

Named Entity Recognition in Twitter: A Dataset and Analysis on Short-Term Temporal Shifts

Asahi Ushio¹, Leonardo Neves², Vítor Silva², Francesco Barbieri², Jose Camacho-Collados¹

¹Cardiff NLP, School of Computer Science and Informatics, Cardiff University, United Kingdom

{UshioA,CamachoColladosJ}@cardiff.ac.uk

²Snap Inc., Santa Monica, CA, United States

{lneves,vsilvasousa,fbarbieri}@snap.com

Abstract

Recent progress in language model pre-training has led to important improvements in Named Entity Recognition (NER). Nonetheless, this progress has been mainly tested in well-formatted documents such as news, Wikipedia, or scientific articles. In social media the landscape is different, in which it adds another layer of complexity due to its noisy and dynamic nature. In this paper, we focus on NER in Twitter, one of the largest social media platforms, and construct a new NER dataset, *TweetNER7*, which contains seven entity types annotated over 11,382 tweets from September 2019 to August 2021. The dataset was constructed by carefully distributing the tweets over time and taking representative trends as a basis. Along with the dataset, we provide a set of language model baselines and perform an analysis on the language model performance on the task, especially analyzing the impact of different time periods. In particular, we focus on three important temporal aspects in our analysis: short-term degradation of NER models over time, strategies to fine-tune a language model over different periods, and self-labeling as an alternative to lack of recently-labeled data. *TweetNER7* is released publicly¹ along with the models fine-tuned on it².

1 Introduction

Named Entity Recognition (NER) is a long-standing NLP task that consists of identifying an entity in a sentence or document, and classifying it into an entity-type from a fixed typeset. One of the most common and successful types of NER system is achieved by fine-tuning pre-trained language models (LMs) on a human-annotated NER dataset

¹<https://huggingface.co/datasets/tner/tweetner7>

²NER models have been integrated into TweetNLP (Camacho-Collados et al., 2022) and can be found at https://github.com/asahi417/tner/tree/master/examples/tweetner7_paper

with token-wise classification (Peters et al., 2018; Howard and Ruder, 2018; Radford et al., 2018, 2019; Devlin et al., 2019). Remarkably, LM fine-tuning based NER models (Yamada et al., 2020; Li et al., 2020) already achieve over 90% F1 score in standard NER datasets such as CoNLL2003 (Tjong Kim Sang and De Meulder, 2003) and OntoNotes5 (Hovy et al., 2006). However, NER is far from being solved, specialized domains such as financial news (Salinas Alvarado et al., 2015), biochemical (Collier and Kim, 2004), or biomedical (Wei et al., 2015; Li et al., 2016) still pose additional challenges (Ushio and Camacho-Collados, 2021). Lower performance in these domains may be attributed to various factors such as the usage specific terminologies within those domains, which LMs have not seen while pre-training (Lee et al., 2020).

Among recent studies, social media has been acknowledged as one of the most challenging domains for NER (Derczynski et al., 2016, 2017). Social media texts are generally more noisy and less formal than conventional written languages in addition to its vocabulary specificity. In social media, there is another particular feature that needs to be addressed, which is the presence of (quick) temporal shifts in the text semantics (Rijhwani and Preotiuc-Pietro, 2020), where the meaning of words is constantly changing or evolving over time. This is a general issue with language models (Lazaridou et al., 2021), but it is especially relevant given the dynamic landscape and immediacy present in social media (Del Tredici et al., 2019). There have been a few specific approaches to deal with the temporal shifts in social media. For instance, Loureiro et al. (2022) addressed this issue by pre-training language models on a large tweet collection from different time period, highlighting the importance of having an up-to-date language model. Agarwal and Nenkova (2022) studied the temporal-shift in various NLP tasks including NER and analyzed methods to overcome the temporal-

shift with strategies such as self-labeling.

In this paper, we propose a new NER dataset for Twitter (*TweetNER7* henceforth). *TweetNER7* contains tweets from diverse topics that are distributed uniformly from September 2019 to August 2021. It contains 11,382 annotated tweets in total, spanning seven entity types (*person, location, corporation, creative work, group, product, and event*). To the best of our knowledge, *TweetNER7* is the largest Twitter NER datasets with a high coverage of entity types TTC (Rijhwani and Preotiuc-Pietro, 2020) contains about same amount of annotation yet with three entity types, while WNUT17 (Derczynski et al., 2017) has six entity types yet suffer from very small annotations. The tweets for *TweetNER7* were collected by querying tweets with weekly trending keywords so that the tweet collection covers various topics within the period, and we further removed near-duplicated tweets and irrelevant tweets without any specific topics in order to improve the quality of tweets. We provide baseline results with language model fine-tuning that showcases the difficulty of *TweetNER7*, especially when dealing with time shifts. Finally, we provide a temporal analysis with different strategies including self-labeling, which does not prove highly beneficial in our context, and provide insights in the model inner working and potential biases.

2 Related Work

There is a large variety of NER datasets in the literature. CoNLL2003 (Tjong Kim Sang and De Meulder, 2003) and OntoNotes5 (Hovy et al., 2006) are widely used common NER datasets in the literature, where the texts are collected from public news, blogs, and dialogues. WikiAnn (Pan et al., 2017) and MultiNERD (Tedeschi and Navigli, 2022) are both multilingual NER datasets where the training set is constructed by distant-supervision on Wikipedia and BabelNet. As far as domain-specific NER datasets are concerned, FIN (Salinas Alvarado et al., 2015) is a NER dataset of financial news, while BioNLP2004 (Collier and Kim, 2004) and BioCreative (Wei et al., 2015; Li et al., 2016) are both constructed from scientific documents of the biochemical and biomedical domains. However, none of these datasets address the same challenges posed by the social media domain.

