

Annotating the Tweebank Corpus on Named Entity Recognition and Building NLP Models for Social Media Analysis

Hang Jiang, Yining Hua, Doug Beeferman, Deb Roy

MIT Center for Constructive Communication

75 Amherst St, Cambridge, MA 02139

{hjia42, ninghua, dougb5, dkroy}@mit.edu

Abstract

Social media data such as Twitter messages (“tweets”) pose a particular challenge to NLP systems because of their short, noisy, and colloquial nature. Tasks such as Named Entity Recognition (NER) and syntactic parsing require highly domain-matched training data for good performance. To date, there is no complete training corpus for both NER and syntactic analysis (e.g., part of speech tagging, dependency parsing) of tweets. While there are some publicly available annotated NLP datasets of tweets, they are only designed for individual tasks. In this study, we aim to create Tweebank-NER, an English NER corpus based on Tweebank V2 (TB2), train state-of-the-art (SOTA) Tweet NLP models on TB2, and release an NLP pipeline called Twitter-Stanza. We annotate named entities in TB2 using Amazon Mechanical Turk and measure the quality of our annotations. We train the Stanza pipeline on TB2 and compare with alternative NLP frameworks (e.g., FLAIR, spaCy) and transformer-based models. The Stanza tokenizer and lemmatizer achieve SOTA performance on TB2, while the Stanza NER tagger, part-of-speech (POS) tagger, and dependency parser achieve competitive performance against non-transformer models. The transformer-based models establish a strong baseline in Tweebank-NER and achieve the new SOTA performance in POS tagging and dependency parsing on TB2. We release the dataset and make both the Stanza pipeline and BERTweet-based models available “off-the-shelf” for use in future Tweet NLP research. Our source code, data, and pre-trained models are available at: <https://github.com/social-machines/TweebankNLP>.

Keywords: text annotation, noisy text, NLP toolkit, Twitter, named entity recognition, tokenization, lemmatization, part-of-speech tagging, dependency parsing

1. Introduction

Researchers use text data from social media platforms such as Twitter and Reddit for a wide range of studies including opinion mining, socio-cultural analysis, and language variation. Messages posted to such platforms are typically written in a less formal style than what are found in conventional data sources for NLP models, namely news articles, papers, websites, and books. Processing the noisy and informal language of social media is challenging for traditional NLP tools because such messages are usually short in length and irregular in spelling and structure. In response, the NLP community has been constructing language resources and building NLP pipelines for social media data, especially for Twitter.

Annotating social media language resources is important to the development of NLP tools. Foster et al. (2011) is the one of the earliest attempts to annotate tweets in the Penn Treebank (PTB) format. Following a similar PTB-style convention suggested by Schneider et al. (2013), Kong et al. (2014) created Tweebank V1. However, the PTB annotation guidelines leave many annotation decisions unspecified and are therefore unsuitable for informal and user-generated text. After Universal Dependencies (UD) (Nivre et al., 2016) was introduced to enable consistent annotation across different languages and gen-

res, Liu et al. (2018) introduced a new tweet-based Tweebank V2 in UD, including tokenization, part-of-speech (POS) tags, and (labeled) Universal Dependencies. Besides syntactic annotation, NLP researchers have also annotated tweets on named entities. Ritter et al. (2011) first introduced this English Twitter NER task and found that NER systems trained on the news perform poorly on tweets. Since then, the noisy user-generated text (WNUT) workshop has proposed a few benchmark datasets including WNUT15 (Xu et al., 2015), WNUT16 (Xu et al., 2015), and WNUT17 (Derczynski et al., 2017) for Twitter lexical normalization and named entity recognition (NER). However, these benchmarks are not based upon TB2, which contains high-quality UD annotations. Annotating named entities in TB2 fills a gap in NLP research, allowing researchers to train multi-task learning models in NER, POS tagging, and dependency parsing, and study the linguistic relationship between syntactic labels and named entities in the Twitter domain.

Many researchers have invested in building better NLP pipelines for tokenization, POS tagging, parsing, and NER. The earliest work focuses on Twitter POS taggers (Gimpel et al., 2010; Owoputi et al., 2013) and NER (Ritter et al., 2011). Later, Kong et al. (2014) published TweepoParser on Tweebank V1 to include tokenization, POS tagging, and dependency parsing. Liu et al. (2018) further improved the whole pipeline based on TB2. The current state-of-the-art (SOTA) pipeline

The first two authors contribute equally. YH is also affiliated with Harvard Medical School.

in POS tagging and NER is based on BERT pre-trained on a large number of tweets Nguyen et al. (2020). However, these efforts (1) are often no longer maintained (Ritter et al., 2011; Kong et al., 2014), (2) do not contain publicly available NLP models (e.g., NER, POS tagger) (Nguyen et al., 2020), (3) are written in C/C++ or R with complicated dependencies and installation process (e.g., Twpipe (Liu et al., 2018) and UD-Pipe (Straka et al., 2016)), making them difficult to be integrated into Python frameworks and to be used in an “off-the-shelf” fashion. Many modern NLP tools in Python such as spaCy¹, Stanza (Qi et al., 2020), and FLAIR (Akbik et al., 2019) have been developed for standard NLP benchmarks but have never been adapted to Tweet NLP tasks. In this study, we choose Stanza over other NLP frameworks because (1) the Stanza framework achieves SOTA or competitive performance on many NLP tasks across 66 languages (Qi et al., 2020), (2) Stanza supports both CPU and GPU training and inference while transformer-based models (e.g., BERTweet) need GPU, (3) Stanza shows superior performance against spaCy in our experiments despite slower speeds, (4) Stanza is competitive in speed compared with FLAIR of similar accuracy (Qi et al., 2020), but the FLAIR dependency parser is still under development.

