

FiNER: Financial Numeric Entity Recognition for XBRL Tagging

Lefteris Loukas^{1,2}, Manos Fergadiotis^{1,2}, Ilias Chalkidis³, Eirini Spyropoulou¹,
Prodromos Malakasiotis^{1,2}, Ion Androutsopoulos^{1,2}, George Paliouras¹

¹Institute of Informatics and Telecommunications, NCSR “Demokritos”

²Department of Informatics, Athens University of Economics and Business

³Department of Computer Science, University of Copenhagen

Abstract

Publicly traded companies are required to submit periodic reports with eXtensive Business Reporting Language (XBRL) word-level tags. Manually tagging the reports is tedious and costly. We, therefore, introduce XBRL tagging as a new entity extraction task for the financial domain and release FiNER-139, a dataset of 1.1M sentences with gold XBRL tags. Unlike typical entity extraction datasets, FiNER-139 uses a much larger label set of 139 entity types. Most annotated tokens are numeric, with the correct tag per token depending mostly on context, rather than the token itself. We show that subword fragmentation of numeric expressions harms BERT’s performance, allowing word-level BILSTMs to perform better. To improve BERT’s performance, we propose two simple and effective solutions that replace numeric expressions with pseudo-tokens reflecting original token shapes and numeric magnitudes. We also experiment with FIN-BERT, an existing BERT model for the financial domain, and release our own BERT (SEC-BERT), pre-trained on financial filings, which performs best. Through data and error analysis, we finally identify possible limitations to inspire future work on XBRL tagging.

1 Introduction

Natural language processing (NLP) for finance is an emerging research area (Hahn et al., 2019; Chen et al., 2020; El-Haj et al., 2020). Financial data are mostly reported in tables, but substantial information can also be found in textual form, e.g., in company filings, analyst reports, and economic news. Such information is useful in numerous financial intelligence tasks, like stock market prediction (Chen et al., 2019; Yang et al., 2019), financial sentiment analysis (Malo et al., 2014; Wang et al.,

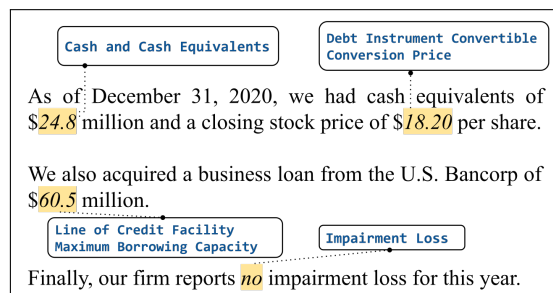


Figure 1: Sentences from FiNER-139, with XBRL tags on numeric and non-numeric tokens. XBRL tags are actually XML-based and most tagged tokens are numeric.

2013; Akhtar et al., 2017), economic event detection (Ein-Dor et al., 2019; Jacobs et al., 2018; Zhai and Zhang, 2019), and causality analysis (Tabari et al., 2018; Izumi and Sakaji, 2019). In this work, we study how financial reports can be automatically enriched with word-level tags from the eXtensive Business Reporting Language (XBRL), a tedious and costly task not considered so far.¹

To promote transparency among shareholders and potential investors, publicly traded companies are required to file periodic financial reports. These comprise multiple sections, including financial tables and text paragraphs, called *text notes*. In addition, legislation in the US, the UK, the EU and elsewhere requires the reports to be annotated with tags of XBRL, an XML-based language, to facilitate the processing of financial information. The annotation of tables can be easily achieved by using company-specific pre-tagged table templates, since the structure and contents of the tables in the reports of a particular company rarely change. On the other hand, the unstructured and dynamic nature of text notes (Figure 1) makes adding XBRL tags to them much more difficult. Hence, we focus on automatically tagging text notes. Tackling this task could facilitate the annotation of new and old reports (which may not include XBRL tags), e.g.,

Source code: <https://github.com/nlpaueb/finer>
Correspondence: eleftheriosloukas@aueb.gr

¹See <https://www.xbrl.org/the-standard/what/an-introduction-to-xbrl/> for an introduction to XBRL.

Dataset	Domain	Entity Types
CONLL-2003	Generic	4
ONTONOTES-v5	Generic	18
ACE-2005	Generic	7
GENIA	Biomedical	36
Chalkidis et al. (2019)	Legal	14
Francis et al. (2019)	Financial	9
FiNER-139 (ours)	Financial	139

Table 1: Examples of previous entity extraction datasets. Information about the first four from Tjong Kim Sang and De Meulder (2003); Pradhan et al. (2012); Doddington et al. (2004); Kim et al. (2003).

by inspecting automatically suggested tags.

Towards this direction, we release FiNER-139, a new dataset of 1.1M sentences with gold XBRL tags, from annual and quarterly reports of publicly traded companies obtained from the US Securities and Exchange Commission (SEC). Unlike other entity extraction tasks, like named entity recognition (NER) or contract element extraction (Table 1), which typically require identifying entities of a small set of common types (e.g., persons, organizations), XBRL defines approx. 6k entity types. As a first step, we consider the 139 most frequent XBRL entity types, still a much larger label set than usual.

Another important difference from typical entity extraction is that most tagged tokens (~91%) in the text notes we consider are numeric, with the correct tag per token depending mostly on context, not the token itself (Figure 1). The abundance of numeric tokens also leads to a very high ratio of out-of-vocabulary (OOV) tokens, approx. 10.4% when using a custom WORD2VEC (Mikolov et al., 2013a) model trained on our corpus. When using subwords, e.g., in models like BERT (Devlin et al., 2019), there are no OOV tokens, but numeric expressions get excessively fragmented, making it difficult for the model to gather information from the fragments and correctly tag them all. In our experiments, this is evident by the slightly better performance of stacked BILSTMs (Graves et al., 2013; Lample et al., 2016) operating on word embeddings compared to BERT. The latter improves when using a CRF (Lafferty et al., 2001) layer, which helps avoid assigning nonsensical sequences of labels to the fragments (subwords) of numeric expressions.

To further improve BERT’s performance, we propose two simple and effective solutions that replace numeric expressions with pseudo-tokens reflecting the original token shapes and magnitudes. We also experiment with FIN-BERT (Yang et al., 2020),

an existing BERT model for the financial domain, and release our own family of BERT models, pre-trained on 200k financial filings, achieving the best overall performance.