In the social media domain, the pioneering Broad Twitter Corpus (BTC) NER dataset (Derczynski

et al., 2016) included users with different demographics with the aim to investigate spatial and temporal shift of semantics in NER. More recently, the test set of WNUT2017 (Derczynski et al., 2017) contained unseen entities in the training set from broader social media including Twitter, Reddit, YouTube, and StackExchange. The recent TweepBankNER dataset (Jiang et al., 2022) annotated TweepBank (Liu et al., 2018) with entity labels to investigate the interaction between syntax and NER.

The most similar dataset to ours is the Temporal Twitter Corpus (TTC) NER dataset. (Rijhwani and Preotiuc-Pietro, 2020), which was also aimed at analysing the temporal effects of NER in social media. For this dataset, 2,000 tweets every year from 2014 to 2019 were annotated. In general, however, these social media datasets suffer from limited data, non-uniform distribution over time, or limited entity types (see Subsection 3.3 for more details). In this paper, we contribute with a new NER dataset (*TweetNER7*) based on recent data until 2021, which is specifically designed to analyze temporal shifts in social media.

3 TweetNER7: Dataset Construction, Statistics and Baselines

In this section, we present our time-aware NER dataset from publicly available tweets with seven general entity types, which we refer as *TweetNER7*. In the following subsections, we describe the data collection (Subsection 3.1) and annotation (Subsection 3.2) processes. We also share relevant statistics (Subsection 3.3) and baseline results (Subsection 3.4) of our dataset.

3.1 Data Collection

This NER dataset annotates a similar tweet collection used to construct *TweetTopic* (Antypas et al., 2022). The main data consists of tweets from September 2019 to August 2021 with roughly same amount of tweets in each month. This collection period makes it suitable for our purpose of evaluating short-term temporal-shift of NER on Twitter. The original tweets were filtered by leveraging weekly trending topics as well as by various other types of filtering see Antypas et al. (2022) for more details on the collection and filtering process). The collected tweets were then split into two periods: September 2019 to August 2020 (2020-set) and September 2020 to August 2021 (2021-set).

3.2 Dataset Annotation

Annotation. To attain named-entity annotations over the tweets, we conducted a manual annotation on Amazon Mechanical Turk with the interface shown in Figure 1. We split tweets into two periods: September 2019 to August 2020 (2020-set) and September 2020 to August 2021 (2021-set), and randomly sampled 6,000 tweets from each period, which were annotated by three annotators, collecting 36,000 annotations in total. As the entity types, we employed seven labels: *person*, *location*, *corporation*, *creative work*, *group*, *product*, and *event*. We followed Derczynski et al. (2017) for the selection of the first six labels, and additionally included *event*, as we found a large amount of entities for events in our collected tweets.

Pre-processing. We pre-process tweets before the annotation to normalize some artifacts, converting URLs into a special token `{{URL}}` and non-verified usernames into `{{USERNAME}}`. For verified usernames, we replace its display name with symbols @. For example, a tweet

```
Get the all-analog Classic Vinyl Edition
of "Takin' Off" Album from @herbiehancock
via @bluenoterecords link below:
http://bluenote.lnk.to/AlbumOfTheWeek
```

is transformed into the following text.

```
Get the all-analog Classic Vinyl Edition
of "Takin' Off" Album from {@Herbie Hancock@}
via {{USERNAME}} link below: {{URL}}
```

We ask annotators to ignore those special tokens but label the verified users' mentions.

Quality Control. Since we have three annotations per tweet, we control the quality of the annotation by taking the agreement into account. We disregard the annotation if the agreement is 1/3, and manually validate the annotation if it is 2/3, which happens for roughly half of the instances.

3.3 Statistics

This subsection provides an statistical analysis of (i) our dataset, (ii) our dataset in comparison with other Twitter NER datasets, and (iii) our dataset distribution over time.

Statistics of TweetNER7. TweetNER7 contains 5,768 and 5,614 tweets annotated in each period of 2020 and 2021, which are then split into training / validation / test sets for each year. Since the 2020-set is for model development, we consider 80% of the dataset as training set and 10% for validation and test sets. Meanwhile, the 2021-set is

Period Split	2020-set			2021-set		
	Train	Valid	Test	Train	Valid	Test
Number of Entities						
- corporation	1,700	203	191	902	102	900
- creative work	1,661	208	179	690	74	731
- event	2,242	256	265	968	131	1,097
- group	2,242	227	311	1,313	227	1,516
- location	1,259	181	165	697	72	716
- person	4,666	598	596	2,362	283	2,712
- product	1,850	241	220	926	111	972
All	15,620	1,914	1,927	8,864	1,000	8,644
Entity Diversity						
- corporation	69.9	92.6	90.1	72.1	85.3	74.3
- creative work	80.1	92.8	91.6	89.0	93.2	91.0
- event	71.1	90.6	84.2	75.9	89.3	70.9
- group	66.7	86.8	81.7	66.0	86.3	66.2
- location	66.4	80.7	81.2	67.9	88.9	64.9
- person	68.4	85.6	83.6	77.3	90.1	77.7
- product	56.2	71.4	76.4	60.3	79.3	56.6
Number of Tweets	4,616	576	576	2,495	310	2,807

Table 1: Number of entities, tweets, and entity diversity in each data split and period, where the 2020-set is from September 2019 to August 2020, while the 2021-set is from September 2020 to August 2021.

mainly devised for model evaluation to measure the temporal adaptability, so we take the majority of the 2021-set (50%) as the test set and split the rest into training and validation set with the same ratio of training and validation set of the 2020-set. Table 1 summarizes the number of the entities as well as the instances in each subset of TweetNER7. We can observe a large gap between frequent entity types such as *person* and rare entity types as *location*, while the distribution of the entities are roughly balanced across subsets. We also report entity diversity, which we define as the percentage of unique entities with respect to the total number of entities. Entity types such as *product* contain a relatively large number of duplicates (ranging between 56.2% and 76.4% entity diversity scores), while other types such as creative work are more diverse (ranging between 80.1% and 93.2%).