In this paper, we annotate Tweebank V2 on NER to create Tweebank-NER and also build Tweet NLP models based on Stanza and transformer models. We run additional experiments to answer the following questions: (1) How is the quality of the NER annotations? (2) Do NER models trained on existing Twitter NER data perform well on Tweebank-NER? (3) How do Stanza models perform compared with other NLP frameworks on the core Tweet NLP tasks? (4) How do transformer-based models perform compared with traditional models on these tasks? Our contributions are as follows:

- We annotate Tweebank V2, the main treebank for English Twitter NLP tasks, on NER. This annotation not only provides a new benchmark (Tweebank-NER) for Twitter NER but also makes Tweebank a complete dataset for both syntactic tasks and NER, making it suitable for training multi-task learning models in POS tagging, dependency parsing, and NER.
- We leverage the Stanza framework to present an accurate and fast Tweet NLP pipeline called Twitter-Stanza. It includes NER, tokenization, lemmatization, POS tagging, and dependency parsing modules, and it supports both CPU and GPU computation.
- We compare Twitter-Stanza against existing models for each presented NLP task, confirming that Stanza’s simple neural architecture

is effective and suitable for tweets. Among non-transformer models, the Twitter-Stanza tokenizer and lemmatizer achieve SOTA performance on TB2, and its POS tagger, dependency parser, and NER model obtain competitive performance.

- We also train transformer-based models to establish a strong baseline on the Tweebank-NER benchmark and SOTA performance in POS tagging and dependency parsing on TB2. We upload the BERTweet-based NER and POS taggers to the Hugging Face Hub: <https://huggingface.co/TweebankNLP>
- We release our data, models, and code. Our Twitter-Stanza pipeline is highly compatible with Stanza’s Python interface and is simple to use in an “off-the-shelf” fashion. We hope that our Twitter-Stanza and Hugging Face BERTweet models can serve as a convenient NLP tool and a strong baseline for future research and applications of Tweet analytic tasks.

2. Dataset and Annotation Scheme

In this study, we primarily work on the Tweebank V2 dataset and develop its NER annotations through rigorous annotation guidelines. We also evaluate the quality of our annotations, showing that it has a good F1 inter-annotator agreement score.

2.1. Datasets and Annotation Statistics

Tweebank V2 (TB2) (Kong et al., 2014; Liu et al., 2018) is a collection of 3,550 labeled anonymous English tweets annotated in Universal Dependencies. It is a commonly used corpus for the training and fine-tuning of NLP systems on social media texts. A summary of TB2 is shown in Table 1.

Dataset	Train	Dev	Test
Tweets	1,639	710	1,201
Tokens	24,753	11,742	19,112
Avg. token per tweet	15.1	16.6	15.9
Annotated spans	979	425	750
Annotated tokens	1,484	675	1183
Avg. token per span	1.5	1.6	1.6

Table 1: Annotated corpus statistics.

2.2. Annotation Guidelines

We follow the CoNLL 2003 guidelines² to annotate named entities. We are aware that some NER annotations (e.g., English OntoNotes) have more than four classes. We adopt the standard four-class CoNLL 2003 NER guidelines for two reasons. One one hand, adopting a more fine-grained annotation scheme is more challenging for human annotators. The 4-class scheme

¹<https://spacy.io/>

²<https://www.clips.uantwerpen.be/conll2003/ner/>

is already quite challenging for humans since the inter-annotator agreement is low for the MISC class. On the other hand, Tweebank is relatively small, with only 3,550 tweets. An annotation scheme with more classes than that will mean fewer instances per class, and greater difficulty for NER models to learn efficiently. To help annotators understand the guidelines, we provide multiple examples for each rule and ask annotators to read them before the task. Our task focuses on the following four named entities:

- **PER**: persons (e.g., Joe Biden, joe biden, Ben, 50 Cent, Jesus)
- **ORG**: organizations (e.g., Stanford University, stanford, IBM, Black Lives Matter, WHO, Boston Red Sox, Science Magazine, NYT)
- **LOC**: locations (e.g., United States, usa, China, Boston, Bay Area, CA, MT Washington)
- **MISC**: named entities which do not belong to the previous three. (e.g., Chinese, chinese, World Cup 2002, Democrat, Just Do It, Top 10, Titanic, The Shining, All You Need Is Love)

To handle challenges in tweets, we also add requirements consistent with (Ritter et al., 2011): (1) ignore numerical entities (MONEY, NUMBER, ORDINAL, PERCENT), (2) ignore temporal entities (DATE, TIME, DURATION, SET), (3) "At mentions" are not named entities (e.g., allow "Donald Trump" but not @DonaldTrump), (4) #hashtags are not named entities (e.g., allow "BLM" but not "#BLM"), (5) URLs are not named entities (e.g., disallow <https://www.google.com/>).