Our key contributions are:

1. We introduce XBRL tagging, a new financial NLP task for a real-world need, and we release FiNER-139, the first XBRL tagging dataset.²
2. We provide extensive experiment BILSTMs and BERT with generic or in-domain pre-training, which establish strong baseline results for future work on FiNER-139.
3. We show that replacing numeric tokens with pseudo-tokens reflecting token shapes and magnitudes significantly boosts the performance of BERT-based models in this task.
4. We release a new family of BERT models (SEC-BERT, SEC-BERT-NUM, SEC-BERT-SHAPE) pre-trained on 200k financial filings that obtains the best results on FiNER-139.^{3,4,5}

2 Related Work

Entity extraction: XBRL tagging differs from NER and other previous entity extraction tasks (Table 1), like contract element extraction (Chalkidis et al., 2019). Crucially, in XBRL tagging there is a much larger set of entity types (6k in full XBRL, 139 in FiNER-139), most tagged tokens are numeric (~91%), and the correct tag highly depends on context. In most NER datasets, numeric expressions are classified in generic entity types like ‘amount’ or ‘date’ (Bikel et al., 1999); this can often be achieved with regular expressions that look for common formats of numeric expressions, and the latter are often among the easiest entity types in NER datasets. By contrast, although it is easy to figure out that the first three highlighted expressions of Figure 1 are amounts, assigning them the correct XBRL tags requires carefully considering their context. Contract element extraction (Chalkidis et al., 2019) also requires considering the context of dates, amounts etc. to distinguish, for example, start dates from end dates, total amounts from other mentioned amounts, but the number of entity types in FiNER-139 is an order of magnitude larger (Table 1) and the full tag set of XBRL is even larger (6k).

²<https://huggingface.co/datasets/nlpaueb/finer-139>

³<https://huggingface.co/nlpaueb/sec-bert-base>

⁴<https://huggingface.co/nlpaueb/sec-bert-num>

⁵<https://huggingface.co/nlpaueb/sec-bert-shape>

Financial NER: Previous financial NER applications use at most 9 (generic) class labels. Salinas Alvarado et al. (2015) investigated NER in finance to recognize organizations, persons, locations, and miscellaneous entities on 8 manually annotated SEC financial agreements using CRFs. Francis et al. (2019) experimented with transfer learning by unfreezing different layers of a BiLSTM with a CRF layer, pre-trained on invoices, to extract 9 entity types with distinct morphological patterns (e.g., IBAN, company name, date, total amount). Also, Hampton et al. (2015, 2016) applied a Maximum Entropy classifier, CRFs, and handcrafted rules to London Stock Exchange filings to detect 9 generic entity types (e.g., person, organization, location, money, date, percentages). Finally, Kumar et al. (2016) extended the work of Finkel et al. (2005) and built a financial entity recognizer of dates, numeric values, economic terms in SEC and non-SEC documents, using numerous handcrafted text features. By contrast, FiNER-139 uses a specialized set of 139 highly technical economic tags derived from the real-world need of XBRL tagging, and we employ no handcrafted features.

Numerical reasoning: Neural numerical reasoning studies how to represent numbers to solve numeracy tasks, e.g., compare numbers, understand mathematical operations mentioned in a text etc. Zhang et al. (2020) released NUMBERT, a Transformer-based model that handles numerical reasoning tasks by representing numbers by their scientific notation and applying subword tokenization. On the other hand, GENBERT (Geva et al., 2020) uses the decimal notation and digit-by-digit tokenization of numbers. Both models attempt to deal with the problem that word-level tokenization often turns numeric tokens to OOVs (Thawani et al., 2021). This is important, because numerical reasoning requires modeling the exact value of each numeric token. In FiNER-139, the correct XBRL tags of numeric tokens depend much more on their contexts and token shapes than on their exact numeric values (Fig. 1). Hence, these methods are not directly relevant. GENBERT’s digit-by-digit tokenization would also lead to excessive fragmentation, which we experimentally find to harm performance.

3 Task and Dataset

Traditionally, business filings were simply rendered in plain text. Thus, analysts and researchers needed to manually identify, copy, and paste each

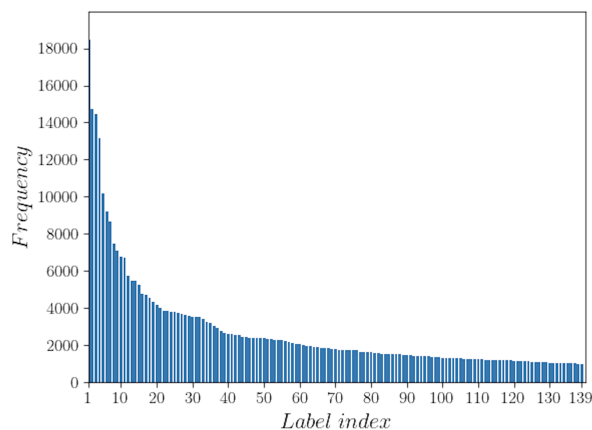


Figure 2: Frequency distribution of the 139 XBRL tags used in this work over the entire FiNER-139 dataset. Label indices shown instead of tag names to save space.

amount of interest (e.g., from filings to spreadsheets). With XBRL-tagged filings, identifying and extracting amounts of interest (e.g., to spreadsheets or databases) can be automated. More generally, XBRL facilitates the machine processing of financial documents. Hence, XBRL-tagged financial reports are required in several countries, as already noted (Section 1). However, manually tagging reports with XBRL tags is tedious and resource-intensive. Therefore, we release FiNER-139 to foster research towards automating XBRL tagging.

FiNER-139 was compiled from approx. 10k annual and quarterly English reports (filings) of publicly traded companies downloaded from SEC’s EDGAR system.⁶ The downloaded reports span a 5-year period, from 2016 to 2020. They are annotated with XBRL tags by professional auditors and describe the performance and projections of the companies. We used regular expressions to extract the text notes from the *Financial Statements Item* of each filing, which is the primary source of XBRL tags in annual and quarterly reports.

XBRL taxonomies have many different attributes, making XBRL tagging challenging even for humans (Baldwin et al., 2006; Hoitash and Hoitash, 2018). Furthermore, each jurisdiction has its own XBRL taxonomy. Since we work with US documents, our labels come from US-GAAP.⁷ Since this is the first effort towards automatic XBRL tagging, we chose to work with the most essential and informative attribute, the *tag names*, which populate our label set. Also, since XBRL tags change periodically, we selected the 139 (out of 6,008) most frequent XBRL tags with at least 1,000 appearances in FiNER-139.