Comparison with other Twitter NER Datasets.

In Table 2, we compare TweetNER7 against existing NER datasets for Twitter, which highlights the large number of annotations of TweetNER7 for our covered period. TweetNER7 and TTC are the overall largest datasets with more than 10k annotations, but TTC covers only three entities, which may be insufficient for certain practical use cases given the diversity of text in social media context (Derczynski et al., 2017). In contrast, TweetNER7 has the highest coverage of entity types among all

READ THE GUIDELINE BEFORE START!! (Click to collapse)

In this project we aim at labelling named entities which are words that belong to specific domains from Twitter. You will need to annotate these special words when you encounter them. There are seven classes of entity in this task: *Person, Location, Corporation, Product, Creative work, Group, Event*.

IMPORTANT NOTES:

- If the entity consists of multiple words such as "[CBS] [Sports] [Radio]" and "[St] [,] [Patrick] [s] [Day]", you **need to annotate all the words composing the entity**.
- The verified twitter usernames are replaced by their displayed name with highlights "{@displayed name@}" (e.g. "@Cristiano was on fire!!" -> "{@Cristiano Ronaldo@} was on fire!!"). You **need to annotate those twitter usernames**.

DON'T ANNOTATE...

- **more than one entity types** on single entity. For example, given a text "*I went to Disney Store*", you are not allowed to annotate "*Disney Store*" as both of *Location* and *Corporation*, but you have to choose the most appropriate entity type, that is *Location* in this example.
- entities **without their own/unique name** as they are not named-entities (e.g. "*school*", "*teacher*", "*pencil*").
- **{(USERNAME)}** and **{(URL)}** as they are custom tokens for *non-verified* twitter username and web url.
- **non-English** words.
- the **hashtag #**.

Description for each entity type

Person (p)	Names of people (e.g. Virginia Wade). Include punctuation in the middle of names. Fictional people can be included, as long as they're referred to by name (e.g. 'Harry Potter').
Location (l)	Names that are locations (e.g. France). Include punctuation in the middle of names. Fictional locations can be included, as long as they're referred to by name (e.g. 'Hogwarts').
Group (g)	Names of groups (e.g. Nirvana, San Diego Padres). There may be no groups mentioned by name in the sentence at all - that's OK. Fictional groups can be included, as long as they're referred to by name.
Event (e)	Names of events (e.g. Christmas, Super Bowl). There may be no events mentioned by name in the sentence at all - that's OK. Fictional events can be included, as long as they're referred to by name.
Product (d)	Name of products (e.g. iPhone). Include punctuation in the middle of names. There may be no products mentioned by name in the sentence at all - that's OK. Fictional products can be included, as long as they're referred to by name (e.g. 'Everlasting Gobstopper'). It's got to be something you can touch, and it's got to be the official name.
Creative work (w)	Names of creative works (e.g. Bohemian Rhapsody). Include punctuation in the middle of names. The work should be created by a human, and referred to by its specific name.
Corporation (c)	Names of corporations (e.g. Google). Include punctuation in the middle of names.

Please also check more detailed guideline [here](#).

Figure 1: The instructions shown to the annotators during the annotation phase.

Dataset	Annotations	Entities	Domain	Year
BTC	9,339	3	Twitter	2009-2015
WNUT2017	5,690	6	Twitter+	2010-2017
TTC	11,969	3	Twitter	2014-2019
TweeBankNER	3,547	4	Twitter	2016
TweetNER7	11,382	7	Twitter	2019-2021

Table 2: Number of annotated instances in TweetNER7 and comparison NER datasets for Twitter.

NER datasets in Twitter, including all the entity types from existing datasets. In addition to the large amount of annotations and a high coverage of entity types, TweetNER7 includes recent tweets from 2019 to 2021, from which most corpus used in pre-training language models do not contain any text (Devlin et al., 2019; Liu et al., 2019; Nguyen et al., 2020). Assuming we tackle NER by language model fine-tuning, this fact makes the task further challenging, since language models have never seen the emerging entities from the period during its pre-training phase.

Distribution over Time. One of the TweetNER7's focus is the temporal shift in Twitter similar to BTC

	Jan	Feb	Mar	Apr	May	Jun
BTC	2,308	68	502	862	1,074	1,056
TTC	945	1,014	1,307	1,089	764	694
TweetNER7	957	943	939	937	951	931
	Jul	Aug	Sep	Oct	Nov	Dec
BTC	1,321	850	342	419	23	21
TTC	760	754	889	958	958	866
TweetNER7	924	928	956	968	975	973

Table 3: The number of tweets in each month from BTC, TTC, and our TweetNER7 (the counts are cumulated across years). The normalized standard deviation across month is 7.5% (BTC), 1.6% (TTC), and 0.2% (TweetNER7).

and TTC datasets. Retaining uniform distribution over time is essential for temporal analysis, since the amount of training instances should have an effect to the metric if it is not uniform. Table 3 shows the distribution of the instances across each month and we can confirm that TweetNER7 has a very similar amount of tweets each month, while BTC and TTC have higher variation than TweetNER7. Moreover, Table 4 compares the number

	2009	2010	2011	2012	2013	2014	2015
BTC	3	5	127	2,414	275	6,022	0
TTC	0	0	0	0	0	2,000	2,000
TweetNER7	0	0	0	0	0	0	0
	2016	2017	2018	2019	2020	2021	
BTC	0	0	0	0	0	0	
TTC	2,000	2,000	2,000	2,000	0	0	
TweetNER7	0	0	0	1,936	5,768	3,678	

Table 4: The number of tweets in each year from BTC, TTC, and our TweetNER7 dataset.

of instances per year for each dataset. TweetNER7 has an uneven distribution here due to the selected range for each period (i.e., September 2019 to August 2021), which results in more tweets in 2020 than 2019 and 2021.