2.3. Annotation Logistics

We use the Qualtrics platform to design the sequence labeling task and Amazon Mechanical Turk to recruit annotators. We first launch a pilot study, annotate each of the 100 tweets, and discuss tweets with divergent annotations. Based on the pilot study, we develop a series of annotation rules and precautions. During the recruiting process, each annotator is given an overview of annotation conventions and our guidelines, after which they are asked to complete the qualification test. The qualification test consists of 7 tweets that are selected from the pilot study. An annotator must make fewer than 2 errors and not make any significant error in order to pass the qualification test. We consider a significant error to be one in which any URL, @USER, or hashtag is labeled as a named entity; or one in which the PERSON, LOCATION, and ORG categories are confused with each other.

After all tweets have been annotated by at least 3 annotators, we merge the annotation results and create the Tweebank-NER dataset in the BIO format (Ratinov and Roth, 2009). In the merging process, if at least two annotators give the annotation result for a tweet, we use that result as the final annotation. Otherwise, we

discuss and re-annotate the tweet to reach a consensus. We identify 178 span annotations whose three annotations are different from each other and decide their gold annotations collectively by two authors. We find that one of the three annotators' answers is the same as the final annotation for 155 out of the 178 annotations.

2.4. Annotation Quality

We first evaluate the quality of the annotations using a measure of inter-annotator agreement (IAA). For NER, Cohen's Kappa is not the best measure because it needs the number of negative cases, but NER is a sequence tagging task. Therefore, we follow previous work (Hripcsak and Rothschild, 2005; Grouin et al., 2011; Brandsen et al., 2020) to use the token-level pairwise F1 score calculated without the O label as a better measure for IAA in NER (Deleger et al., 2012). In Table 2, we observe that PER, LOC, and ORG have higher F1 agreement than MISC, showing that MISC is more difficult to annotate than the other classes. We also provide the additional Kappa measure ($\kappa = 0.347$) on annotated tokens to provide some insights, although it significantly underestimates IAA for NER. Finally, we calculate the scores by comparing the crowdsourced annotators against our own internal annotations on 100 sampled examples, obtaining a similar F1 score (0.71).

Label	Quantity	F1
PER	777	84.6
LOC	317	74.4
ORG	541	71.9
MISC	519	50.9
Overall	2,154	70.7

Table 2: Number of span annotations per entity type and Inter-annotator agreement scores in pairwise F1.

We analyzed the 178 annotations passed to the merge step, finding that the proportion of each label is 8.4% (LOC), 15.2% (PER), 29.2% (ORG), and 47.2% (MISC). These numbers show that MISC is the most challenging class for human annotators and ORG is also relatively difficult compared to LOC and PER. This confirms the IAA measured in pairwise F1 in Table 2 because the MISC has the lowest F1 (50.9%) and ORG has the second lowest F1 (71.9%).

In the future, we suggest a few ways to improve the annotation quality. The first way is to increase the number of annotators per tweet in both the initial and merge stages. Second, hiring a small number of experienced annotators instead of using crowdsourcing platforms will make the annotations more consistent. Third, adopting a human-in-the-loop approach allows annotators to focus on difficult instances from MISC and ORG, which can reduce the cost and improve the performance of the models at the same time.

3. Methods for NLP Modeling

Stanza is a state-of-the-art and efficient framework for many NLP tasks (Qi et al., 2020; Zhang et al., 2021) and it supports both NER and syntactic tasks. We use Stanza to train NER models as well as syntactic models (tokenization, lemmatization, POS tagging, dependency parsing) on TB2. For more detailed information on Stanza, we refer the readers to the Stanza paper (Qi et al., 2020) and its current website³. We use Twitter GloVe embeddings (Pennington et al., 2014) with 100 dimensions in our experiments and the default parameters in Stanza for training.

Alternative NLP frameworks such as spaCy, FLAIR, transformers, and spaCy-transformers are compared with Stanza. Both spaCy and FLAIR are open-source NLP frameworks for NER and syntactic tasks. Transformers is a library of pre-trained transformer models for NLP and it provides a TokenClassification module⁴, which is adopted for NER and POS tagging. We denote these models as HuggingFace-BERTweet in our experiments. The spaCy-transformers framework provides the spaCy interface to combine pre-trained representations from transformer-based language models and its own NLP models via Hugging Face’s transformers. To train spaCy, we adopt the default NER setting⁵ and the default syntactic NLP pipeline⁶. For FLAIR, we train its NER and syntactic modules with the default settings as well. For spaCy-transformers models, we fine-tune BERTweet-base and XLM-RoBERTa-base language models via spaCy-transformers for NER, POS Tagging, and dependency parsing⁷. We denote them as spaCy-BERTweet and spaCy-XLM-RoBERTa in the paper. BERTweet (Nguyen et al., 2020) is the first public large-scale language model for English tweets based on RoBERTa and XLM-RoBERTa-base is a multilingual version of RoBERTa-base. All transformer-based models show strong performance in Tweet NER and POS tagging (Nguyen et al., 2020). The architecture and training details of the models above can be found at our public repository.

3.1. Named Entity Recognition

In this paper, we adopt the four-class convention to define NER as a task to locate and classify named entities mentioned in unstructured text into four pre-defined categories: PER, ORG, LOC, and MISC (Sang

³<https://stanfordnlp.github.io/stanza/>

⁴<https://github.com/huggingface/transformers/tree/main/examples/legacy/token-classification>

⁵https://github.com/explosion/projects/blob/v3/pipelines/ner_wikiner/configs/default.cfg

⁶https://github.com/explosion/projects/tree/v3/benchmarks/ud_benchmark

⁷https://github.com/explosion/projects/blob/v3/benchmarks/ud_benchmark/configs/transformer.cfg

and De Meulder, 2003). We use the Stanza NER architecture for training and evaluation, which is a contextualized string representation-based sequence tagger (Akbik et al., 2018). This model contains a forward and a backward character-level LSTM language model to extract token-level representations and a BiLSTM-CRF sequence labeler to predict the named entities. We also train the default NER models for SpaCy, FLAIR, HuggingFace-BERTweet, and spaCy-BERTweet for comparison.

