

Cross-Register Projection for Headline Part of Speech Tagging

Adrian Benton, Hanyang Li, Igor Malioutov

Bloomberg

731 Lexington Ave

New York, NY 10022 USA

{abenton10, hli762, imalioutov}@bloomberg.net

Abstract

Part of speech (POS) tagging is a familiar NLP task. State of the art taggers routinely achieve token-level accuracies of over 97% on news body text, evidence that the problem is well-understood. However, the register of English news headlines, “headlines”, is very different from the register of long-form text, causing POS tagging models to underperform on headlines. In this work, we automatically annotate news headlines with POS tags by projecting predicted tags from corresponding sentences in news bodies. We train a multi-domain POS tagger on both long-form and headline text and show that joint training on both registers improves over training on just one or naïvely concatenating training sets. We evaluate on a newly-annotated corpus of over 5,248 English news headlines from the Google sentence compression corpus, and show that our model yields a 23% relative error reduction per token and 19% per headline. In addition, we demonstrate that better headline POS tags can improve the performance of a syntax-based open information extraction system. We make POSH, the POS-tagged Headline corpus, available to encourage research in improved NLP models for news headlines.

1 Introduction

News headlines were identified as a unique register of written language at least as far back as [Straumann \(1935\)](#), and the term *headlines* was coined specifically to refer to the unnatural register used in headlines. Hallmarks of headlines include frequent omission of articles and auxiliary verbs, stand-alone nominals and adverbials, and infinitival forms of the main verb when referring to the future tense ([Mårdh, 1980](#)). In spite of the clear differences in syntax between news headlines and bodies, the NLP community has invested little effort in building headline-specific models for predicting syntactic annotations such as part of speech (POS) tags and dependency parse trees. This limits the use

of syntax-based models on headlines. For example, the Open Domain Information Extraction (Open IE) system, PredPatt ([White et al., 2016](#)), cannot be expected to perform well on headlines, if existing models cannot accurately tag headlines with POS tags and dependency relations. More broadly, headline processing provides an important signal for downstream applications such as summarization ([Bambrick et al., 2020](#)), sentiment/stance classification ([Strapparava and Mihalcea, 2007](#); [Ferreira and Vlachos, 2016](#)), semantic clustering ([Wities et al., 2017](#); [Laban et al., 2021](#)), and information retrieval ([Lee et al., 2010](#)), among other applications.

In this work we take a first step toward developing strong NLP models for headlines by focusing on improving POS taggers. We propose a projection technique (inspired by work on cross-lingual projection ([Yarowsky et al., 2001](#))) to generate silver training data for each headline based on POS tags from the lead sentence of the news body. To facilitate evaluation, we crowdsource and manually adjudicate gold POS tag annotations for 5,248 headlines from the Google sentence compression corpus ([Filippova and Altun, 2013](#)). To the best of our knowledge, this is the first headline dataset with gold-annotated POS tags.

We evaluate a range of neural POS taggers trained on long-form text, silver-labeled headlines, and the concatenation of these datasets, and find that multi-domain models which train a separate decoder layer per domain outperform models trained on data from either headlines or long-form text. We finally show that more accurate POS taggers directly lead to more precise Open IE extractions, yielding an extracted tuple precision of 56.9% vs. 29.8% when using a baseline POS tagger trained on the English Web Treebank alone. We release our gold-annotated data as POSH, the POS-tagged Headline corpus, at <https://zenodo.org/record/5495668>.

Japanese Firms Push Posh Car Showrooms
 Twinkies are Back but Smaller
 World Briefing — Africa: Angola: New Bid To End War
 Ebola Fear Stalks Home Hunt for Quarantined Now Released
 Madonna Addicted to Sweat Dance Plugs Toronto Condos: Mortgages

Figure 1: Examples of headlines exhibiting register-specific phenomena: article/auxiliary omission, multiple “decks” or independent segments, and garden paths following from the need to fit a large amount of information in a small space.

Contributions

1. A gold-annotated evaluation set of over 5,000 English headlines to promote stronger NLP tools for news headlines.
2. An error analysis, confirming that many of the errors made by taggers trained on body text are due to headline-specific phenomena, such as article and copula omission. This demonstrates that existing POS taggers are not effective at headline POS tagging.
3. By training models on headlines tagged using a simple projection of POS tags from the corresponding news body’s lead sentence, we can outperform a tagger trained on gold-labeled long-form text data. Multi-domain taggers that are trained on both gold-annotated long-form text and silver-labeled headlines outperform taggers trained on either, reducing the relative token error rate by 23%.
4. Finally, we demonstrate that more accurate headline POS tags translate to more precise tuple extractions in a state-of-the-art multilingual Open IE system.

The proposed projection technique and models can also be applied to other sequence tagging tasks such as chunking or named entity recognition. These are also applicable in other domains where one has access to long-form sentence and a parallel reduced sentence (e.g., simplified English (Coster and Kauchak, 2011) or writing from English as a second language students (Dahlmeier et al., 2013)).

