

KPI-EDGAR: A Novel Dataset and Accompanying Metric for Relation Extraction from Financial Documents

Tobias Deußer^{*†}, Syed Musharraf Ali^{†‡}, Lars Hillebrand^{*†}, Desiana Nurchalifah^{†‡}, Basil Jacob[†],

Christian Bauckhage^{*†}, Rafet Sifa[†]

^{*} University of Bonn, Bonn, Germany

[†] Fraunhofer IAIS, Sankt Augustin, Germany

[‡] Hochschule Bonn-Rhein-Sieg, Sankt Augustin, Germany

tobias.deusser@iais.fraunhofer.de

ORCID iD: 0000-0003-4685-0847

Abstract—We introduce KPI-EDGAR, a novel dataset for Joint Named Entity Recognition and Relation Extraction building on financial reports uploaded to the Electronic Data Gathering, Analysis, and Retrieval (EDGAR) system, where the main objective is to extract Key Performance Indicators (KPIs) from financial documents and link them to their numerical values and other attributes. We further provide four accompanying baselines for benchmarking potential future research. Additionally, we propose a new way of measuring the success of said extraction process by incorporating a word-level weighting scheme into the conventional F_1 score to better model the inherently fuzzy borders of the entity pairs of a relation in this domain.

Index Terms—text mining, natural language processing, relation extraction, named entity recognition, machine learning

I. INTRODUCTION

Financial documents like income statements, business plans, initial public offering prospectuses, or whole annual financial reports hold a large amount of information for the reader willing to spend a considerable time reading through and analyzing them. Financial auditors are employed to check their validity and financial analysts utilize them to evaluate companies. Thus, creating an algorithm that automatically extracts information is likely to give the user a competitive edge by saving them extensive time usually spend on in-depth reading.

Arguably, one of the most important pieces of information contained in such documents are key performance indicators (KPIs), which are defined as “vital navigation instruments used by managers to understand whether their business is on a successful voyage or whether it is veering off the prosperous path” [1], and are mostly quantitative measures. Well-known examples are revenue, profit, and loss.

In a previous paper, [2] introduced KPI-BERT, a system which employed novel methods of named entity recognition

This research will be published in the 2022 proceedings of the IEEE International Conference on Machine Learning and Applications (ICMLA) and has been funded by the Federal Ministry of Education and Research of Germany and the state of North-Rhine Westphalia as part of the Lamarr-Institute for Machine Learning and Artificial Intelligence, LAMARR22B.

(NER) and relation extraction (RE) to extract and link key performance indicators of companies from real-world German financial documents. This model was trained on undisclosed annotations in public financial documents owned by a major auditing company, and thus, unavailable to the research community.

To alleviate this flaw and to facilitate reproducibility as well as further research, we introduce KPI-EDGAR, a dataset based on publicly available data and manually annotated by us, with the main difference being that the language changed from German to English. It was collected from the Electronic Data Gathering, Analysis, and Retrieval system¹, usually shortened to EDGAR, which is maintained by the U.S. Securities and Exchange Commission and is made up of annual comprehensive financial reports, so-called *10-K* reports. It will be published with the release of this paper.

Furthermore, we present an adjusted F_1 metric, that considers an entity and thus relation not as a unified, single object, but as a collection of words that should be evaluated as such. This is based on the observation that a relation extraction algorithm is often partly correct and missing only a minor part of the whole relation, regularly a rather unimportant word. The conventional F_1 metric would still assign zero accuracy to this prediction, hence not reflecting the actual performance. These prediction errors probably stem from the fact that the word boundaries of a particular KPI might not be clearly defined and even financial auditors can be unsure where the exact confines of a KPI lie. To mitigate this issue, we incorporate a weighting scheme into the F_1 metric to better capture these cases.

To sum up, our contributions are:

- We introduce a novel dataset, named KPI-EDGAR, for joint named entity recognition and relation extraction in the financial text domain based on actual corporate reports submitted to the EDGAR database. In contrast to

¹At the time of writing, the EDGAR database can be accessed at this link: sec.gov/edgar

other resources scraped from EDGAR (see [3], [4], [5], and [6]), we provide manual annotations along with the corpus to allow for joint named entity recognition and relation extraction.

- We propose a new metric for measuring the performance of the relation extraction task to better capture when predictions are partially correct.