3.2. Syntactic NLP Tasks

3.2.1. Tokenization

Tokenizers predict whether a given character in a sentence is the end of a token. The Stanza tokenizer jointly works on tokenization and sentence segmentation, by modeling them as a tagging problem over character sequences. In accordance with previous work (Gimpel et al., 2010; Liu et al., 2018), we focus on the performance in tokenization, as tweets are usually short with a single sentence.

To compare with spaCy, we train a spaCy tokenizer named char_pretokenizer.v1. FLAIR uses spaCy’s tokenizer, so we exclude it from comparison. We also include baselines mentioned in previous work (Kong et al., 2014; Liu et al., 2018). Twokenizer (O’Connor et al., 2010) is a regex-based tokenizer and does not adapt to the UD tokenization scheme. Stanford CoreNLP (Manning et al., 2014), spaCy, and UDPipe v1.2 (Straka and Straková, 2017) are three popular NLP frameworks re-trained on TB2. Twipipe tokenizer (Liu et al., 2018) is similar to UDPipe, but replaces GRU in UDPipe with an LSTM and uses a larger hidden unit number. We do not compare with transformer-based models because they use subword-level tokenization schemes like WordPiece (Wu et al., 2016) and BPE (Sennrich et al., 2015).

3.2.2. Lemmatization

Lemmatization is the process of recovering each word in a sentence to its canonical form. We train the Stanza lemmatizer on TB2, which is implemented as an ensemble model of a dictionary-based lemmatizer and a neural seq2seq lemmatizer. We compare the Stanza lemmatizer against three lemmatizers from spaCy, NLTK, and FLAIR (Table 7). Both NLTK and spaCy lemmatizer are rule-based and use a dictionary to look up the canonical form given a word and its POS tag. The FLAIR lemmatizer is a char-level seq2seq model. We provide gold POS tags for lemmatization.

3.2.3. POS Tagging

POS tagging assigns each token in a sentence a POS tag. We train the Stanza POS tagger, a bidirectional long short-term memory network as the basic architecture to predict the universal POS (UPOS) tags. We ignore the language-specific POS (XPOS) tags because TB2 only contains UPOS tags.

We also train the default POS taggers for SpaCy, FLAIR, HuggingFace-BERTweet, spaCy-BERTweet, spaCy-XLM-RoBERTa. We include performance from existing work in Tweet POS tagging: (1) Stanford CoreNLP tagger, (2) Owoputi et al. (2013)’s word cluster-enhanced greedy tagger, (3) Owoputi et al. (2013)’s word cluster-enhanced tagger with CRF, (4) Ma and Hovy (2016)’s neural tagger, (5) BERTweet-based POS tagger (Nguyen et al., 2020). The first four models were re-trained on the combination of TB2 and UD_English-EWT (Ann Bies, Justin Mott, Colin Warner, Seth Kulick, 2012) training sets, whereas the BERTweet-based tagger was fine-tuned solely on TB2. HuggingFace-BERTweet has the same architecture implementation as Nguyen et al. (2020).

3.2.4. Dependency Parsing

Dependency parsing predicts a syntactic structure for a sentence, where every word in the sentence is assigned a syntactic head that points to either another word in the sentence or an artificial root symbol. Stanza’s dependency parser combines a Bi-LSTM-based deep biaffine neural parser (Dozat and Manning, 2017) and two linguistic features, which can significantly improve parsing accuracy (Qi et al., 2018). Gold-standard tokenization and automatic POS tags are used.

We also re-train spaCy, spaCy-BERTweet, and spaCy-RoBERTa dependency parsers with their default parser architectures⁸. We compare our Stanza models with previous work: (1) Kong et al. (2014)’s graph-based parser with lexical features and word cluster and it uses dual decomposition for decoding, (2) Dozat and Manning (2017)’s neural graph parser with biaffine attention, (3) Ballesteros et al. (2015)’s neural greedy stack LSTM parser, (4) an ensemble model of 20 transition-based parsers (Liu et al., 2018), (5) A distilled graph-based parser of the previous ensemble model (Liu et al., 2018). These models are all trained on TB2+UD_English-EWT. We are aware that Stymne (2020) trained a transition-based uparser (de Lhoneux et al., 2017) on a combination of TB2, UD_English-EWT, and more out-of-domain data (English GUM (Zeldes, 2017), LinES (Ahrenberg, 2007), ParTUT (Sanguinetti and Bosco, 2015)) to further boost model performance, but we do not experiment with this data combination to be consistent with Liu et al. (2018).

4. Evaluation

We train the NER and syntactic NLP models described above with 1) TB2 training data (the default data setting), 2) TB2 training data + extra Twitter data (the combined data setting). For the combined data setting, we add the training and dev sets from other data sources to TB2’s training and dev sets respectively. Specifically, we add WNUT17⁹ (Derczynski

⁸FLAIR and Hugging Face’s transformers do not contain dependency parsing by default.

