheli Bhattacharya, Kaustubh Hiware, Subham Rajgaria, Nilay Pochhi, Kripabandhu Ghosh, and Saptarshi Ghosh. 2019. A comparative study of summarization algorithms applied to legal case judgments. In *Advances in Information Retrieval: 41st European Conference on IR Research, ECIR 2019, Cologne, Germany, April 14–18, 2019, Proceedings, Part I 41*, pages 413–428. Springer.
- Ilias Chalkidis, Emmanouil Fergadiotis, Prodromos Malakasiotis, and Ion Androutsopoulos. 2019. [Large-scale multi-label text classification on EU legislation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6314–6322, Florence, Italy. Association for Computational Linguistics.
- Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Katz, and Nikolaos Aletras. 2022. [LexGLUE: A benchmark dataset for legal language understanding in English](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4310–4330, Dublin, Ireland. Association for Computational Linguistics.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20:37 – 46.
- 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*.
- Sumam Francis, Jordy Van Landeghem, and Marie-Francine Moens. 2019. Transfer learning for named entity recognition in financial and biomedical documents. *Information*, 10(8):248.
- Peter Hampton, Hui Wang, William Blackburn, and Zhiwei Lin. 2016. Automated sequence tagging: Applications in financial hybrid systems. In *Research and Development in Intelligent Systems XXXIII: Incorporating Applications and Innovations in Intelligent Systems XXIV 33*, pages 295–306. Springer.
- Aman Kumar, Hassan Alam, Tina Werner, and Manan Vyas. 2016. [Experiments in candidate phrase selection for financial named entity extraction - a demo](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: System Demonstrations*, pages 45–48, Osaka, Japan. The COLING 2016 Organizing Committee.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Jiao Li, Yueping Sun, Robin J Johnson, Daniela Sciaky, Chih-Hsuan Wei, Robert Leaman, Allan Peter Davis, Carolyn J Mattingly, Thomas C Wieggers, and Zhiyong Lu. 2016. Biocreative v cdr task corpus: a resource for chemical disease relation extraction. *Database*, 2016.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Lefteris Loukas, Manos Fergadiotis, Ion Androutsopoulos, and Prodromos Malakasiotis. 2021. Edgar-corpus: Billions of tokens make the world go round. *arXiv preprint arXiv:2109.14394*.
- Lefteris Loukas, Manos Fergadiotis, Ilias Chalkidis, Eirini Spyropoulou, Prodromos Malakasiotis, Ion Androutsopoulos, and Georgios Paliouras. 2022. [Finer: Financial numeric entity recognition for xbrl tagging](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4419–4431.
- Chris Quirk and Hoifung Poon. 2016. Distant supervision for relation extraction beyond the sentence boundary. *arXiv preprint arXiv:1609.04873*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Yuan Yao, Deming Ye, Peng Li, Xu Han, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, Lixin Huang, Jie Zhou, and Maosong Sun. 2019. Docred: A large-scale document-level relation extraction dataset. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 764–777.