The rest of this paper is structured as follows. We first review related work. In the next section, we describe our novel dataset, KPI-EDGAR, in detail and how it was gathered and pre-processed. Section IV defines the adjusted metric that incorporates a weighting scheme and introduces and specifies the four baselines we provide to researchers to benchmark their approaches. Thereafter, we evaluate these four baselines on KPI-EDGAR and report the conventional and adjusted F_1 metric as well as a few example predictions to highlight strengths and weaknesses of the metric. We close this paper with concluding remarks and take an outlook into conceivable future work.

II. RELATED WORK

Similar to our topic, many recent studies have focused on named entity recognition ([7], [8], [9], [10]), relation extraction ([11], [12], [13]), and the joint combination of both ([14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24]). Furthermore, [25] as well as [26] have shown that hierarchically using both tasks in a single pipeline can significantly improve the performance thus indicating that information from one task can be exploited by the other. Additionally, methods described in recent publications are usually based on modern pre-trained natural language models ([27], [28], [29], [30]) and these methods have produced state-of-the-art results as expected.

For extracting information from financial reports, [31] provided pre-available accounting variables to a multilayer perceptron as input to build interpretable structures similar to accounting ratios. In comparison, [32] automatically extracted these variables and developed a NER model using a rule-based approach. [33] developed a hybrid AI tool that extracts KPIs by first detecting paragraphs and tables with computer vision techniques and then leverage supervised decision tree-based classifiers to accomplish the final extraction process. [2] applied a pre-trained language model combined with conditional label masking to extract key performance indicators from German financial reports. [34] used a joint entity and relation extraction approach to cross-check different financial formulas in Chinese documents.

Regarding the herein proposed metric, [35] proposed a similar setup, but for Named Entity Recognition. Furthermore, [36] suggested dividing entities into buckets and evaluating the performance on each of these buckets.

Due to EDGAR’s popularity, many researchers ([3], [4], [6]) have developed methods to extract data from EDGAR and used it in their own research. [5] even released a comprehensive corpus comprising annual reports from all the publicly traded companies in the US spanning a period of more than

25 years and accompanied it with a word2vec [37] model titled EDGAR-W2V, which we will also use as a baseline in our experiments. Furthermore, [38], [39], and [40] have applied machine learning methods to EDGAR data and have provided useful results that can be used in real-world financial applications.

III. DATA

Our overall dataset is comprised of 81 manually annotated 10-K reports, with a total of 1355 sentences holding 4522 entities and 3841 relations. A 10-K report is a comprehensive report filed annually by a publicly-traded company about its financial performance. The dataset was scraped from the EDGAR (Electronic Data Gathering, Analysis, and Retrieval system) database, which is maintained by the U.S. Securities and Exchange Commission and where all in the U.S. publicly listed companies have to submit a variety of financial reports.

As a first pre-processing step we tokenize the reports on a sentence level and then on a word level. Thereafter, we identify monetary values, the scale they are in (e.g. billion) and which currency unit is attached to it (usually U.S. Dollars) through rule-based string matching heuristics. This allows us to filter our reports for sentences containing such entities, as the focus of our dataset is to extract numeric monetary values and their corresponding key performance indicator. Given these sentences, we annotate word-level entities and relations.

These annotations were generated by a team of four annotators, led by an auditing expert. The various entities and their annotations guideline, shown in Table I, as well as the allowed relation matrix, which restricts which entities can be linked together, shown in Table II, were defined in consultation with a wider team of auditors, who provided valuable feedback during the earlier stages of our project. After completing the annotations, the aforementioned senior auditing expert re-annotated 41 documents previously annotated by the other three annotators. The resulting inter-annotator agreement in terms of *Cohen’s Kappa*² [42] on the word level is 0.7037. We further discuss this value and further insights regarding this inter-annotator agreement in Section V and the accompanying Table V, which documents various *Cohen’s Kappa* scores.

We randomly split the pre-processed dataset on the document level into a training, validation, and test set, encompassing 969, 146, and 240 sentences each.

The KPI-EDGAR dataset and code to reproduce the results herein are available at GitHub³.

IV. METHODOLOGY

In this section, we shortly review the key points of KPI-BERT, a model tailored for key performance indicator extraction, introduced and illustrated in much greater detail in [2] Thereafter, one span-level approach is briefly touched

²The range of Cohen’s Kappa lies between 1 (complete agreement of annotators) and -1 (complete disagreement of annotators), with 0 being no agreement between annotators or only what is to be expected by chance. See [41] for an in-depth explanation on how to interpret *Cohen’s Kappa*.