⁹We map both “group” and “corporation” to “ORG”, and both “creative work” and “product” to “MISC”.

et al., 2017) for NER. For syntactic NLP tasks, we add UD_English-EWT (Ann Bies, Justin Mott, Colin Warner, Seth Kulick, 2012). We pick the best models based on the corresponding dev sets and report their performance on their TB2 test sets. For each task, we compare Stanza models with existing studies and alternative NLP frameworks.

4.1. Performance in NER

Systems	F1
spaCy (TB2)	52.20
spaCy (TB2+W17)	53.89
FLAIR (TB2)	62.12
FLAIR (TB2+W17)	59.08
HuggingFace-BERTweet (TB2)	73.71
HuggingFace-BERTweet (TB2+W17)	74.35
spaCy-BERTweet (TB2)	73.79
spaCy-BERTweet (TB2+W17)	74.15
Stanza (TB2)	60.14
Stanza (TB2+W17)	62.53

Table 3: NER comparison on the TB2 test set in entity-level F1. “TB2” indicates to use the TB2 train set for training. “TB2+W17” indicates to combine TB2 and WNUT17 train sets for training.

4.1.1. Main Findings

The NER experiments presented in Table 3 show that the Stanza NER model (TB2+W17) achieves the best performance among all non-transformer models. At the same time, the Stanza model is up to 75% smaller than the second-best FLAIR model (Qi et al., 2020). For transformer-based approaches, spacy-BERTweet and HuggingFace-BERTweet have close performance to each other. The HuggingFace-BERTweet approach trained on TB2+W17 achieves the highest performance (74.35%) on Tweepbank-NER, establishing a strong benchmark for future research. We also find that combining the training data from both WNUT17 and TB2 improves the performance of spaCy, FLAIR, Stanza, and BERTweet-based models.

4.1.2. Confusion Matrix Analysis

In Figure 1, we plot a confusion matrix for all four entity types and “O”, the label for tokens that do not belong to any of these types. The diagonal and the vertical blue lines are expected because the cells on the diagonal are when the algorithm predicts the correct entity and the vertical line is when the algorithm mistakes an entity for the “O” entity, which is the most common error for NER. We notice that MISC entities are easily mistaken as “O”, which corresponds to the annotation statistics in Table 2, where MISC has the lowest IIA score in pairwise F1. Thus, MISC is the most challenging of the four types for both humans and machines.

Error type	weet example
PER → O	The 50 % Return Method Billionaire Investor Warren Buffet Wishes He Could Use
LOC → O	Getting ready ... @ Pasco Ephesus Seventh - day Adventist Church
ORG → O	#bargains #deals 10.27.10 Guess Who “American Woman” Guhhh deehh you !
MISC → O	RT @USER1508 : Do you ever realize Sounds Live Feels Live Starts this month and just

Table 4: Common mistakes made by the Stanza (W17+TB2) NER model for each error type. “X → O” means the model predicts X entity to be O by mistake. Green and red texts are gold annotations of the corresponding type in each row. Correct predictions are in bold green and gold annotations missed by the model are in bold red.

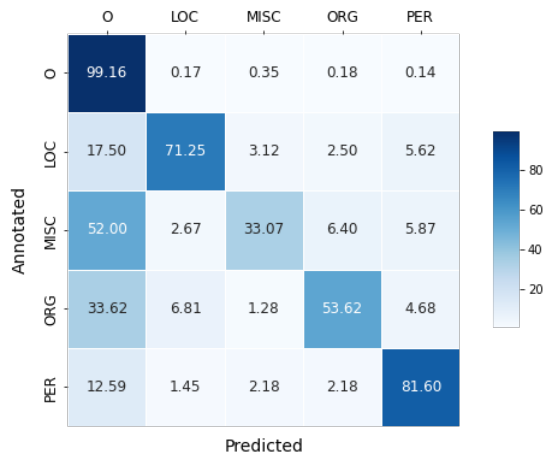


Figure 1: Confusion matrix generated by the Stanza (TB2+W17) model to show percentages for each combination of predicted and true entity types.

4.1.3. Error Analysis

We identify the most common error types that Stanza (TB2+W17)¹⁰ makes on the TB2 test in Figure 1: predicting PER, LOC, ORG, MISC to be O. We pick some representative examples for each error type, shown in Table 4. For the *PER* → *O* error type, every first letter in a word is capitalized and the model fails to recognize the famous investor “Warren Buffet” in such a context. We find that person entities with abbreviations (e.g., “GD” for “G-dragon”), lower case (e.g., “kush” for “Kush”), or irregular contextual capitalization are challenging to the NER system. For the *LOC* → *O* error type, the structure to encode location is complicated and sometimes interrupted by the parentheses and dashes (e.g., “- day Adventist Church”). In this case, it is caused by the fact that “Seventh-day” is tokenized into three words in TB2. For the *ORG/MISC* → *O* examples, “Guess Who” is a rock band and “Sounds Live Feels Live” is a concert tour by Australian pop-rock band 5 Seconds of Summer. These named entities tend to contain common English verbs with their first letters capitalized. It is difficult to annotate them correctly if the model does not have access to world and

¹⁰We pick Stanza over BERTweet for error analysis because we only aimed to publish the Stanza pipeline at the beginning. We eventually publish the BERTweet models too.

domain knowledge. Our analysis points to the future Twitter NER research to introduce text perturbations into training and to encode commonsense knowledge into NER modeling.