3.4 Baseline Results

Finally, we introduce a couple of baselines with language model fine-tuning on the TweetNER7 in temporal-shift setup, where we develop models with the training and the validation set from the 2020-set, and evaluate the models on the test set of the 2021-set. In this setup, models are required to generalize to the text from newer period, which the model has not seen in the fine-tuning phase.

Experimental Setting. We consider masked language model fine-tuning with the following LMs: BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) as general-purpose LMs, and BERTweet (Nguyen et al., 2020), and TimeLMs (Loureiro et al., 2022) as Twitter-specific LMs. TimeLMs are based on a RoBERTa_{BASE} architecture pre-trained on tweets collected continuously until different years: 2019, 2020, and 2021. Model weights are taken from HuggingFace (Wolf et al., 2020).³ As evaluation metrics, we consider micro/macro F1 score and type-ignored F1 score (Ushio and Camacho-Collados, 2021), in which the entity type of the prediction is not considered in the evaluation (i.e., this metric only assesses whether the predicted entity is an entity or not). The F1 scores measure the NER systems’ entire performance, while the type-ignored F1 score measures the ability of identifying whether a span of text is

³We use bert-base-cased and bert-large-cased for BERT, roberta-base and roberta-large for RoBERTa, vinai/bertweet-base and vinai/bertweet-large for BERTweet, and cardiffnlp/twitter-roberta-base-2019-90m, cardiffnlp/twitter-roberta-base-dec2020, and cardiffnlp/twitter-roberta-base-dec2021 for TimeLMs.

Model	Micro F1	Macro F1	Type-ig. F1
	2021 / 2020	2021 / 2020	2021 / 2020
BERT _{BASE}	60.1 / 60.9	54.7 / 56.5	75.6 / 72.4
BERT _{LARGE}	61.4 / 62.2	56.1 / 58.1	75.9 / 73.8
BERTweet _{BASE}	64.1 / <u>66.4</u>	59.4 / 62.4	<u>77.9</u> / <u>77.7</u>
BERTweet _{LARGE}	64.0 / 65.9	59.5 / <u>62.6</u>	78.3 / 77.4
RoBERTa _{BASE}	64.2 / 64.2	59.1 / 60.2	77.9 / 74.8
RoBERTa _{LARGE}	64.8 / 65.7	60.0 / 61.9	78.4 / 76.1
TimeLM ₂₀₁₉	64.3 / 65.4	59.3 / 61.1	77.9 / 76.6
TimeLM ₂₀₂₀	62.9 / 64.4	58.3 / 60.3	76.5 / 75.7
TimeLM ₂₀₂₁	64.2 / 65.4	59.5 / 61.1	77.4 / 76.4

Table 5: Result of temporal-shift NER on TweetNER7 where micro and macro F1 score as well as type-ignored F1 score on the test set of the 2021-set / 2020-set are reported. The best results in each of the 2021-set / 2020-set are highlighted in bold character / underline in each metric.

an entity or not. LM fine-tuning on NER relies on the T-NER library (Ushio and Camacho-Collados, 2021) and to find the best combination of hyperparameters to fine-tune LMs on NER, we run two-phase grid search. First, we fine-tune a model on every possible configuration from the search space for 10 epochs. The top-5 models in terms of micro F1 score on the validation set are selected to continue fine-tuning until their performance plateaus, and then the model that achieves the highest micro F1 score on the validation set is employed as the final model. The search space contains 24 configurations, which consist of the following variations: learning rates from [0.000001, 0.00001, 0.0001]; ratio of total training step for linear warm up of learning rate from [0.15, 0.3]; whether to normalize the gradient norm or not; and whether to add conditional random field (CRF) on top of the output logit of LM.⁴

Results. We report the NER results on TweetNER7 in Table 5, where RoBERTa_{LARGE} is the best across metrics. We should note, however, that the overall metrics (micro F1 lower than 65% in all cases on the 2021 test set) are lower than those in standard NER datasets (Ushio and Camacho-Collados, 2021), which highlights the difficulty of the social media and temporal-shift components in TweetNER7. RoBERTa is also the best model among the _{BASE} models but interestingly the TimeLM₂₀₂₀ performs worse than other RoBERTa models. This can be explained by the fact that TimeLM₂₀₂₀ was pre-trained over tweets until the end of 2020. This may have let the model to over-fit to the training

⁴Other parameters are fixed: random seed is 0 and batch size is 32.

corpus and makes it hard to generalize on the newer test set. Instead, TimeLM₂₀₂₁ shows a better performance. Table 7 also reports the metrics on the 2020 test set for completeness. While that is not our primary aim, we can find an interesting result which is the superior performance of BERTweet in this case. This implies that a model that performs well in the same period of the training set does not guarantee an equally strong performance on an unseen period.

Breakdown by entity type. Figure 2 shows a comparison of entity-wise F1 scores over the language models, and we can see an important gap across entity types. According to Table 1, *person* is the most frequent entity type and its F1 score is equally high (around 80%), while *creative work* and *location* are the rarest entity types and hence their F1 scores are relatively low (around 40% for *creative work* and 60% for *location*). The reason why the performance for *location* is better than for *creative work* may be attributable to their differences in entity diversity. As we could see from Table 1, *creative work*'s diversity is higher than *location*, which means *creative work* contains more variation of entities than *location* while having the same amount of entities in both types, which entails a higher degree of difficulty. This seems a consistent trend that lower entity diversity results in lower F1 score as can be seen for *event* and *corporation* as well, which also have a low entity diversity score. To overcome such entity imbalance, strategies such as balancing the instances of each class could be explored (Li et al., 2020).