³See github.com/tobideusser/kpi-edgar

Entity	Support	Description / Annotation Guideline
kpi	1341	Key Performance Indicators expressible in numerical and monetary value, e.g. revenue or net sales.
cy	1211	Current Year monetary value of a KPI .
py	619	Prior Year monetary value of a KPI.
py1	307	2 Year Past Value of a KPI
increase	35	Increase of a KPI from the previous year to the current year.
increase-py	15	Analogous to increase, but from py1 to py.
decrease	23	Decrease of a KPI from the previous year to the current year.
decrease-py	11	Analogous to decrease, but from py1 to py.
thereof	507	Represents a <i>subordinate</i> KPI, i.e. if a KPI is part of another, broader KPI.
attr	272	Attribute that further describes a KPI.
kpi-coref	11	A co-reference to a KPI mentioned in a previous sentence. Captures tokens that are similar to other entities, but are explicitly not one of them, e.g. when the writer of the report forecasts next year’s revenue.
false-positive	170	

TABLE I: Description and support of all entity types in the complete dataset, excluding the *none* type.

upon, namely the SpERT model introduced by [16]. We then provide two further baselines building on EDGAR–W2V [5] and GloVe [43], leveraging a similar setup like [2]. These four models will be the baselines we provide to other researchers to benchmark their model against on KPI-EDGAR.

Finally, we describe our proposed adjustments to the prevalent F_1 metric to better capture the actual performance of machine learning models on relation extraction from financial documents.

A. KPI-BERT

The KPI-BERT [2] model has three main building blocks: A BERT-based Sentence Encoder, a Named Entity Recognition (NER) Decoder, and a Relation Extraction (RE) Decoder.

Given a WordPiece [44] tokenized input sentence we leverage a pre-trained BERT [27] model to obtain our encoded token embedding. Thereafter, we apply a pooling function, which creates word representations by combining their individual subword embeddings. Both the more basic pooling functions `max` and `mean` as well as the specifically by [2] introduced trainable recurrent neural network (RNN) pooling mechanism building on a bidirectional gated recurrent unit (GRU) [45] are evaluated in this context.

The NER decoder module classifies named entities within the sentence building on the IOBES annotation scheme. IOBES tagging functions by prepending all entity classes with the prefixes *I-* (inside), *B-* (begin), *E-* (end) and *S-* (single), or *O* (outside) for tokens not belonging to the entity. Therefore, the number of possible IOBES entity tags is

$$|\mathcal{E}_{\text{IOBES}}| = 4(|\mathcal{E}| - 1) + 1, \quad (1)$$

where $|\mathcal{E}|$ is the number of entities, including the *none* entity, as described in Table I. The inherently sequential nature of this task leads us to leverage a GRU in combination with conditional label masking to sequentially predict IOBES tags considering the previous predictions.

To briefly elaborate, an embedding matrix $\mathbf{W}_{\text{label}} \in \mathbb{R}^{|\mathcal{E}_{\text{IOBES}}| \times u}$ holding learnable u -dimensional embeddings of

	kpi	cy	py	py1	increase	increase-py
kpi	-	1:1	1:1	1:1	1:1	1:1
cy	1:1	-	-	-	-	-
py	1:1	-	-	-	-	-
py1	1:1	-	-	-	-	-
increase	1:1	-	-	-	-	-
increase-py	1:1	-	-	-	-	-
decrease	1:1	-	-	-	-	-
decrease-py	1:1	-	-	-	-	-
thereof	n:1	1:1	1:1	1:1	1:1	1:1
attr	n:1	-	-	-	-	-
kpi-coref	-	1:1	1:1	1:1	1:1	1:1
false-positive	-	-	-	-	-	-

(a)

	decrease	decrease-py	thereof	attr	kpi-coref	false-positive
kpi	1:1	1:1	1:n	1:n	-	-
cy	-	-	1:1	-	1:1	-
py	-	-	1:1	-	1:1	-
py1	-	-	1:1	-	1:1	-
increase	-	-	1:1	-	1:1	-
increase-py	-	-	1:1	-	1:1	-
decrease	-	-	1:1	-	1:1	-
decrease-py	-	-	1:1	-	1:1	-
thereof	1:1	1:1	-	-	n:1	-
attr	-	-	-	-	n:1	-
kpi-coref	1:1	1:1	1:n	1:n	-	-
false-positive	-	-	-	-	-	-