⁶<https://www.sec.gov/edgar/>

⁷www.xbrl.us/xbrl-taxonomy/2020-us-gaap/

Subset	Sentences (S)	Avg. Tokens/S	Avg. Tags/S
Train	900,384	44.7 ± 33.9	1.8 ± 1.2
Dev	112,494	45.4 ± 35.9	1.7 ± 1.2
Test	108,378	46.5 ± 38.9	1.7 ± 1.1

Table 2: FiNER-139 statistics, using SPACY’s tokenizer and the 139 tags of this work (± standard deviation).

The distribution of these tags seems to follow a power law (Figure 2), hence most of the 6k XBRL tags that we did not consider are very rare. We used the IOB2 annotation scheme to distinguish tokens at the beginning, inside, or outside of tagged expressions, which leads to 279 possible token labels.

We split the text notes into 1.8M sentences, the majority of which (~90%) contained no tags.⁸ The sentences are also HTML-stripped, normalized, and lower-cased. To avoid conflating trivial and more difficult cases, we apply heuristic rules to discard sentences that can be easily flagged as almost certainly requiring no tagging; in a real-life setting, the heuristics, possibly further improved, would discard sentences that do not need to be processed by the tagger. The heuristic rules were created by inspecting the training subset and include regular expressions that look for amounts and other expressions that are typically annotated. Approx. 40% of the 1.8M sentences were removed, discarding only 1% of tagged ones. We split chronologically the remaining sentences into training, development, and test sets with an 80/10/10 ratio (Table 2).

4 Baseline Models

SPACY (Honnibal et al., 2020) is an open-source NLP library.⁹ It includes an industrial NER that uses word-level Bloom embeddings (Serrà and Karatzoglou, 2017) and residual Convolutional Neural Networks (CNNs) (He et al., 2016). We trained SPACY’s NER from scratch on FiNER-139.

BILSTM: This baseline uses a stacked bidirectional Long-Short Term Memory (LSTM) network (Graves et al., 2013; Lample et al., 2016) with residual connections. Each token t_i of a sentence S is mapped to an embedding and passed through the BILSTM stack to extract the corresponding contextualized embedding. A shared multinomial logistic regression (LR) layer operates on top of each contextualized embedding to predict the correct label. We use the WORD2VEC embeddings (Mikolov et al., 2013a,b) of Loukas et al. (2021).

⁸We use NLTK (Bird et al., 2009) for sentence splitting.

⁹We used SPACY v.2.3; see <https://spacy.io/>.

BERT: This is similar to BILSTM, but now we fine-tune BERT-BASE (Devlin et al., 2019) to extract contextualized embeddings of subwords. Again, a multinomial LR layer operates on top of the contextualized embeddings to predict the correct label of the corresponding subword.

CRFs: In this case, we replace the LR layer of the previous two models with a Conditional Random Field (CRF) layer (Lafferty et al., 2001), which has been shown to be beneficial in several token labeling tasks (Huang et al., 2015; Lample et al., 2016; Chalkidis et al., 2020b).¹⁰

5 Baseline Results

We report micro- F_1 (μ - F_1) and macro- F_1 (m - F_1) at the entity level, i.e., if a gold tag annotates a multi-word span, a model gets credit only if it tags the exact same span. This allows comparing more easily methods that label words vs. subwords.

Table 3 shows that SPACY performs poorly, possibly due to the differences from typical token labeling tasks, i.e., the large amount of entity types, the abundance of numeric tokens, and the fact that in FiNER-139 the tagging decisions depend mostly on context. Interestingly enough, BILSTM (with word embeddings) performs slightly better than BERT. However, when a CRF layer is added, BERT achieves the best results, while the performance of BILSTM (with word embeddings) deteriorates significantly, contradicting previous studies.

Baseline methods	μ - F_1	m - F_1
SPACY (words)	48.6 ± 0.4	37.6 ± 0.2
BILSTM (words)	<u>77.3</u> ± 0.6	<u>73.8</u> ± 1.8
BILSTM (subwords)	71.3 ± 0.2	68.6 ± 0.2
BERT (subwords)	75.1 ± 1.1	72.6 ± 1.4
BILSTM (words) + CRF	69.4 ± 1.2	67.3 ± 1.6
BILSTM (subwords) + CRF	76.2 ± 0.2	73.4 ± 0.3
BERT (subwords) + CRF	78.0 ± 0.5	75.2 ± 0.6

Table 3: Entity-level μ - F_1 and m - F_1 (%), avg. of 3 runs with different random seeds, ± std. dev.) on test data.

We hypothesize that the inconsistent effect of CRFs is due to tokenization differences. When using BERT’s subword tokenizer, there are more decisions that need to be all correct for a tagged span to be correct (one decision per subword) than when using word tokenization (one decision per word). Thus, it becomes more difficult for subword models to avoid nonsensical sequences of token labels,

¹⁰We use a linear-chain CRF layer with log-likelihood optimization and Viterbi decoding.

e.g., labeling two consecutive subwords as beginning and inside of different entity types, especially given the large set of 279 labels (Table 1). The CRF layer on top of subword models helps reduce the nonsensical sequences of labels.

On the other side, when using words as tokens, there are fewer opportunities for nonsensical label sequences, because there are fewer tokens. For instance, the average number of subwords and words per gold span is 2.53 and 1.04, respectively. Hence, it is easier for the BILSTM to avoid predicting nonsensical sequences of labels and the CRF layer on top of the BILSTM (with word embeddings) has less room to contribute and mainly introduces noise (e.g., it often assigns low probabilities to acceptable, but less frequent label sequences). With the CRF layer, the model tries to maximize both the confidence of the BILSTM for the predicted label of each word and the probability that the predicted sequence of labels is frequent. When the BILSTM on its own rarely predicts nonsensical sequences of labels, adding the CRF layer rewards commonly seen sequences of labels, even if they are not the correct labels, without reducing the already rare nonsensical sequences of labels.

To further support our hypothesis, we repeated the BILSTM experiments, but with subword (instead of word) embeddings, trained on the same vocabulary with BERT. Without the CRF, the subword BILSTM performs much worse than the word BILSTM (6 p.p drop in $\mu\text{-}F_1$), because of the many more decisions and opportunities to predict nonsensical label sequences. The CRF layer substantially improves the performance of the subword BILSTM (4.9 p.p. increase in $\mu\text{-}F_1$), as expected, though the word BILSTM (without CRF) is still better, because of the fewer opportunities for nonsensical predictions. A drawback of CRFs is that they significantly slow down the models both during training and inference, especially when using large label sets (Goldman and Goldberger, 2020), as in our case. Hence, although BERT with CRF was the best model in Table 3, we wished to improve BERT’s performance further without employing CRFs.