2 Data

We rely on two main sources of data for training and evaluating headline POS taggers: version 2.6 of the Universal Dependencies English Web Treebank

(EWT; Silveira et al. 2014)¹ and headlines from the Google sentence compression corpus (GSC).

In this work, we only consider the UD POS tag set. UD has become a dominant tag set for tagging and parsing, with many treebanks available across many languages (Zeman et al., 2017). In addition, we chose to annotate headlines with UD tags as the coarser granularity made it easier for non-experts to label. We leave experiments with finer representation granularity as future work, but do not expect the choice of representation to alter the underlying conclusions.

2.1 Headline Evaluation Set

We construct the GSC headline evaluation set (GSCh) by sampling uniformly at random from headlines in the GSC where the headline was a (possibly non-contiguous) subsequence of the associated lead sentence. Headlines were first tokenized using the Stanford CoreNLP PTBTokenizer (Manning et al., 2014b) with default settings, and were annotated by a pool of six annotators who were all proficient in English. Each headline was annotated independently by three annotators from this pool. Annotators were warm-started with POS tags generated by an EWT-trained BiLSTM model and were trained to follow the UD 2.0 POS tagging guidelines. When in doubt, annotators referred to similar examples in the UD 2.6 EWT. Full annotation instructions and a sample of the interface are given in Appendix A.

We gave the annotators feedback by monitoring their performance on a set of 250 unique test examples randomly inserted as tasks (annotated by the authors). We identified common mistakes on

¹https://github.com/UniversalDependencies/UD_English-EWT/tree/r2.6




Figure 2: Unigram POS tag distribution for the EWT and GSCh datasets.

these questions and provided feedback to annotators throughout the annotation process, after every 500-1000 examples.

Annotators achieved unanimous agreement on 70.7% of headlines, and the majority vote label sequence agreed with our test examples 88% of the time. We annotated an additional 200 samples from the unanimously agreed set, and found that 97% of labeled examples completely agreed with our manual annotation. We thus focused on those headlines without unanimous agreement, and manually reviewed the majority vote tag sequence for all headlines where there was no unanimous agreement (1534 out of 5248 examples). We corrected at least one POS tag in 24.97% of these headlines. Common mistakes include labeling a nominalized verb as VERB, labeling a NOUN as ADJ when it is a nominal modifier, confusing ADP and SCONJ, and labeling “to” as ADP instead of PART.

2.2 Headline vs. Body POS Tag Distribution

The unigram tag distribution for EWT and GSCh exhibit clear, expected differences (Figure 2). For instance, the lack of DET is due to article dropping and the lack of ADV and increase in frequency of other open class tags follows from the function that news headlines serve: maximizing the relevance of the article to a reader under strict space constraints (Dor, 2003). Also, PROPN is more frequent in headlines, while PRON is almost non-existent. This naturally follows from the fact that headlines do not offer context from which an antecedent for the PRON can be found. Type-token ratio is higher for GSCh than EWT (0.273 vs. 0.076), and GSCh headlines tend to be shorter than EWT sentences

(7 vs. 15 mean token length).

2.3 Auxiliary Datasets

We use the EWT as our baseline training set as the corpus contains over 250,000 annotated words of web text drawn from a wide variety of sources. In addition, we use the Revised English News Text Treebank (Bies et al., 2015), converted to Universal Dependencies using the Stanford CoreNLP Toolkit (Manning et al., 2014a), GUM (Zeldes, 2017), and English portions of the LinES (Ahrenberg, 2015) as additional training data for a subset of taggers.

In addition to the GSCh, we collected POS tag annotations for 500 additional headlines: 271 sampled uniformly at random from the GSC and 229 from The New York Times Annotated Corpus (NYT; Sandhaus 2008). NYT headlines were restricted to those with 4-12 words (10th-90th percentile of length distribution). No subsequence constraint was imposed on any of these headlines, as we used this set to evaluate how well models performed on general headlines.

3 Methods

Here we describe the architecture of the POS taggers and our approach for learning a headline POS tagger without direct supervision.

3.1 Models

In all experiments, we use a bidirectional recurrent neural network with gated recurrent units, followed by a linear-chain conditional random field layer (BiGRU-CRF). We consider two flavors of tagger: *non-contextual* and *contextual*. The contextual taggers use the BERT (Devlin et al., 2019) base uncased model pretrained on Wikipedia and the Google books corpus as a token encoder,² and continue to finetune model weights during tagger training. We represent each word by the embedding of its initial subword as done in Devlin et al. (2019). The non-contextual taggers use 50-dimensional GloVe word embeddings (Pennington et al., 2014) concatenated with a 50-dimensional cased character embedding generated by a single-layer BiGRU.³ Non-contextual taggers use two-layer BiGRUs whereas contextual taggers only con-

²We use the BERT implementation and model exposed by <https://huggingface.co/>.