4 Temporal Analysis

To better understand the effect of the temporal-shift, we conduct three additional comparative experiments: (i) temporal vs. random splits, (ii) joint vs. continuous fine-tuning, and (iii) self-labeling as a solution to deal with temporal shifts.

4.1 Short-Term Temporal Effect

If TweetNER7 does not suffer temporal-shift, how is the model performance changed? This is a question we aim to answer in this analysis, and we create new training and validation split without temporal-shift for this purpose. Concretely, temporal-shift usually occurs in a situation where the training and the validation sets do not contain any texts from the test period, so we keep the amount of the training/validation split as the

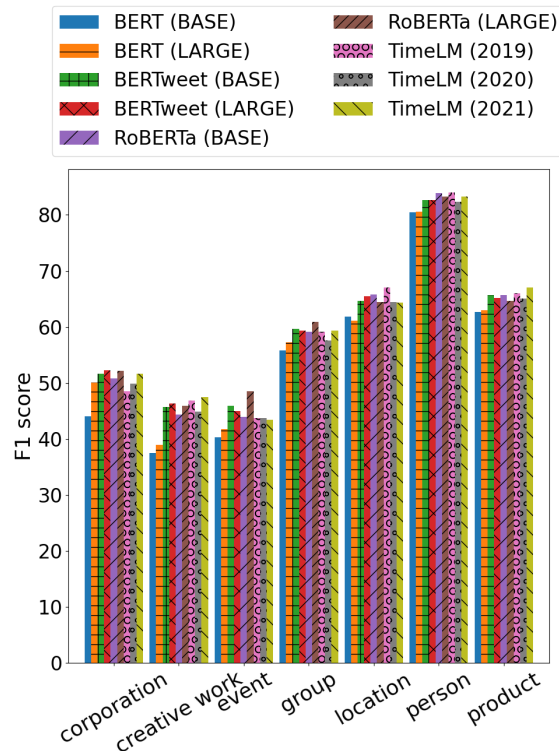


Figure 2: Entity-wise F1 score breakdown from the baseline results in the 2021 test set (Table 5).

same in Subsection 3.4, but randomly sample from the full period of September 2019 to August 2021 instead of the first half period instead. Note that we do not change the test set and make sure that each month has roughly the same amount of instances at the sampling of the new training/validation sets, to make it fair comparison with the temporal-shift result in Subsection 3.4.

Table 6 shows the variations of results between the random and temporal splits. As expected, the F1 scores on the 2021 test set are generally improved across all LMs, while the F1 scores on the 2020 test set are decreased. The increase of accuracy in 2021 is achieved with the inclusion of training/validation set from 2021, and the decrease of accuracy in 2020 is caused by the reduced number of the training/validation set from the same 2020 period. This result further highlights the benefit of having a human annotated training set from the test period, even if the time period differs in a year only. Interestingly, the results for the time-specific pre-trained TimeLMs models differ across years. Since in this paper we did not focus on the analysis of the pre-training corpora, we leave further analysis about this result for future exploration.

Model	2021-set			2020-set		
	Mi. F1	Ma. F1	T-i. F1	Mi. F1	Ma. F1	T-i. F1
BERT _{BASE}	+0.8	+1.2	+0.1	+0.1	+0.3	+0.3
BERT _{LARGE}	+1.0	+1.4	+0.6	-0.7	-1.0	-0.5
BERTweet _{BASE}	+1.5	+0.2	-0.1	-2.5	-3.8	-3.3
BERTweet _{LARGE}	+0.9	+1.0	+0.1	+0.1	+0.1	-0.2
RoBERTa _{BASE}	-0.2	+0.1	+0.1	-0.1	-0.4	-0.5
RoBERTa _{LARGE}	+1.5	+1.0	+0.6	-1.3	-1.8	-0.6
TimeLM ₂₀₁₉	-1.0	-0.8	-0.5	-1.1	-0.4	-0.4
TimeLM ₂₀₂₀	+1.8	+1.7	+1.8	+0.3	+0.2	+0.2
TimeLM ₂₀₂₁	-1.0	-1.1	-0.4	-1.7	-1.3	-1.0

Table 6: Absolute performance improvement when evaluating on the random split result over the original temporal split reported in Table 5. Positive improvements are in blue and negative drops are in red.

4.2 Continuous vs. Joint Fine-Tuning

In the previous experiments we have shown the differences between training and testing on the time period or not. Instead, this analysis comes under the assumption that a labeled 2021 training set is available. Thus, the main aim of this analysis is to explore different strategies to improve the original model. In addition to fine-tuning LMs on the combined set of the 2020-set and 2021-set as in Subsection 4.1, we employed a continuous fine-tuning scheme, where we first fine-tune LMs on the 2020-set and then continue fine-tuning on the 2021-set. Table 7 shows the results of all strategies for different language models. As can be observed, continuous fine-tuning provides the best results in terms of micro F1 and type-ignored F1 in the 2021 test sets in most cases, although the differences with respect to the concatenation of sets are not substantial.

4.3 Self-Labeling

In both Subsections 4.1 and 4.2, we compared different strategies when a human-annotated training dataset from the test period was considered, namely the training and the validation sets from the 2021-set. This shows that improvements can be obtained when the time between training and test data is reduced. However, in many cases and real-world applications this is not practical as it requires a large amount of human resources to annotate newer tweets whenever. Thus, we consider an alternative approach to rely on distantly annotated tweets by the already fine-tuned model. This solution was explored by Agarwal and Nenkova (2022) in a similar setting, with promising results. In this paper, we reproduced their experiments in our TweetNER7 dataset focusing on short-term temporal shift.