6 Fragmentation in BERT

In FiNER-139, the majority (91.2%) of the gold tagged spans are numeric expressions, which cannot all be included in BERT’s finite vocabulary; e.g., the token ‘9,323.0’ is split into five subword units, [‘9’, ‘##’, ‘##323’, ‘##.’, ‘##0’], while the token

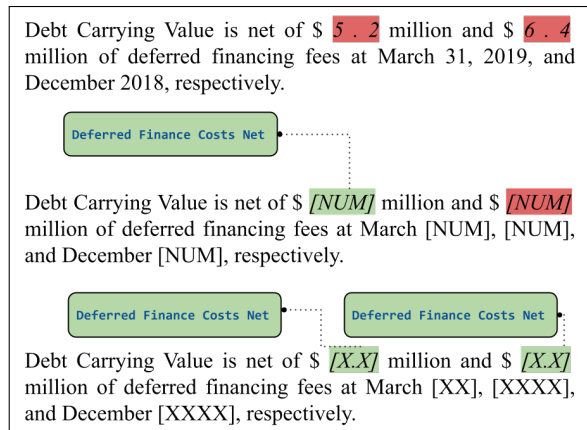


Figure 3: XBRL tag predictions of BERT (top), BERT + [NUM] (middle), BERT + [SHAPE] (bottom) for the same sentence. BERT tags incorrectly the amounts in red. BERT + [NUM] and BERT + [SHAPE] tag them more successfully (green indicates correct tags).

‘12.78’ is split into [‘12’, ‘##.’, ‘##78’]. The excessive fragmentation of numeric expressions, when using subword tokenization, harms the performance of the subword-based models (Table 3), because it increases the probability of producing nonsensical sequences of labels, as already discussed. We, therefore, propose two simple and effective solutions to avoid the over-fragmentation of numbers.

BERT + [NUM]: We detect numbers using regular expressions and replace each one with a single [NUM] pseudo-token, which cannot be split. The pseudo-token is added to the BERT vocabulary, and its representation is learned during fine-tuning. This allows handling all numeric expressions in a uniform manner, disallowing their fragmentation.

BERT + [SHAPE]: We replace numbers with pseudo-tokens that cannot be split and represent the number’s shape. For instance, ‘53.2’ becomes [XX.X], and ‘40,200.5’ becomes [XX,XXX.X]. We use 214 special tokens that cover all the number shapes of the training set. Again, the representations of the pseudo-tokens are fine-tuned, and numeric expressions (of known shapes) are no longer fragmented. The shape pseudo-tokens also capture information about each number’s magnitude; the intuition is that numeric tokens of similar magnitudes may require similar XBRL tags. Figure 3 illustrates the use of [NUM] and [SHAPE] pseudo-tokens.

7 In-domain Pre-training

Driven by the recent findings that pre-training language models on specialized domains is beneficial for downstream tasks (Alsentzer et al., 2019; Belt-

	DEVELOPMENT			TEST		
	μ -P	μ -R	μ -F ₁	μ -P	μ -R	μ -F ₁
BERT	74.9 ± 1.5	82.0 ± 1.3	78.2 ± 1.4	71.5 ± 1.1	79.6 ± 1.4	75.1 ± 1.1
BERT + CRF	<u>78.3</u> ± 0.8	<u>83.6</u> ± 0.4	<u>80.9</u> ± 0.3	<u>75.0</u> ± 0.9	<u>81.2</u> ± 0.2	<u>78.0</u> ± 0.5
BERT + [NUM]	79.4 ± 0.8	83.0 ± 0.9	81.2 ± 0.9	76.0 ± 0.6	80.7 ± 0.8	78.3 ± 0.7
BERT + [SHAPE]	<u>82.1</u> ± 0.6	82.6 ± 0.4	<u>82.3</u> ± 0.2	<u>78.7</u> ± 0.5	80.1 ± 0.2	<u>79.4</u> ± 0.2
FIN-BERT	73.9 ± 1.3	81.4 ± 0.7	77.5 ± 1.0	70.2 ± 1.2	78.7 ± 0.7	74.0 ± 1.1
FIN-BERT + [NUM]	81.1 ± 0.1	82.5 ± 1.2	81.8 ± 0.1	77.9 ± 0.1	79.9 ± 0.7	78.8 ± 0.3
FIN-BERT + [SHAPE]	<u>82.3</u> ± 1.7	<u>84.0</u> ± 1.2	<u>83.2</u> ± 1.4	<u>79.0</u> ± 1.6	<u>81.2</u> ± 1.1	<u>80.1</u> ± 1.4
SEC-BERT (ours)	75.2 ± 0.4	82.7 ± 0.5	78.8 ± 0.1	71.6 ± 0.4	80.3 ± 0.5	75.7 ± 0.1
SEC-BERT-NUM (ours)	82.5 ± 2.1	84.4 ± 1.2	83.7 ± 1.7	79.0 ± 1.9	82.0 ± 0.9	80.4 ± 1.4
SEC-BERT-SHAPE (ours)	84.8 ± 0.2	85.8 ± 0.2	85.3 ± 0.0	81.0 ± 0.2	83.2 ± 0.1	82.1 ± 0.1

Table 4: Entity-level micro-averaged P, R, F₁ ± std. dev. (3 runs) on the dev. and test data for BERT-based models.

agy et al., 2019; Yang et al., 2020; Chalkidis et al., 2020b), we explore this direction in our task which is derived from the financial domain.

FIN-BERT: We fine-tune FIN-BERT (Yang et al., 2020), which is pre-trained on a financial corpus from SEC documents, earnings call transcripts, and analyst reports.¹¹ The 30k subwords vocabulary of FIN-BERT is built from scratch from its pre-training corpus. Again, we utilize FIN-BERT with and without our numeric pseudo-tokens, whose representations are learned during fine-tuning.

SEC-BERT: We also release our own family of BERT models. Following the original setup of Devlin et al. (2019), we pre-trained BERT from scratch on EDGAR-CORPUS, a collection of financial documents released by Loukas et al. (2021). The resulting model, called SEC-BERT, has a newly created vocabulary of 30k subwords. To further examine the impact of the proposed [NUM] and [SHAPE] special tokens, we also pre-trained two additional BERT variants, SEC-BERT-NUM and SEC-BERT-SHAPE, on the same corpus, having replaced all numbers by [NUM] or [SHAPE] pseudo-tokens, respectively. In this case, the representations of the pseudo-tokens are learned during pre-training and they are updated during fine-tuning.

8 Improved BERT Results

Table 4 reports micro-averaged precision, recall, and F₁ on development and test data. As with Table 3, a LR layer is used on top of each embedding to predict the correct label, unless specified otherwise.