(b)

TABLE II: Comprehensive overview of all allowed relations and their uniqueness. “1:1”: One entity of type 1 can only be linked to one entity of type 2, “1:n”: One entity of type 1 can be linked to many entities of type 2. “-”: No relation possible.

all IOBES entity types is defined and is concatenated to pooled BERT word embeddings (e_1, e_2, \dots, e_m) , yielding the decoding input representation of any word j ,

$$\mathbf{z}_j = [e_j; \mathbf{w}_{j-1}^{\text{label}}], \quad (2)$$

where $\mathbf{w}_{j-1}^{\text{label}} \in \mathbb{R}^u$ is the embedding of the formerly predicted IOBES tag and $\mathbf{w}_0^{\text{label}}$ is defined as the *O* (outside) label embedding.

Afterwards, \mathbf{z}_j and the previous hidden state \mathbf{h}_{j-1} is passed to a GRU, resulting in

$$\mathbf{h}_j = \text{GRU}(\mathbf{z}_j, \mathbf{h}_{j-1}). \quad (3)$$

We linearly transform \mathbf{h}_j combined with masking out impossible tag predictions, e.g., an *I* (inside) token without a previous *B* (begin) token, and applying softmax to arrive at the IOBES tag predictions for word j :

$$\hat{\mathbf{y}}_j = \text{softmax}(\text{mask}(\mathbf{W}_{\text{seq}}\mathbf{h}_j + \mathbf{b}_{\text{seq}})). \quad (4)$$

Finally, the RE Decoder takes the binary decision of whether two entities match or not. The entity sampling for this process is constrained by only passing allowed pairs, as defined in Table II, to the decoder. The inputs to this decoder are the embeddings, defined for each entity s as

$$e(s) := [\text{pool}(e_j, e_{j+1}, \dots, e_{j+k-1}); \mathbf{w}_k^{\text{width}}], \quad (5)$$

where $e(s) \in \mathbb{R}^{d+v}$ and $\mathbf{w}_k^{\text{width}}$ is a span size embedding, which is taken from an embedding matrix $\mathbf{W}_{\text{width}} \in \mathbb{R}^{l \times v}$ holding fixed-size learned embeddings of dimensionality v .

B. Span-Level Baseline

At the core of this model, we utilize the same architecture as shown in [16], where they titled the model SpERT, which we will use for this approach as well. Passing a byte-pair encoded input sentence of length n through BERT [27], we obtain $n+1$ token embeddings, $(c, e_1, e_2, \dots, e_n)$, where c represents the context embedding for the whole sentence.

In contrast to the approaches in sections IV-A and IV-C, we classify entire entity spans instead of individual tokens. Thus, we create all possible token subsequences from the encoded BERT output, up to a maximum span length l , set to $l = 10$ to keep true to the original configuration of [16]. For example, given the input (alea iacta est) we form the spans (alea), (iacta), (est), (alea, iacta), (iacta, est), (alea, iacta, est). Thereafter, each span is classified into entity types, and each possible entity pair receives a prediction whether or not it is linked together.

Different from [16], we additionally filter overlapping entity spans by removing the ones with a lower classification score, as overlapping spans are simply not possible in our application of extracting KPI’s from financial documents. To illustrate, take the example above and let us assume that the span (alea) and (alea, iacta) are both classified as an entity, the first with a score of 0.75 and the latter with 0.5. In this approach, we would only keep the entity (alea).

C. Further Token-Level Baselines

As defined in the methodology section IV, the introduced joint NER and RE model consists of three building blocks: a sentence encoder, NER decoder, and RE decoder. For generating further baselines, the sentence encoder of the described model is replaced by EDGAR-W2V [5] and GloVe [43]. The rest of the architecture is left unchanged. For these baselines, the sentences are tokenized into words and not WordPiece tokens.

The word embeddings are extracted from the available pre-trained embeddings uploaded by their respective authors. For those words which are not present in the vocabulary, a random vector is used as a substitute.

D. Adjusted F_1 Metric

To better capture the actual performance of predictions during the relation extraction from financial documents, we introduce an adjusted F_1 metric that allows for relations to be partially correct.