Training data	TB2	WNUT17	F1 Drop
spaCy	52.20	44.93	7.27↓
FLAIR	62.12	55.11	7.01↓
HgFace-BERTweet	73.71	59.43	14.28↓
spaCy-BERTweet	73.79	60.77	13.02↓
Stanza	60.14	56.40	3.74↓

Table 5: Comparison among NER models trained on TB2 vs. WNUT17 on TB2 test in entity-level F1. “HgFace” stands for “HuggingFace”.

4.1.4. NER Models Trained on WNUT17

We train spaCy, FLAIR, Stanza, HuggingFace-BERTweet, and spaCy-BERTweet NER models on the four-class version of WNUT17 and evaluate their performance on the TB2 test. In Table 5, we compare the performance of these models trained on WNUT17 against the ones trained on TB2. We show that the performance of all the models drops significantly if we use the pre-trained model from WNUT17, meaning the Tweepbank-NER dataset is still challenging for current NER models and can be used as an additional benchmark to evaluate NER models.

4.2. Performance in Syntactic NLP Tasks

Apart from NER, we train and evaluate Stanza models for tokenization, lemmatization, POS tagging, and dependency parsing by leveraging TB2 and UD.English-EWT. For each task, we compare our models against previous work on the TB2 test set.

4.2.1. Tokenization Performance

In Table 6, we observe that the Stanza model trained on TB2 outperforms Twpipe tokenizer, the previous SOTA model, and it achieves slightly higher performance than the spaCy tokenizer. We also find that blending TB2 and UD.English-EWT for training brings down the tokenization performance slightly. This is probably because the data source of UD.English-EWT, which is collected from weblogs, newsgroups, emails, reviews, and Yahoo! Answers, represents a different dialect from Twitter English.

System	F1
Twokenizer	94.6
Stanford CoreNLP	97.3
UDPipe v1.2	97.4
Twpipe	98.3
spaCy (TB2)	98.57
spaCy (TB2+EWT)	95.57
Stanza (TB2)	98.64
Stanza (TB2+EWT)	98.59

Table 6: Tokenizer comparison on the TB2 test set. “TB2” indicates to use TB2 for training. “TB2+EWT” indicates to combine TB2 and UD English-EWT for training. Note that the first four results are rounded to one decimal place by Liu et al., (2018).

4.2.2. Lemmatization Performance

None of the previous Twitter NLP work reports the lemmatization performance on TB2. As shown in Table 7, the Stanza model outperforms the other two rule-based (NLTK and spaCy) and one neural (FLAIR) baseline approaches on TB2. This is not surprising because the Stanza ensemble lemmatizer makes good use of both ruled-based dictionary lookup and seq2seq learning. Similar to what we observe in the tokenization experiments, the combined data setting brings down the performance of FLAIR and Stanza models.

System	F1
NLTK	88.23
spaCy	85.28
Flair (TB2)	96.18
Flair (TB2+EWT)	84.54
Stanza (TB2)	98.25
Stanza (TB2+EWT)	85.45

Table 7: Lemmatization results on the TB2 test set. “TB2” is to use TB2 for training. “TB2+EWT” is to combine TB2 and UD English-EWT for training.

4.2.3. POS Tagging Performance

As shown in Table 8, HuggingFace-BERTweet (TB2) replicates the SOTA performance from BERTweet (Nguyen et al., 2020) in terms of accuracy. When trained on the combined data of TB2 and UD_English-EWT, HuggingFace-BERTweet achieves the best performance (95.38%) in accuracy out of all the models. Compared to HuggingFace-BERTweet, spaCy-transformers models perform worse. The spaCy-XLM-RoBERTa trained on TB2 is 1.3% lower than Nguyen et al. (2020). We conjecture that the difference is mainly due to the implementations of the POS tagging layer between spaCy and HuggingFace-BERTweet, which is the same as Nguyen et al. (2020). Among the non-transformer models, Stanza achieves competitive performance compared with Owoputi et al. (2013)’s tagger with CRF (93.53% vs. 94.6%). Stanza outperforms all other non-transformer baselines includ-

ing Stanford CoreNLP, spaCy, FLAIR, and Ma and Hovy (2016). Interestingly, we observe that adding UD_English-EWT for training improves the performance of non-transformer models and HuggingFace-BERTweet but slightly brings down the performance of spaCy-transformers models.

System	UPOS
Stanford CoreNLP	90.6
Owoputi et al. (2013) (greedy)	93.7
Owoputi et al. (2013) (CRF)	94.6
Ma and Hovy (2016)	92.5
BERTweet (Nguyen et al., 2020)	95.2
spaCy (TB2)	86.72
spaCy (TB2+EWT)	88.84
FLAIR (TB2)	87.85
FLAIR (TB2+EWT)	88.19
HuggingFace-BERTweet (TB2)	95.21
HuggingFace-BERTweet (TB2+EWT)	95.38
spaCy-BERTweet (TB2)	87.61
spaCy-BERTweet (TB2+EWT)	86.31
spaCy-XLM-RoBERTa (TB2)	93.90
spaCy-XLM-RoBERTa (TB2+EWT)	93.75
Stanza (TB2)	93.20
Stanza (TB2+EWT)	93.53

Table 8: POS Tagging comparison in accuracy on the TB2 test set. “TB2” is to use TB2 for training. “TB2+EWT” is to combine TB2 and UD English-EWT for training. Please note that the first five results are rounded to one decimal place by Liu et al., (2018).