¹¹We use the FINBERT-FINVOcab-UNCASED version from <https://github.com/yya518/FinBERT>.

Focusing on the second zone, we observe that the [NUM] pseudo-token improves BERT’s results, as expected, since it does not allow numeric expressions to be fragmented. The results of BERT + [NUM] are now comparable to those of BERT + CRF. Performance improves further when utilizing the shape pseudo-tokens (BERT + [SHAPE]), yielding 79.4 μ -F₁ and showing that information about each number’s magnitude is valuable in XBRL tagging.

Interestingly, FIN-BERT (3rd zone) performs worse than BERT despite its pre-training on financial data. Similarly to BERT, this can be attributed to the fragmentation of numbers (2.5 subwords per gold tag span). Again, the proposed pseudo-tokens ([NUM], [SHAPE]) alleviate this problem and allow FIN-BERT to leverage its in-domain pre-training in order to finally surpass the corresponding BERT variants, achieving an 80.1 μ -F₁ test score.

Our new model, SEC-BERT (last zone), which is pre-trained on SEC reports, performs better than the existing BERT and FIN-BERT models, when no numeric pseudo-tokens are used. However, SEC-BERT is still worse than BERT with numeric pseudo-tokens (75.7 vs. 78.3 and 79.4 test μ -F₁), suffering from number fragmentation (2.4 subwords per gold tag span). SEC-BERT (without pseudo-tokens) also performs worse than the BILSTM with word embeddings (75.7 vs. 77.3 μ -F₁, cf. Table 3). However, when the proposed pseudo-tokens are used, SEC-BERT-NUM and SEC-BERT-SHAPE achieve the best overall performance, boosting the test μ -F₁ to 80.4 and 82.1, respectively. This indicates that learning to handle numeric expressions during model pre-training is a better strategy than trying to acquire this knowledge only during fine-tuning.

9 Additional Experiments

9.1 Subword pooling

An alternative way to bypass word fragmentation is to use subword pooling for each word. Ács et al. (2021) found that for NER tasks, it is better to use the *first* subword only, i.e., predict the label of an entire word from the contextualized embedding of its first subword only; they compared to several other methods, such as using only the *last* subword of each word, or combining the contextualized embeddings of all subwords with a self-attention mechanism. Given this finding, we conducted an ablation study and compare (i) our best model (SEC-BERT) with *first* subword pooling (denoted SEC-BERT-FIRST) to (ii) SEC-BERT with our special tokens (SEC-BERT-NUM, SEC-BERT-SHAPE), which avoid segmenting numeric tokens.

Table 5 shows that, in XBRL tagging, using the proposed special tokens is comparable (SEC-BERT-NUM) or better (SEC-BERT-SHAPE) than performing *first* pooling (SEC-BERT-FIRST). It might be worth trying other pooling strategies as well, like *last*-pooling or subword self-attention pooling. It’s worth noting, however, that the latter will increase the training and inference times.

	μ -F ₁	m-F ₁
SEC-BERT	78.8 ± 0.1	72.6 ± 0.4
SEC-BERT-FIRST	79.9 ± 1.2	77.1 ± 1.7
SEC-BERT-NUM	80.4 ± 1.4	78.9 ± 1.3
SEC-BERT-SHAPE	82.1 ± 0.1	80.1 ± 0.2

Table 5: Entity-level μ -F₁ and m-F₁ (% , avg. of 3 runs with different random seeds, ± std. dev.) on test data using different ways to alleviate fragmentation.

9.2 Subword BILSTM with [NUM] and [SHAPE]

To further investigate the effectiveness of our pseudo-tokens, we incorporated them in the BILSTM operating on subword embeddings (3rd model of Table 3). Again, we replace each number by a single [NUM] pseudo-token or one of 214 [SHAPE] pseudo-tokens, for the two approaches, respectively. These replacements also happen when pre-training WORD2VEC subword embeddings; hence, an embedding is obtained for each pseudo-token.

Table 6 shows that BILSTM-NUM outperforms the BILSTM subword model. BILSTM-SHAPE further improves performance and is the best BILSTM subword model overall, surpassing the subword BILSTM with CRF, which was the best subword

BILSTM model in Table 3. These results further support our hypothesis that the [NUM] and [SHAPE] pseudo-tokens help subword models successfully generalize over numeric expressions, with [SHAPE] being the best of the two approaches, while also avoiding the over-fragmentation of numbers.

	μ -F ₁	m-F ₁
BILSTM (subwords)	71.3 ± 0.2	68.6 ± 0.2
BILSTM (subwords) + CRF	76.2 ± 0.2	73.4 ± 0.3
BILSTM-NUM (subwords)	75.6 ± 0.3	72.7 ± 0.4
BILSTM-SHAPE (subwords)	76.8 ± 0.2	74.1 ± 0.3

Table 6: Entity-level μ -F₁ and m-F₁ (% , avg. of 3 runs with different random seeds, ± std. dev.) on test data for BILSTM models with [NUM] and [SHAPE] tokens.

9.3 A Business Use Case

Since XBRL tagging is derived from a real-world need, it is crucial to analyze the model’s performance in a business use case. After consulting with experts of the financial domain, we concluded that one practical use case would be to use an XBRL tagger as a recommendation engine that would propose the k most probable XBRL tags for a specific token selected by the user. The idea is that an expert (e.g., accountant, auditor) knows beforehand the token(s) that should be annotated and the tagger would assist by helping identify the appropriate tags more quickly. Instead of having to select from several hundreds of XBRL tags, the expert would only have to inspect a short list of k proposed tags.

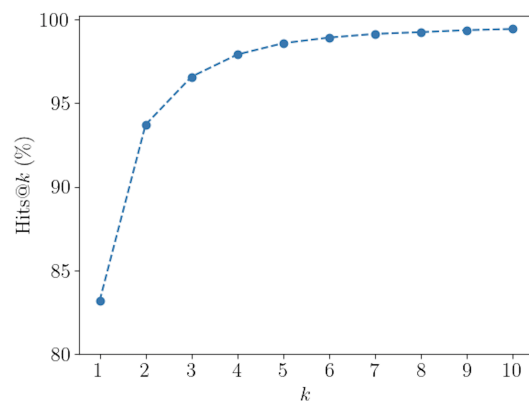


Figure 4: Hits@ k results (% , avg. of 3 runs with different random seeds) on test data, for different k values. Standard deviations were very small and are omitted.