	Dataset	Micro F1	Macro F1	Type-ig. F1
BERT _{BASE}	2020	60.1 / 60.9	54.7 / 56.5	75.6 / 72.4
	2021	60.7 / 58.4	55.5 / 54.2	75.7 / 70.9
	2020 + 2021	62.3 / 62.1	57.6 / 57.7	76.6 / 73.0
	2020 → 2021	61.8 / 61.4	56.8 / 57.1	76.5 / 72.5
BERT _{LARGE}	2020	61.4 / 62.2	56.1 / 58.1	75.9 / 73.8
	2021	59.7 / 56.6	53.9 / 51.0	75.0 / 70.7
	2020 + 2021	63.6 / 62.5	59.0 / 58.6	77.2 / 73.6
	2020 → 2021	63.2 / 62.5	57.7 / 57.9	76.0 / 72.5
BERTweet _{BASE}	2020	64.1 / 66.4	59.4 / 62.4	77.9 / 77.7
	2021	63.1 / 62.1	57.4 / 57.2	77.9 / 76.0
	2020 + 2021	65.4 / 65.7	60.5 / 61.6	79.0 / 76.9
	2020 → 2021	65.8 / 65.2	61.0 / 61.4	79.1 / 76.8
BERTweet _{LARGE}	2020	64.0 / 65.9	59.5 / 62.6	78.3 / 77.4
	2021	62.9 / 61.6	58.1 / 56.8	76.5 / 74.5
	2020 + 2021	66.5 / 66.8	61.9 / 63.1	79.5 / 77.6
	2020 → 2021	66.4 / 65.9	61.7 / 61.8	79.0 / 76.4
RoBERTa _{BASE}	2020	64.2 / 64.2	59.1 / 60.2	77.9 / 74.8
	2021	61.8 / 60.5	57.0 / 56.1	76.9 / 73.8
	2020 + 2021	65.2 / 65.3	60.8 / 61.7	78.9 / 75.2
	2020 → 2021	65.5 / 65.1	60.0 / 60.8	78.1 / 75.0
RoBERTa _{LARGE}	2020	64.8 / 65.7	60.0 / 61.9	78.4 / 76.1
	2021	64.0 / 63.4	59.1 / 59.1	77.7 / 74.4
	2020 + 2021	65.7 / 66.3	61.2 / 63.0	78.8 / 76.4
	2020 → 2021	66.0 / 66.3	60.9 / 62.4	79.1 / 76.4
TimeLM ₂₀₁₉	2020	64.3 / 65.4	59.3 / 61.1	77.9 / 76.6
	2021	63.2 / 61.9	56.7 / 56.1	75.7 / 73.0
	2020 + 2021	65.7 / 65.5	61.0 / 61.2	78.9 / 76.4
	2020 → 2021	65.9 / 64.8	61.1 / 60.6	78.4 / 75.5
TimeLM ₂₀₂₀	2020	62.9 / 64.4	58.3 / 60.3	76.5 / 75.7
	2021	64.0 / 63.1	58.9 / 58.5	77.9 / 75.3
	2020 + 2021	65.3 / 65.4	60.7 / 61.4	78.7 / 75.9
	2020 → 2021	65.5 / 65.3	60.6 / 61.3	78.0 / 75.9
TimeLM ₂₀₂₁	2020	64.2 / 65.4	59.5 / 61.1	77.4 / 76.4
	2021	63.5 / 62.3	58.7 / 57.9	77.5 / 74.1
	2020 + 2021	64.5 / 65.8	59.8 / 61.9	77.9 / 76.5
	2020 → 2021	65.1 / 64.9	60.0 / 60.7	78.1 / 75.8

Table 7: Results of different strategies to ingest the training set of the 2021-set in TweetNER7 for different language models (→: continuous fine-tuning; +: concatenation of datasets). The best results in each model of the 2021-set / 2020-set are highlighted in bold character / underline in each metric.

4.3.1 Evaluation

Experimental Setting. For our experiments we focused on the best model in our previous experiments, which is RoBERTa_{LARGE}. We collected extra (unlabeled) tweets following the same procedure described in (Antypas et al., 2022), that results in 93,594 and 878,80 tweets from the period of 2020-set and 2021-set, respectively. Over those extra tweets, we use the RoBERTa_{LARGE} NER model fine-tuned on the 2020-set to predict labels.

Results. Table 8 shows the result of self-labeling, where we report three patterns of model fine-tuning: (i) fine-tuning only on the pseudo dataset (e.g., 2020-extra); (ii) fine-tuning on the joint dataset

Training Set	Micro F1		Macro F1		Type-ig. F1	
	2021 / 2020	2021 / 2020	2021 / 2020	2021 / 2020	2021 / 2020	2021 / 2020
2020	64.8 / <u>65.7</u>	60.0 / <u>61.9</u>	78.4 / 76.1			
2020-extra	64.6 / 65.5	59.3 / 61.4	78.6 / 76.2			
2020 + 2020-extra	64.7 / 65.2	59.6 / 61.0	78.7 / 76.8			
2020 → 2020-extra	64.6 / 65.5	59.5 / 61.5	78.6 / 76.4			
2021-extra	64.2 / <u>65.7</u>	59.3 / 61.8	78.2 / <u>76.9</u>			
2020 + 2021-extra	64.3 / 65.6	59.3 / 61.7	78.4 / <u>76.9</u>			
2020 → 2021-extra	64.5 / 65.5	59.5 / 61.4	78.6 / 76.3			

Table 8: Results of the self-labeling experiment with different strategies for RoBERTa_{LARGE} model (→: continuous fine-tuning; +: concatenation of datasets) where micro and macro F1 score as well as type-ignored F1 score on the test set of 2021-set / 2020-set are reported. The best results in each of the 2021-set / 2020-set are highlighted in bold character / underline in each metric.

of the training set of the 2020-set and the pseudo dataset (e.g., 2020 + 2020-extra); and (iii) continuous fine-tuning of the 2020-set fine-tuned model on the pseudo dataset (e.g., 2020 → 2020-extra). In general, we can not find any major improvement by self-labeling, regardless of the strategy. In a way, this contradicts the self-labeling experiment on the TTC dataset performed by Agarwal and Nenkova (2022).⁵ This may suggest that the temporal-shift of TweetNER7 is more challenging to mitigate than TTC, and self-labeling is not enough in itself to overcome the temporal shift.