Given the following example sentence and ground truth relations from a financial document,

“In 2021 and 2020 the total net revenue was \$100 million and \$80 million, respectively.”

total net revenue – 100, total net revenue – 80,

we found that predictions often omit a part of non-numerical entities like the KPI entity. In this example, predicting *net revenue* instead of *total net revenue* leads to a strict, conventional

Model	Relation F_1 in %	Adjusted Relation F_1 in %
KPI-BERT	22.68	43.76
SpERT	20.95	40.04
EDGAR-W2V	6.13	19.71
GloVe	5.11	17.18

TABLE III: Model performances⁴, measured in the conventional as well as adjusted (see subsection IV) F_1 score, of the models described in sections IV-A, IV-C, and IV-B on our novel dataset KPI-EDGAR.

F_1 score of zero for both relations, even though the essential information was correctly extracted.

First, we define o_i as the overlap/intersection of an entity prediction i and its ground truth:

$$o_i := |e_{i,\text{pred}} \cap e_{i,\text{gt}}|, \quad (6)$$

where $e_{i,\text{pred}}$ and $e_{i,\text{gt}}$ is the set of all token identifiers for the entity prediction and ground truth, respectively, and the operation $|\cdot|$ calculates the size of a given set.

The *true positives* (tp), *false negatives* (fn), and *false positives* (fp) of a relation r between entity i and j are then calculated by:

$$\text{tp}_r = \frac{1}{2} \left(\frac{o_i}{n_{i,\text{gt}}} + \frac{o_j}{n_{j,\text{gt}}} \right) \quad (7)$$

$$\text{fn}_r = 1 - \text{tp}_r \quad (8)$$

$$\text{fp}_r = \frac{1}{2} \left(\frac{n_{i,\text{pred}} - o_i}{n_{i,\text{pred}}} + \frac{n_{j,\text{pred}} - o_j}{n_{j,\text{pred}}} \right), \quad (9)$$

where $n_{i,\cdot} := |e_{i,\cdot}|$. With this, we can calculate precision, recall, and the F_1 score of the relation r in the conventional way:

$$\text{precision}_r = \frac{\text{tp}_r}{\text{tp}_r + \text{fp}_r} \quad (10)$$

$$\text{recall}_r = \frac{\text{tp}_r}{\text{tp}_r + \text{fn}_r} \quad (11)$$

$$F_{1,r} = 2 \cdot \frac{\text{precision}_r \cdot \text{recall}_r}{\text{precision}_r + \text{recall}_r}. \quad (12)$$

V. EXPERIMENTS

Herein, we briefly evaluate the unmodified KPI-BERT [2] (see subsection IV-A), the same structure but with word embeddings from EDGAR-W2V [5] and GloVe [43] instead of embeddings from BERT [27] (see subsection IV-C for these two), and a competing span-based approach (see subsection IV-B), as introduced by [16] and named SpERT, on our novel dataset. The results should be seen as a baseline on which further research can be benchmarked. We also look into a few

⁴Please note that no extensive grid search and finetuning was performed in this context, as the aim of this work is to introduce a new dataset and a better suited metric and not to find the model best suited for this task. Accordingly, these results should be seen as baselines for further research.

	Sentence with predicted Entities	Relations
(a)	Entity boundaries are debatable, leading to arguably correct predictions but a conventional F_1 score of zero	
1	(a) [[Unrealized gains] totaled] \$[96] million in 2020, \$[88] million in 2019 and \$[73] million in 2018 [...].	kpi - cy kpi - cy kpi - py kpi - py kpi - py1 kpi - py1
2	As of December 31, 2020 and 2019, the Company's [Medicare Part D [receivables]] amounted to \$[2.9] billion and \$[2.3] billion, respectively.	kpi - cy kpi - cy kpi - py kpi - py
(b)	Difference in prediction and ground truth, however both are viable options. Nevertheless, the conventional and adjusted F_1 score will assign 0 here.	
3	As a result, we recognized \$[50] million of [costs] primarily related to [employee termination expenses and losses] from closing certain stores impacting both segments	kpi - cy kpi - cy

TABLE IV: Several example sentences from the test set of KPI-EDGAR with joint named entity recognition and relation extraction results. Green, Blue and Orange represent “true positive”, “false positive”, and “false negative” entity and relation classifications, respectively. The predictions were generated by KPI-BERT [2].

prediction examples and study the strength and weaknesses of our proposed adjustment to the F_1 score.