We evaluate our best model, SEC-BERT-SHAPE, in this use case using Hits@ k . We use the model to return the k most probable XBRL tags for each token that needs to be annotated, now assuming that

the tokens to be annotated are known. If the correct tag is among the top k , we increase the number of hits by one. Finally, we divide by the number of tokens to be annotated. Figure 4 shows the results for different values of k . The curve is steep for $k = 1$ to 5 and saturates as k approaches 10, where Hits@ k is nearly perfect (99.4%). In practice, this means that a user would have to inspect 10 recommended XBRL tags instead of hundreds for each token to be annotated; and in most cases, the correct tag would be among the top 5 recommended ones.

9.4 Error Analysis

We also performed an exploratory data and error analysis to unveil the peculiarities of FiNER-139, extract new insights about it, and discover the limitations of our best model. Specifically, we manually inspected the errors of SEC-BERT-SHAPE in under-performing classes (where $F_1 < 50\%$) and identified three main sources of errors.

Specialized terminology: In this type of errors, the model is able to understand the general financial semantics, but does not fully comprehend highly technical details. For example, *Operating Lease Expense* amounts are sometimes missclassified as *Lease And Rental Expense*, i.e., the model manages to predict that these amounts are about expenses in general, but fails to identify the specific details that distinguish operating lease expenses from lease and rental expenses. Similarly, *Payments to Acquire Businesses (Net of Cash Acquired)* amounts are mostly missclassified as *Payments to Acquire Businesses (Gross)*. In this case, the model understands the notion of business acquisition, but fails to differentiate between net and gross payments.

Financial dates: Another interesting error type is the misclassification of financial dates. For example, tokens of the class *Debt Instrument Maturity Date* are mostly missclassified as not belonging to any entity at all ('O' tag). Given the previous type of errors, one would expect the model to missclassify these tokens as a different type of financial date, but this is not the case here. We suspect that errors of this type may be due to annotation inconsistencies by the financial experts.

Annotation inconsistencies: Even though the gold XBRL tags of FiNER-139 come from professional auditors, as required by the Securities & Exchange Commission (SEC) legislation, there are still some discrepancies. We provide an illustrative

Excluding short-term leases, the Company's leases have initial terms ranging from 1 to 16 years, most of which include options the Company may elect to extend or renew the lease for 0.1 to 6 years, and some of which may include options to terminate the leases with notice periods of up to 1 year.

Figure 5: A manually inspected sentence from FiNER-139 showing some inconsistencies in the gold XBRL tags of the auditors. The green '1' is correctly annotated with the XBRL tag *Lessee Operating Lease Term Of Contract*. The red '16' should have also been annotated with the same tag, but is not, possibly because the annotator thought the (same) tag was obvious. The orange numbers '0.1' and '6' lack XBRL annotations; they should have both been annotated as *Lessee Operating Lease Renewal Term*. We can only speculate that the auditor might not have been aware that there is an XBRL tag for lease renewal terms, in which case the recommendation engine of Section 9.3 might have helped.

example in Figure 5. We believe that such inconsistencies are inevitable to occur and they are a part of the real-world nature of the problem.

We hope that this analysis inspires future work on XBRL tagging. For example, the specialized terminology and financial date errors may be alleviated by adopting hierarchical classifiers (Chalkidis et al., 2020a; Manginas et al., 2020), which would first detect entities in coarse classes (e.g., expenses, dates) and would then try to classify the identified entities into finer classes (e.g., lease vs. rent expenses, instrument maturity dates vs. other types of dates). It would also be interesting to train classifiers towards detecting wrong (or missing) gold annotations, in order to help in quality assurance checks of XBRL-tagged documents.

10 Conclusions and Future Work

We introduced a new real-word NLP task from the financial domain, XBRL tagging, required by regulatory commissions worldwide. We released FiNER-139, a dataset of 1.1M sentences with XBRL tags. Unlike typical entity extraction tasks, FiNER-139 uses a much larger label set (139 tags), most tokens to be tagged are numeric, and the correct tag depends mostly on context rather than the tagged token. We experimented with several neural classifiers, showing that a BILSTM outperforms BERT due to the excessive numeric token fragmentation of the latter. We proposed two simple and effective solutions that use special tokens to generalize over the shapes and magnitudes of numeric expressions. We also experimented with FIN-BERT, an existing

BERT model for the financial domain, which also benefits from our special tokens. Finally, we pre-trained and released our own domain-specific BERT model, SEC-BERT, both with and without the special tokens, which achieves the best overall results with the special tokens, without costly CRF layers.

In future work, one could hire experts to re-annotate a subset of the dataset to measure human performance against the gold tags. Future work could also consider less frequent XBRL tags (few- and zero-shot learning) and exploit the hierarchical dependencies of XBRL tags, possibly with hierarchical classifiers, building upon our error analysis.

References

- Judit Ács, Ákos Kádár, and Andras Kornai. 2021. [Sub-word pooling makes a difference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2284–2295, Online.
- Md Shad Akhtar, Abhishek Kumar, Deepanway Ghosal, Asif Ekbal, and Pushpak Bhattacharyya. 2017. [A multilayer perceptron based ensemble technique for fine-grained financial sentiment analysis](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 540–546, Copenhagen, Denmark.
- Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, USA.
- Amelia Annette Baldwin, Carol E. Brown, and Brad S. Trinkle. 2006. Xbrl: An impacts framework and research challenge. *Journal of Emerging Technologies in Accounting*, 3:97–116.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3615–3620, Hong Kong, China.
- Daniel M. Bikel, Richard Schwartz, and Ralph M. Weischedel. 1999. An algorithm that learns what’s in a name. *Machine Learning*, 34(1-3):211–231.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python*, 1st edition. O’Reilly Media, Inc.
- Ilias Chalkidis, Manos Fergadiotis, Sotiris Kotitsas, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020a. [An empirical study on large-scale multi-label text classification including few and zero-shot labels](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020*, pages 7503–7515.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020b. [LEGAL-BERT: The muppets straight out of law school](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2898–2904, Online.
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, and Ion Androutsopoulos. 2019. [Neural contract element extraction revisited](#). In *Workshop on Document Intelligence at NeurIPS 2019*.
- Chung-Chi Chen, Hen-Hsen Huang, Hiroya Takamura, and Hsin-Hsi Chen, editors. 2020. *Proceedings of the Second Workshop on Financial Technology and Natural Language Processing*. Kyoto, Japan.
- Deli Chen, Shuming Ma, Keiko Harimoto, Ruihan Bao, Qi Su, and Xu Sun. 2019. [Group, extract and aggregate: Summarizing a large amount of finance news for forex movement prediction](#). In *Proceedings of the Second Workshop on Economics and Natural Language Processing*, pages 41–50, Hong Kong.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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)*, pages 4171–4186, Minneapolis, Minnesota.
- George Doddington, Alexis Mitchell, Mark Przybocki, Lance Ramshaw, Stephanie Strassel, and Ralph Weischedel. 2004. [The automatic content extraction \(ACE\) program – tasks, data, and evaluation](#). In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC’04)*, Lisbon, Portugal.
- Liat Ein-Dor, Ariel Gera, Orith Toledo-Ronen, Alon Halfon, Benjamin Sznajder, Lena Dankin, Yonatan Bilu, Yoav Katz, and Noam Slonim. 2019. [Financial event extraction using Wikipedia-based weak supervision](#). In *Proceedings of the Second Workshop on Economics and Natural Language Processing*, pages 10–15, Hong Kong.
- Dr Mahmoud El-Haj, Dr Vasiliki Athanasakou, Dr Sira Ferradans, Dr Catherine Salzedo, Dr Ans Elhag, Dr Houda Bouamor, Dr Marina Litvak, Dr Paul Rayson, Dr George Giannakopoulos, and Nikiforos Pittaras, editors. 2020. *Proceedings of the 1st Joint Workshop on Financial Narrative Processing and MultiLing Financial Summarisation*. COLING, Barcelona, Spain (Online).
- Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. [Incorporating non-local information into information extraction systems by Gibbs](#)