4.3.2 Contextual Prediction Analysis

To explore the reason why self-labeling does not help to mitigate temporal-shift in TweetNER7, we conducted an analysis over the self-labeled tweets. Inspired by recent semi-parametric approach in information retrieval (Lewis et al., 2021), we considered a retrieval module that fetches relevant tweets given a target entity from the self-labeled corpus and see the portion of retrieved tweets containing the true prediction. To be precise, we first ran the NER model prediction on target tweets, and for each of the predicted entities. Then, we queried tweets from the extra tweet corpus used in Subsection 4.3 to compute the ratio of correct predictions within the retrieved predictions, which we call contextualized predictions. Since we are interested in the error of the original prediction, we focus only on the entities where the original prediction is incorrect.

Figure 3 describes the whole pipeline and we

⁵While in our setting we extract a larger number of tweets, this trend does not change with less self-labeled training data.

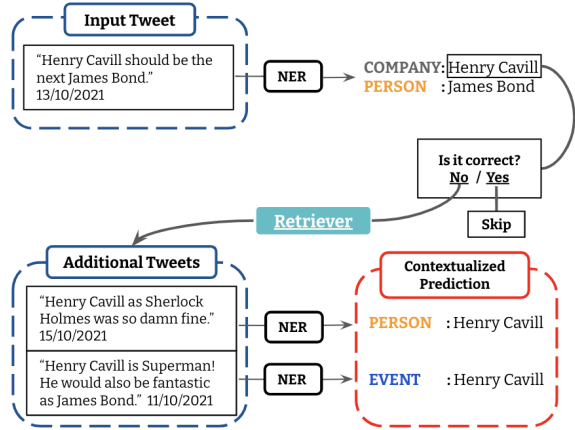


Figure 3: Overview of the pipeline to retrieve contextualized prediction.

use Whoosh library⁶ for search engine where the query is always the entity name, constraining the search result by the number of days from the query tweet.⁷ Similarly to the analysis in § 4.3, we used the RoBERTa_{LARGE} fine-tuned on the 2020-set of TweetNER7 and evaluated the contextualized predictions on the 2021 test set.

Figure 4 shows the ratio of positive and negative predictions in the contextualized tweets. These are further broken into two error types whether it is the same prediction as the original prediction or not, along with the days we set as a search constraint. Most frequent predictions are usually the same as the original predictions, which means that the original language model tends to output similar predictions for the same entities, irrespective of the context. As far as the time variable is concerned, the ratio is almost consistent over time, which suggests that the possible original bias of the model does not change over time. Nonetheless, the second most frequent predictions are on average the correct ones, with a large gap with respect to other types of error. This implies there may still be a useful signal to improve the original prediction in the self-labeled corpus.

5 Conclusion

In this paper, we have constructed TweetNER7, a new NER dataset for Twitter, in which we annotated 11,382 tweets with seven entity types. The collected tweets are distributed uniformly over time from September 2019 to August 2021, which facil-

⁶<https://pypi.org/project/Whoosh/>

⁷Setting days as 7 means the search results should be in the range of 7 days before/after was made.

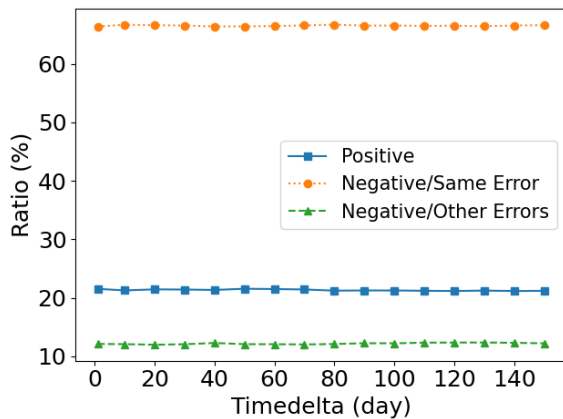


Figure 4: Ratio of positive and negative predictions in the contextualized tweets, split into two error types: same prediction as the original prediction or not. The X-axis represents the days from the original tweet (0=same date as the original tweet) and results are broken on 20-day chunks..

itates temporal analysis in NER for social media. The dataset is diverse topic-wise, as we leveraged weekly trending topics to query tweets and near-duplicated and irrelevant tweets were dropped. To establish baselines on TweetNER7, we fine-tuned standard LMs including a few Twitter-specific LMs. Moreover, we performed a few targeted temporal-related analyses in order to better understand the short-term temporal effect. Finally, we show that self-labeling is not enough to mitigate the temporal-shift and had no noticeable improvement over the baseline vanilla fine-tuning, which further highlights the challenging nature of the dataset.

6 Limitations and Future Work

The TweetNER7 dataset was constructed on English tweets so it is limited to English, as most of the existing NER datasets for social media (Derczynski et al., 2016). In the future we are planning to apply a similar methodology to extend it to languages other than English. Given the dynamic nature of social media, TweetNER7 is designed to study short-term temporal-shift (e.g., monthly) but would not be suitable for analysing longer temporal shifts (e.g., yearly) (Rijhwani and Preotiuc-Pietro, 2020). We selected Twitter as the data source but temporal-shift is a common problem in social media generally. As a future work, we are planning to add more data from other social media platforms as in WNUT17 (Derczynski et al., 2017) to give us more general insights to understand temporal shift phenomena in social media more generally.