4.2.4. Dependency Parsing Performance

For dependency parsing experiments, spaCy-XLM-RoBERTa (TB2) achieves the SOTA performance (Table 9), surpassing Liu et al. (2018) (Ensemble) by 0.42% in UAS¹¹. Besides that, the Stanza parser achieves the same UAS score and has a close LAS score (−0.3%) compared to this best non-transformer performance (UAS 82.1% + LAS 77.9%) reported by the distilled parser. As Liu et al. (2018) mentioned, the ensemble model is 20 times larger in size compared to the Stanza parser, although the former performs better. Finally, we confirm that the combination of TB2 and UD_English-EWT training sets boost the performance for non-transformer models (Liu et al., 2018). The data combination brings down the performance of transformer-based models, which is consistent with our observations in tokenization, POS tagging, and dependency parsing.

5. Conclusion

In this paper, we introduce four-class named entities to Tweepbank V2, a popular Twitter dataset within the Universal Dependencies framework, creating a new

¹¹It is difficult to compare their LAS with ours due to the difference in decimal places.

System	UAS	LAS
Kong et al. (2014)	81.4	76.9
Dozat et al. (2017)	81.8	77.7
Ballesteros et al. (2015)	80.2	75.7
Liu et al. (2018) (Ensemble)	83.4	79.4
Liu et al. (2018) (Distillation)	82.1	77.9
spaCy (TB2)	66.93	58.79
spaCy (TB2 + EWT)	72.06	63.84
spaCy-BERTweet (TB2)	76.32	71.72
spaCy-BERTweet (TB2+EWT)	76.18	69.28
spaCy-XLM-RoBERTa (TB2)	83.82	79.39
spaCy-XLM-RoBERTa (TB2+EWT)	81.02	75.43
Stanza (TB2)	79.28	74.34
Stanza (TB2 + EWT)	82.10	77.60

Table 9: Dependency parsing comparison on the TB2 test set. “TB2” indicates to use TB2 for training. “TB2+EWT” indicates to combine TB2 and UD English-EWT for training. Note that the first six results are rounded to one decimal place by Liu et al., (2018).

NER benchmark called `Tweebank-NER`. We evaluate our annotations and observe good inter-annotator agreement score in pairwise F1 for NER annotation. We train Twitter-specific NLP models (NER, tokenization, lemmatization, POS tagging, dependency parsing) on the dataset with Stanza and compare our models against existing work and NLP frameworks. Our Stanza models show SOTA performance on tokenization and lemmatization and competitive performance in NER, POS tagging, and dependency parsing on TB2. We also train BERT-based methods to establish a strong benchmark on `Tweebank-NER` and achieve SOTA performance in POS tagging and dependency parsing on TB2. Finally, we publish our dataset and release the Stanza pipeline `Twitter-Stanza`, which is easy to download and use with Stanza’s Python interface. We also release the BERTweet-based NER and POS tagger on Hugging Face Hub. We hope that our research not only contributes annotations to an important dataset but also enables other researchers to use off-the-shelf NLP models for social media analysis.

6. Acknowledgements

We would like to thank Alan Ritter, Yuhui Zhang, Zifan Lin, and anonymous reviewers, who gave precious advice and comments on our paper. We also want to thank John Bauer and Yijia Liu for answering questions related to Stanza and Twpipe. Finally, we would like to thank MIT Center for Constructive Communication for funding our research.

7. Bibliographical References

Ahrenberg, L. (2007). LinES: An English-Swedish parallel treebank. In *Proceedings of the 16th Nordic Conference of Computational Linguistics (NODALIDA 2007)*, pages 270–273, Tartu, Estonia, May. University of Tartu, Estonia.

Akbik, A., Blythe, D., and Vollgraf, R. (2018). Contextual string embeddings for sequence labeling. In *Proceedings of the 27th international conference on computational linguistics*, pages 1638–1649.

Akbik, A., Bergmann, T., Blythe, D., Rasul, K., Schweter, S., and Vollgraf, R. (2019). Flair: An easy-to-use framework for state-of-the-art nlp. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 54–59.

Ballesteros, M., Dyer, C., and Smith, N. A. (2015). Improved transition-based parsing by modeling characters instead of words with lstms. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 349–359.

Branden, A., Verberne, S., Lambers, K., Wansleeben, M., Calzolari, N., Béchet, F., Blache, P., Choukri, K., Cieri, C., Declerck, T., et al. (2020). Creating a dataset for named entity recognition in the archaeology domain. In *Conference Proceedings LREC 2020*, pages 4573–4577. The European Language Resources Association.

de Lhoneux, M., Shao, Y., Basirat, A., Kiperwasser, E., Szymne, S., Goldberg, Y., and Nivre, J. (2017). From raw text to universal dependencies-look, no tags! In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 207–217.

Deleger, L., Li, Q., Lingren, T., Kaiser, M., Molnar, K., et al. (2012). Building gold standard corpora for medical natural language processing tasks. In *AMIA Annual Symposium Proceedings*, volume 2012, page 144. American Medical Informatics Association.

Leon Derczynski, et al., editors. (2017). *Proceedings of the 3rd Workshop on Noisy User-generated Text*, Copenhagen, Denmark, September. Association for Computational Linguistics.

Dozat, T. and Manning, C. D. (2017). Deep biaffine attention for neural dependency parsing. *The International Conference on Learning Representations*.

Dozat, T., Qi, P., and Manning, C. D. (2017). Stanford’s graph-based neural dependency parser at the conll 2017 shared task. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 20–30.