- sampling. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pages 363–370, Ann Arbor, Michigan.
- 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.
- Mor Geva, Ankit Gupta, and Jonathan Berant. 2020. Injecting numerical reasoning skills into language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 946–958.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 249–256, Chia Laguna Resort, Italy.
- Eran Goldman and Jacob Goldberger. 2020. Crf with deep class embedding for large scale classification. *Computer Vision and Image Understanding*, 191:102865.
- Alex Graves, Abdel rahman Mohamed, and Geoffrey E. Hinton. 2013. Speech recognition with deep recurrent neural networks. *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6645–6649.
- Udo Hahn, Véronique Hoste, and Zhu Zhang, editors. 2019. *Proceedings of the Second Workshop on Economics and Natural Language Processing*. Association for Computational Linguistics, Hong Kong.
- 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*, pages 295–306. Springer.
- Peter John Hampton, Hui Wang, and William Blackburn. 2015. A hybrid ensemble for classifying and repurposing financial entities. In *Research and Development in Intelligent Systems XXXII*, pages 197–202. Springer.
- Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.
- Rani Hoitash and Udi Hoitash. 2018. Measuring accounting reporting complexity with xbrl. *The Accounting Review*, 93:259–287.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. spaCy: Industrial-strength Natural Language Processing in Python.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *CoRR*, abs/1508.01991.
- Kiyoshi Izumi and Hiroki Sakaji. 2019. Economic causal-chain search using text mining technology. In *Proceedings of the First Workshop on Financial Technology and Natural Language Processing*, pages 61–65, Macao, China.
- Gilles Jacobs, Els Lefever, and Véronique Hoste. 2018. Economic event detection in company-specific news text. In *Proceedings of the First Workshop on Economics and Natural Language Processing*, pages 1–10, Melbourne, Australia.
- J-D Kim, Tomoko Ohta, Yuka Tateisi, and Jun'ichi Tsujii. 2003. Genia corpus—a semantically annotated corpus for bio-textmining. *Bioinformatics*, 19:i180–i182.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, USA, May 7-9, 2015, Conference Track Proceedings*.
- 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.
- John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01*, page 282–289, San Francisco, USA.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California.
- Lefteris Loukas, Manos Fergadiotis, Ion Androutsopoulos, and Prodromos Malakasiotis. 2021. EDGAR-CORPUS: Billions of tokens make the world go round. In *Proceedings of the Third Workshop on Economics and Natural Language Processing*, pages 13–18, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Pekka Malo, Ankur Sinha, Pekka J. Korhonen, Jyrki Wallenius, and Pyry Takala. 2014. Good debt or bad debt: Detecting semantic orientations in economic texts. *J. Assoc. Inf. Sci. Technol.*, 65(4):782–796.
- Nikolaos Manginas, Ilias Chalkidis, and Prodromos Malakasiotis. 2020. Layer-wise guided training for BERT: Learning incrementally refined document

- representations. In *Proceedings of the Fourth Workshop on Structured Prediction for NLP*, pages 53–61, Online. Association for Computational Linguistics.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. [Efficient estimation of word representations in vector space](#). In *1st International Conference on Learning Representations, ICLR 2013*, Scottsdale, USA.
- Tomás Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013b. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013*, pages 3111–3119, Lake Tahoe, United States.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. [CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes](#). In *Joint Conference on EMNLP and CoNLL - Shared Task*, pages 1–40, Jeju Island, Korea.
- 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, Paramatta, Australia.
- Joan Serrà and Alexandros Karatzoglou. 2017. [Getting deep recommenders fit: Bloom embeddings for sparse binary input/output networks](#). In *Proceedings of the Eleventh ACM Conference on Recommender Systems*, page 279–287, New York, USA.
- Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. [Practical bayesian optimization of machine learning algorithms](#). In *Advances in Neural Information Processing Systems*, volume 25.
- Narges Tabari, Piyusha Biswas, Bhanu Praneeth, Armin Seyeditabari, Mirsad Hadzikadic, and Wlodek Zadrozny. 2018. [Causality analysis of Twitter sentiments and stock market returns](#). In *Proceedings of the First Workshop on Economics and Natural Language Processing*, pages 11–19, Melbourne, Australia.
- Avijit Thawani, Jay Pujara, Filip Ilievski, and Pedro Szekely. 2021. [Representing numbers in NLP: a survey and a vision](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–656, Online.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. In *Proceedings of CoNLL-2003*, pages 142–147. Edmonton, Canada.
- Chuan-Ju Wang, Ming-Feng Tsai, Tse Liu, and Ching-Ting Chang. 2013. [Financial sentiment analysis for risk prediction](#). In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 802–808, Nagoya, Japan.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online.
- Linyi Yang, Ruihai Dong, Tin Lok James Ng, and Yang Xu. 2019. [Leveraging BERT to improve the FEARS index for stock forecasting](#). In *Proceedings of the First Workshop on Financial Technology and Natural Language Processing*, pages 54–60, Macao, China.
- Yi Yang, Mark Christopher Siy UY, and Allen Huang. 2020. [Finbert: A pretrained language model for financial communications](#). *arXiv*, abs/2006.08097.
- Shuang (Sophie) Zhai and Zhu (Drew) Zhang. 2019. [Forecasting firm material events from 8-k reports](#). In *Proceedings of the Second Workshop on Economics and Natural Language Processing*, pages 22–30, Hong Kong.
- Xikun Zhang, Deepak Ramachandran, Ian Tenney, Yanai Elazar, and Dan Roth. 2020. [Do language embeddings capture scales?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4889–4896, Online.