Acknowledgements

Jose Camacho-Collados is supported by a UKRI Future Leaders Fellowship.

References

- Oshin Agarwal and Ani Nenkova. 2022. [Temporal effects on pre-trained models for language processing tasks](#). *Transactions of the Association for Computational Linguistics*, 10:904–921.
- Dimosthenis Antypas, Asahi Ushio, Jose Camacho-Collados, Leonardo Neves, Vitor Silva, and Francesco Barbieri. 2022. [Twitter Topic Classification](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Jose Camacho-Collados, Kiamehr Rezaee, Talayeh Riahi, Asahi Ushio, Daniel Loureiro, Dimosthenis Antypas, Joanne Boisson, Luis Espinosa-Anke, Fangyu Liu, Eugenio Martínez-Cámara, et al. 2022. [Tweetnlp: Cutting-edge natural language processing for social media](#). *arXiv preprint arXiv:2206.14774*.
- Nigel Collier and Jin-Dong Kim. 2004. [Introduction to the bio-entity recognition task at JNLPBA](#). In *Proceedings of the International Joint Workshop on Natural Language Processing in Biomedicine and its Applications (NLPBA/BioNLP)*, pages 73–78, Geneva, Switzerland. COLING.
- Marco Del Tredici, Raquel Fernández, and Gemma Boleda. 2019. [Short-term meaning shift: A distributional exploration](#). 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 2069–2075, Minneapolis, Minnesota. Association for Computational Linguistics.
- Leon Derczynski, Kalina Bontcheva, and Ian Roberts. 2016. [Broad Twitter corpus: A diverse named entity recognition resource](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1169–1179, Osaka, Japan. The COLING 2016 Organizing Committee.
- Leon Derczynski, Eric Nichols, Marieke van Erp, and Nut Limsopatham. 2017. [Results of the WNUT2017 shared task on novel and emerging entity recognition](#). In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 140–147, Copenhagen, Denmark. Association for Computational Linguistics.
- 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. Association for Computational Linguistics.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. [OntoNotes: The 90% solution](#). In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 57–60, New York City, USA. Association for Computational Linguistics.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Hang Jiang, Yining Hua, Doug Beeferman, and Deb Roy. 2022. [Annotating the tweekbank corpus on named entity recognition and building nlp models for social media analysis](#). In *Proceedings of the Language Resources and Evaluation Conference*, pages 7199–7208, Marseille, France. European Language Resources Association.
- Angeliki Lazaridou, Adhi Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Tomas Kocisky, Sebastian Ruder, et al. 2021. [Mind the gap: Assessing temporal generalization in neural language models](#). *Advances in Neural Information Processing Systems*, 34:29348–29363.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. [Biobert: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36(4):1234–1240.
- Patrick Lewis, Yuxiang Wu, Linqing Liu, Pasquale Minervini, Heinrich Küttler, Aleksandra Piktus, Pontus Stenetorp, and Sebastian Riedel. 2021. [PAQ: 65 million probably-asked questions and what you can do with them](#). *Transactions of the Association for Computational Linguistics*, 9:1098–1115.
- 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.
- Xiaoya Li, Xiaofei Sun, Yuxian Meng, Junjun Liang, Fei Wu, and Jiwei Li. 2020. [Dice loss for data-imbalanced NLP tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 465–476, Online. Association for Computational Linguistics.
- Yijia Liu, Yi Zhu, Wanxiang Che, Bing Qin, Nathan Schneider, and Noah A. Smith. 2018. [Parsing tweets into Universal Dependencies](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 965–975, New Orleans, Louisiana. Association for Computational Linguistics.
- 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](#). *CoRR*, abs/1907.11692.
- Daniel Loureiro, Francesco Barbieri, Leonardo Neves, Luis Espinosa Anke, and Jose Camacho-collados. 2022. [TimeLMs: Diachronic language models from Twitter](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 251–260, Dublin, Ireland. Association for Computational Linguistics.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. [BERTweet: A pre-trained language model for English tweets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 9–14, Online. Association for Computational Linguistics.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. [Cross-lingual name tagging and linking for 282 languages](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958, Vancouver, Canada. Association for Computational Linguistics.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving language understanding by generative pre-training](#). *Technical report, OpenAI*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*, 1(8):9.
- Shruti Rijhwani and Daniel Preotiuc-Pietro. 2020. [Temporally-informed analysis of named entity recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7605–7617, Online. Association for Computational Linguistics.
- 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, Parramatta, Australia.
- Simone Tedeschi and Roberto Navigli. 2022. **MultiNERD: A multilingual, multi-genre and fine-grained dataset for named entity recognition (and disambiguation)**. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 801–812, Seattle, United States. Association for Computational Linguistics.
- 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 the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Asahi Ushio and Jose Camacho-Collados. 2021. **T-NER: An all-round python library for transformer-based named entity recognition**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 53–62, Online. Association for Computational Linguistics.
- Chih-Hsuan Wei, Yifan Peng, Robert Leaman, Allan Peter Davis, Carolyn J Mattingly, Jiao Li, Thomas C Wieggers, and Zhiyong Lu. 2015. Overview of the biocreative v chemical disease relation (cdr) task. In *Proceedings of the fifth BioCreative challenge evaluation workshop*, volume 14.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi 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 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. Association for Computational Linguistics.
- Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. **LUKE: Deep contextualized entity representations with entity-aware self-attention**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6442–6454, Online. Association for Computational Linguistics.