Table III shows the conventional as well as the adjusted F_1 score of KPI-BERT, SpERT, and baselines with EDGAR-W2V and GloVe as an encoder on KPI-EDGAR. A first and obvious observation is that approaches without a transformer as an encoder are severely lacking in performance compared to the ones utilizing such an encoder. Given the sequential and context based nature of detecting and extracting key performance indicators and their values, this result is what we believe can be expected from attention based models as their inherent ability is to model such relations.

Looking at the adjusted F_1 score, it is, as expected, significantly higher than the conventional score, which begs the question: does it do a better job in actually measuring the success of the prediction?

To answer this question, Table IV illustrates a few examples taken directly from the test set of KPI-EDGAR. Subtable (a) highlights sentences where our adjusted F_1 score enables the metric to reflect the actual performance of a model. These examples also reveal a critical issue annotators and thus, machine learning models have to face when extracting KPIs from financial documents: Where are the exact boundaries of an entity? What are relevant pieces of information that still belong to said entity and what is less important? Often, two expert auditors will have a differing opinion on where to put these boundaries. This inherently fuzzy and noisy process can also be observed when looking at Table V, where various Cohen’s Kappa scores are reported. As shown there, annotators agree strongly on where numeric entities like *cy* or *py* lie, but the location and especially borders of non-numeric entities like *kpi* or *thereof* are much more up to debate.

Therefore, we believe that our proposed weighted F_1 will

Type	Cohen’s Kappa
All words	0.7037
Only entities	0.4885
Entity: kpi	0.0822
Entity: cy	0.5972
Entity: py	0.6657
Entity: py1	0.6095
Entity: increase	0.7419
Entity: increase-py	0.5556
Entity: decrease	0.7713
Entity: decrease-py	-0.1538
Entity: thereof	0.3053
Entity: attr	-0.1076
Entity: kpi-coref	-0.2592
Entity: false-positive	-0.7587

TABLE V: Cohen’s Kappa [42] scores of various input series. *All words* includes all word tokens. *Only entities* only includes word tokens that belong to an entity annotation. Types starting with *Entity:* are calculated by only considering word tokens that are of such an entity.

at least allow for some noise in the predicted boundaries and thus, be able to capture these fuzzy borders more accurately, especially when compared to the conventional strict F_1 score.

Nevertheless, the adjusted F_1 is unable to fairly capture all variations in annotation and prediction difference, as shown in Subtable (b) of Table IV. However, it is much harder to actually detect that both options are viable in these cases. It would require the metric to have an in-depth auditing knowledge that our proposed weighting scheme simply can not have.

VI. CONCLUSION AND FUTURE WORK

In this paper, we introduce KPI-EDGAR, a financial document dataset for joint named entity recognition and relation extraction, based on annual financial reports submitted to the EDGAR database by publicly-traded companies and manually annotated by us. Moreover, we present an adjusted F_1 score, based on a weighting scheme, that allows for relations to be partially correct in their entity boundaries.

We find that this metric can alleviate the problem of fuzzy boundaries of entities and thus relations during the extraction of key performance indicators from financial reports and is therefore superior to the conventional F_1 score for measuring the success of the relation extraction task in this domain. As these fuzzy boundaries occur frequently in KPI-EDGAR, we propose that other researchers use this adjusted F_1 score when measuring model performances on it.

This study is part of a larger endeavor and long-time research project to advance information extraction processes in the financial text domain, where both machine learning researchers and audit experts work together on finding novel solutions to practical issues encountered during the audit process. In its wake, we plan to further improve availability, quality, and quantity of datasets. As a logical next step, we intend to build a named entity recognition dataset based on available GAAP (general accepted accounting principle)

entities from the EDGAR database. Given the massive amount of freely available financial reports on EDGAR, it is also conceivable to train a custom language model tailored to handle such financial documents, as shown with EDGAR-W2V [5] but with a state-of-the-art transformer model like BERT [27]. Additionally, in a further push to improve performance tracking, we will investigate dynamic ways to weight the metric, e.g., by incorporating the importance of each word that is part of an entity and relation.