Foster, J., Cetinoglu, O., Wagner, J., Le Roux, J., Hogan, S., Nivre, J., Hogan, D., and Van Genabith, J. (2011). #hardtoparse: Pos tagging and parsing the twitterverse.

Gimpel, K., Schneider, N., O’Connor, B., Das, D., Mills, D., Eisenstein, J., Heilman, M., Yogatama, D., Flanigan, J., and Smith, N. A. (2010). Part-of-speech tagging for twitter: Annotation, features, and experiments. Technical report, Carnegie-Mellon Univ Pittsburgh Pa School of Computer Science.

Grouin, C., Rosset, S., Zweigenbaum, P., Fort, K., Galibert, O., and Quintard, L. (2011). Proposal for an extension of traditional named entitites: from guide-

- lines to evaluation, an overview. In *5th Linguistics Annotation Workshop (The LAW V)*, pages 92–100.
- Hripcsak, G. and Rothschild, A. S. (2005). Agreement, the f-measure, and reliability in information retrieval. *Journal of the American medical informatics association*, 12(3):296–298.
- Kong, L., Schneider, N., Swayamdipta, S., Bhatia, A., Dyer, C., and Smith, N. A. (2014). A dependency parser for tweets. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1001–1012.
- Liu, Y., Zhu, Y., Che, W., Qin, B., Schneider, N., and Smith, N. A. (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.
- Ma, X. and Hovy, E. (2016). End-to-end sequence labeling via bi-directional lstm-cnns-crf. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, page 1064–1074.
- Manning, C. D., Surdeanu, M., Bauer, J., Finkel, J. R., Bethard, S., and McClosky, D. (2014). The stanford corenlp natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, pages 55–60.
- Nguyen, D. Q., Vu, T., and Nguyen, A. T. (2020). Bertweet: A pre-trained language model for english tweets. *Association for Computational Linguistics*.
- Nivre, J., De Marneffe, M.-C., Ginter, F., Goldberg, Y., Hajic, J., Manning, C. D., McDonald, R., Petrov, S., Pyysalo, S., Silveira, N., et al. (2016). Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1659–1666.
- O'Connor, B., Krieger, M., and Ahn, D. (2010). Tweetmotif: Tweetmotif: Exploratory search and topic summarization for twitter. In *Fourth International AAAI Conference on Weblogs and Social Media*.
- Owoputi, O., O'Connor, B., Dyer, C., Gimpel, K., Schneider, N., and Smith, N. A. (2013). Improved part-of-speech tagging for online conversational text with word clusters. In *Proceedings of the 2013 conference of the North American chapter of the association for computational linguistics: human language technologies*, pages 380–390.
- Pennington, J., Socher, R., and Manning, C. D. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Qi, P., Dozat, T., Zhang, Y., and Manning, C. D. (2018). Universal dependency parsing from scratch. *CoNLL 2018 UD Shared Task*.
- Qi, P., Zhang, Y., Zhang, Y., Bolton, J., and Manning, C. D. (2020). Stanza: A python natural language processing toolkit for many human languages. pages 101–108.
- Ratinov, L. and Roth, D. (2009). Design challenges and misconceptions in named entity recognition. In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009)*, pages 147–155.
- Ritter, A., Clark, S., Etzioni, O., et al. (2011). Named entity recognition in tweets: An experimental study. In *Proceedings of the 2011 conference on empirical methods in natural language processing*, pages 1524–1534.
- Sang, E. F. and De Meulder, F. (2003). Introduction to the conll-2003 shared task: Language-independent named entity recognition. *arXiv preprint cs/0306050*.
- Sanguinetti, M. and Bosco, C. (2015). Partut: The turin university parallel treebank. In *Harmonization and development of resources and tools for italian natural language processing within the parli project*, pages 51–69. Springer.
- Schneider, N., O'Connor, B., Saphra, N., Bamman, D., Faruqui, M., Smith, N. A., Dyer, C., and Baldridge, J. (2013). A framework for (under) specifying dependency syntax without overloading annotators. *Linguistic Annotation Workshop & Interoperability with Discourse*.
- Sennrich, R., Haddow, B., and Birch, A. (2015). Neural machine translation of rare words with subword units. pages 1715–1725.
- Straka, M. and Straková, J. (2017). Tokenizing, pos tagging, lemmatizing and parsing ud 2.0 with udpipe. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 88–99.
- Straka, M., Hajic, J., and Straková, J. (2016). Udpipeline: trainable pipeline for processing conll-u files performing tokenization, morphological analysis, pos tagging and parsing. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 4290–4297.
- Stymne, S. (2020). Cross-lingual domain adaptation for dependency parsing. In *Proceedings of the 19th Workshop on Treebanks and Linguistic Theories*, pages 62–69.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. (2016). Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Wei Xu, et al., editors. (2015). *Proceedings of the Workshop on Noisy User-generated Text*, Beijing,

China, July. Association for Computational Linguistics.

Zeldes, A. (2017). The GUM corpus: Creating multilayer resources in the classroom. *Language Resources and Evaluation*, 51(3):581–612.

Zhang, Y., Zhang, Y., Qi, P., Manning, C. D., and Languotz, C. P. (2021). Biomedical and clinical english model packages for the stanza python nlp library. *Journal of the American Medical Informatics Association*, 28(9):1892–1899.

8. Language Resource References

Ann Bies, Justin Mott, Colin Warner, Seth Kulick. (2012). *English Web Treebank*. Philadelphia: Linguistic Data Consortium, ISLRN 230-396-178-102-3.