A Experimental Setup

For SPACY, we followed the recommended practices.¹² All other methods were implemented in TENSORFLOW.¹³ Concerning BERT models, we used the implementation of HUGGINGFACE (Wolf et al., 2020). We also use Adam (Kingma and Ba, 2015), Glorot initialization (Glorot and Bengio, 2010), and the categorical cross-entropy loss.

Hyper-parameters were tuned on development data with Bayesian Optimization (Snoek et al., 2012) monitoring the development loss for 15 trials.¹⁴ For the BILSTM encoders, we searched for {1, 2, 3} hidden layers, {128, 200, 256} hidden units, {1e-3, 2e-3, 3e-3, 4e-3, 5e-3} learning rate, and {0.1, 0.2, 0.3} dropout. We trained for 30 epochs using early stopping with patience 4. For BERT, we used grid-search to select the optimal learning rate from {1e-5, 2e-5, 3e-5, 4e-5, 5e-5}, fine-tuning for 10 epochs, using early stopping with patience 2. All final hyper-parameters are shown in Table 7. Training was performed mainly on a DGX station with 4 NVIDIA V100 GPUs and an Intel Xeon CPU E5-2698 v4 @ 2.20GHz.

	<i>Params</i>	<i>L</i>	<i>U</i>	<i>P_{drop}</i>	<i>LR</i>
BILSTM (words)	21M	2	128	0.1	1e-3
BILSTM (subwords)	8M	1	256	0.2	1e-3
BILSTM (words) + CRF	21M	2	128	0.1	1e-3
BILSTM (subwords) + CRF	8M	1	256	0.2	1e-3
BILSTM-NUM (subwords)	1M	1	256	0.2	1e-3
BILSTM-SHAPE (subwords)	0.8M	2	128	0.1	1e-3
BERT	110M	-	-	-	1e-5
BERT + [NUM]	110M	-	-	-	1e-5
BERT + [SHAPE]	110M	-	-	-	1e-5
BERT + CRF	110M	-	-	-	1e-5
FIN-BERT	110M	-	-	-	2e-5
FIN-BERT + [NUM]	110M	-	-	-	2e-5
FIN-BERT + [SHAPE]	110M	-	-	-	2e-5
SEC-BERT	110M	-	-	-	1e-5
SEC-BERT-NUM	110M	-	-	-	1e-5
SEC-BERT-SHAPE	110M	-	-	-	1e-5

Table 7: Number of total parameters (*Params*) and the best hyper-parameter values for each method; i.e., number of recurrent layers (*L*), number of recurrent units (*U*), dropout probability P_{drop} , learning rate (*LR*).

B Additional Results

Table 8 shows micro-averaged Precision, Recall, and F1 for the development and test data, using all

baseline methods. The macro-averaged scores are similar and we omit them for brevity. Using a logistic regression (LR) classification layer, BILSTM (words) surpasses BERT both in Precision and F1 score. However, when using a CRF layer on top, BERT outperforms BILSTM (words) in all measures.

Table 9 shows the micro-averaged Precision, Recall, and F_1 for the development and test data using the BILSTM models operating on subwords with the proposed tokenizations. [NUM] and [SHAPE] tokens help the model to bypass the word fragmentation problem, increasing its scores in all metrics.

¹²<https://spacy.io/usage/v2-3>.

¹³<https://www.tensorflow.org/>

¹⁴We used KERAS TUNER (<https://keras-team.github.io/keras-tuner/documentation/tuners/>)

	DEVELOPMENT			TEST		
	μ -P	μ -R	μ -F ₁	μ -P	μ -R	μ -F ₁
SPACY	38.2 ± 0.4	58.2 ± 0.8	46.1 ± 0.1	40.8 ± 0.8	60.0 ± 0.4	48.6 ± 0.4
BILSTM (words)	78.6 ± 2.4	80.3 ± 1.2	79.4 ± 1.0	75.4 ± 1.9	78.0 ± 0.8	77.3 ± 0.6
BILSTM (subwords)	73.4 ± 0.1	77.2 ± 0.0	75.2 ± 0.1	68.8 ± 0.2	74.1 ± 0.2	71.3 ± 0.2
BERT (subwords)	74.9 ± 1.5	82.0 ± 1.3	78.2 ± 1.4	71.5 ± 1.1	79.6 ± 1.4	75.1 ± 1.1
BILSTM (words) + CRF	73.4 ± 2.0	69.3 ± 0.9	71.3 ± 1.2	70.9 ± 1.8	68.0 ± 0.9	69.4 ± 1.2
BILSTM (subwords) + CRF	80.0 ± 0.3	78.7 ± 0.5	79.3 ± 0.4	76.5 ± 0.2	76.0 ± 0.2	76.2 ± 0.2
BERT (subwords) + CRF	78.3 ± 0.8	83.6 ± 0.4	80.9 ± 0.3	75.0 ± 0.9	81.2 ± 0.2	78.0 ± 0.5

Table 8: Entity-level micro-averaged P, R, F1 ± std. dev. (3 runs) on the dev. and test data for our baselines.

	DEVELOPMENT			TEST		
	μ -P	μ -R	μ -F ₁	μ -P	μ -R	μ -F ₁
BILSTM (subwords)	73.4 ± 0.1	77.2 ± 0.0	75.2 ± 0.1	68.8 ± 0.2	74.1 ± 0.2	71.3 ± 0.2
BILSTM (subwords) + CRF	80.0 ± 0.3	78.7 ± 0.5	79.3 ± 0.4	76.5 ± 0.2	76.0 ± 0.2	76.2 ± 0.2
BILSTM-NUM (subwords)	77.9 ± 0.4	78.6 ± 0.7	78.2 ± 0.6	74.8 ± 0.2	76.5 ± 0.5	75.6 ± 0.3
BILSTM-SHAPE (subwords)	81.1 ± 0.1	81.5 ± 0.3	81.3 ± 0.2	77.5 ± 0.3	78.7 ± 0.5	78.1 ± 0.4

Table 9: Entity-level micro-averaged P, R, F1 ± std. dev. (3 runs) on the dev. and test data for the BILSTM models using the [NUM] and [SHAPE] tokens.