REFERENCES

- [1] B. Marr, *Key Performance Indicators (KPI): The 75 measures every manager needs to know*. Pearson UK, 2012.
- [2] L. Hillebrand, T. Deußler, T. Dilmaghani, B. Kliem, R. Loitz, C. Bauckhage, and R. Sifa, “KPI-BERT: A joint named entity recognition and relation extraction model for financial reports,” in *Proc. ICPDR*, 2022.
- [3] M. Han, Y. Liang, Z. Duan, and Y. Wang, “Mining public business knowledge: A case study in sec’s edgar,” in *Int. Conf. on BDDC, SCN and SCC*. IEEE, 2016, pp. 393–400.
- [4] R. Ashraf, “Scraping EDGAR with python,” *J. Educ. Bus.*, vol. 92, no. 4, pp. 179–185, 2017.
- [5] L. Loukas, M. Fergadiotis, I. Androutsopoulos, and P. Malakasiotis, “EDGAR-CORPUS: Billions of tokens make the world go round,” in *Proc. Third Workshop on Economics and Natural Language Processing*, 2021, pp. 13–18.
- [6] G. Lonare, B. Patil, and N. Raut, “edgar: An R package for the US SEC EDGAR retrieval and parsing of corporate filings,” *SoftwareX*, vol. 16, p. 100865, 2021.
- [7] B. Lester, D. Pressel, A. Hemmeter, S. R. Choudhury, and S. Bangalore, “Constrained decoding for computationally efficient named entity recognition taggers,” in *Proc. EMNLP*, 2020.
- [8] A. Ushio and J. Camacho-Collados, “T-ner: An all-round python library for transformer-based named entity recognition,” in *Proc. EACL*, 2021.
- [9] X. Wang, Y. Jiang, N. Bach, T. Wang, Z. Yu Huang, F. Huang, and K. Tu, “Improving named entity recognition by external context retrieving and cooperative learning,” in *Proc. ACL/IJCNLP*, 2021.
- [10] X. Wang, Y. Jiang, N. Bach, T. Wang, Z. Huang, F. Huang, and K. Tu, “Automated concatenation of embeddings for structured prediction,” in *Proc. ACL*, 2021.
- [11] B. Xu, Q. Wang, Y. Lyu, Y. Zhu, and Z. Mao, “Entity structure within and throughout: Modeling mention dependencies for document-level relation extraction,” in *Proc. AAAI*, 2021.
- [12] S. Zeng, R. Xu, B. Chang, and L. Li, “Double graph based reasoning for document-level relation extraction,” in *Proc. EMNLP*, 2020.
- [13] N. Zhang, X. Chen, X. Xie, S. Deng, C. Tan, M. Chen, F. Huang, L. Si, and H. Chen, “Document-level relation extraction as semantic segmentation,” in *Proc. IJCAI*, 2021.
- [14] Q. Li and H. Ji, “Incremental joint extraction of entity mentions and relations,” in *Proc. ACL*, 2014.
- [15] A. Pappu, R. Blanco, Y. Mehdad, A. Stent, and K. Thadani, “Lightweight multilingual entity extraction and linking,” in *Proc. WSDM*, 2017.
- [16] M. Eberts and A. Ulges, “Span-based joint entity and relation extraction with transformer pre-training,” in *Proc. ECAI*, 2020.
- [17] C. Liang, Y. Yu, H. Jiang, S. Er, R. Wang, T. Zhao, and C. Zhang, “BOND: BERT-assisted open-domain named entity recognition with distant supervision,” in *Proc. KDD*, 2020.
- [18] Y. Shen, X. Ma, Y. Tang, and W. Lu, “A trigger-sense memory flow framework for joint entity and relation extraction,” in *Proc. Web Conference*, 2021.
- [19] J. Wang and W. Lu, “Two are better than one: Joint entity and relation extraction with table-sequence encoders,” in *Proc. EMNLP*, 2020.
- [20] D. Ye, Y. Lin, and M. Sun, “Pack together: Entity and relation extraction with levitated marker,” *arXiv:2109.06067*, 2021.
- [21] Z. Zhong and D. Chen, “A frustratingly easy approach for entity and relation extraction,” in *Proc. NAACL*, 2021.
- [22] V.-H. Tran, V.-T. Phi, A. Kato, H. Shindo, T. Watanabe, and Y. Matsumoto, “Improved decomposition strategy for joint entity and relation extraction,” *J. of NLP*, vol. 28, no. 4, pp. 965–994, 2021.
- [23] T. Hang, J. Feng, L. Yan, Y. Wang, and J. Lu, “Joint extraction of entities and relations using multi-label tagging and relational alignment,” *Neural Comput. Appl.*, pp. 1–16, 2022.
- [24] A. Savary, A. Silvanovich, A.-L. Minard, N. Hiot, and M. Halfeld Ferrari, “Relation extraction from clinical cases for a knowledge graph,” in *Proc. ADBIS*, 2022, pp. 353–365.
- [25] K. Fundel, R. Küffner, and R. Zimmer, “RelEx—relation extraction using dependency parse trees,” *Bioinformatics*, vol. 23, no. 3, 2007.
- [26] H. Gurulingappa, A. Mateen-Rajpu, and L. Toldo, “Extraction of potential adverse drug events from medical case reports,” *J. of Biomedical Semantics*, vol. 3, no. 1, 2012.
- [27] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. NAACL-HLT*, 2019.
- [28] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “RoBERTa: A robustly optimized bert pretraining approach,” *arXiv:1907.11692*, 2019.
- [29] A. Radford and K. Narasimhan, “Improving language understanding by generative pre-training,” 2018. [Online]. Available: <https://www.cs.ubc.ca/~amuham01/LING530/papers/radford2018improving.pdf>
- [30] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, “Language models are unsupervised multitask learners,” in *OpenAI Blog*, Accessed: 2021-10-01, 2019. [Online]. Available: <https://openai.com/blog/better-language-models>
- [31] D. Treiguiros and R. Berry, “The application of neural network based methods to the extraction of knowledge from accounting reports,” in *Proc. HICSS*, 1991.
- [32] D. Farmakiotou, V. Karkaletsis, J. Koutsias, G. Sigletos, C. D. Spyropoulos, and P. Stamatopoulos, “Rule-based named entity recognition for greek financial texts,” in *Proc. COMLEX*, 2000, pp. 75–78.
- [33] E. Brito, R. Sifa, C. Bauckhage, R. Loitz, U. Lohmeier, and C. Pünt, “A hybrid AI tool to extract key performance indicators from financial reports for benchmarking,” in *Proc. DocEng*, 2019, pp. 1–4.
- [34] Y. Cao, H. Li, P. Luo, and J. Yao, “Towards automatic numerical cross-checking: Extracting formulas from text,” in *Proc. WWW*, 2018.
- [35] I. Segura-Bedmar, P. Martínez, and M. Herrero-Zazo, “SemEval-2013 task 9 : Extraction of drug-drug interactions from biomedical texts (DDIExtraction 2013),” in *Proc. SemEval*, 2013, pp. 341–350.
- [36] J. Fu, P. Liu, and G. Neubig, “Interpretable multi-dataset evaluation for named entity recognition,” in *Proc. EMNLP*, 2020, pp. 6058–6069.
- [37] T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” vol. 2013, 01 2013.
- [38] Z. S. Mustafaoglu and N. Bogumill, “LSTM-based analysis of company sentiments regarding cryptocurrencies,” *Metamorphosis*, 2022.
- [39] J. P. Ryans, “Textual classification of SEC comment letters,” *Rev. Account. Stud.*, vol. 26, no. 1, pp. 37–80, 2021.
- [40] F. Li, “The information content of forward-looking statements in corporate filings—a naïve bayesian machine learning approach,” *J. Account. Res.*, vol. 48, no. 5, pp. 1049–1102, 2010.
- [41] J. Sim and C. C. Wright, “The kappa statistic in reliability studies: use, interpretation, and sample size requirements,” *Physical therapy*, vol. 85, no. 3, pp. 257–268, 2005.
- [42] J. Cohen, “A coefficient of agreement for nominal scales,” *Educational and psychological measurement*, vol. 20, no. 1, pp. 37–46, 1960.
- [43] J. Pennington, R. Socher, and C. D. Manning, “Glove: Global vectors for word representation,” in *Proc. EMNLP*, 2014, pp. 1532–1543.
- [44] M. Schuster and K. Nakajima, “Japanese and korean voice search,” in *Proc. ICASSP*, 2012.
- [45] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, “On the properties of neural machine translation: Encoder–decoder approaches,” in *Proc. SSSST-8*, 2014, pp. 103–111.