³<http://nlp.stanford.edu/data/glove.6B.zip>

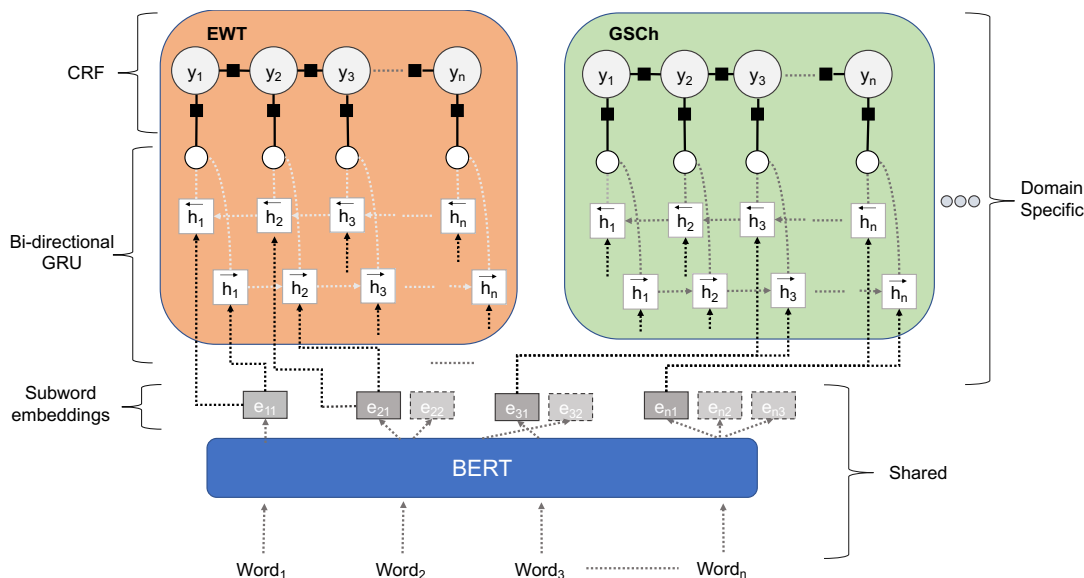


Figure 3: Schematic of the architecture for a multi-domain BERT POS tagger. Word embeddings are given as the first subword embedding from an uncased BERT base language model. The non-contextual variant is identical except it uses GloVe embeddings as input, a 2-layer RNN shared across domains, and only the tag projection and neural CRF layers are domain-specific. Choice of decoder is governed by the corpus the example comes from.

tain single-layer BiGRUs. This is the same architecture used by [Lample et al. \(2016\)](#) for named entity recognition except with a GRU instead of a LSTM as the recurrent unit.⁴ Viterbi decoding is used to generate predictions for all taggers.

Even though large, pretrained language models have become a standard solution for addressing various NLP tasks, we also evaluate RNN taggers as they contain many fewer parameters and may suffice for POS tagging. This is also a new task, and the data come from a domain which BERT was not explicitly trained on (news headlines). Therefore, it is not immediately clear that using a BERT POS tagger will outperform an RNN tagger.

3.1.1 Multi-Domain Taggers

In our experiments, we also train taggers on data from multiple domains. For the purpose of modeling, we designate each dataset as belonging to a separate *domain*, even though the critical distinction between datasets may be due to a variation in register, not the subject matter. In addition to training mixed-domain taggers on a simple concatenation of datasets, we consider tagger variants with domain-specific decoder layers.

Non-contextual, multi-domain taggers have

⁴These settings of number of layers, embedding, and hidden layer dimensionality were found to perform well in preliminary experiments. In addition, we saw little difference in performance between the two types of recurrent cells, and GRUs contain fewer parameters.

domain-specific weights for the tag projection layer and CRF, but share the bidirectional GRU encoder across domains. Contextual taggers share the BERT encoder across domains but learn domain-specific weights for the BiGRU-CRF layers. This multi-domain architecture is a simpler version of that proposed by [Peng and Dredze \(2017\)](#), since all models are trained to perform one task, POS tagging (Figure 3).

3.2 Projecting POS Tags

We construct a silver-labeled training set for GSC headlines by selecting articles from the GSC corpus where the headline is a (possibly non-contiguous) subsequence of the lead sentence after lowercasing. After excluding GSCh examples, we find a total of 44,591 out of 200,000 headlines that satisfy this condition, which we divide into 70% for the training fold and 30% for a validation fold. We run inference with the selected EWT-trained tagger on these lead sentences, and assign the predicted tag for each word in the lead sentence to the corresponding aligned word in the headline. These silver-labeled headlines are used both for training and model selection (GSCproj). See Figure 4 for an example of this projection.

This approach was inspired by work on projecting syntactic and morphological annotations across languages ([Yarowsky et al., 2001](#); [Fossum and Abney, 2005](#); [Täckström et al., 2013](#); [Buys and Botha,](#)




Figure 4: Example of tag projection: tags of lead words are projected onto the corresponding headline word.

2016; Eskander et al., 2020), except that we impose no type constraints on the projected tags in our alignment. Although Filippova and Altun (2013) state that there are relatively few headlines which are extractive summaries of the article’s lead sentence, there were enough of such headlines in the GSC corpus for us to construct a silver-labeled training set of the same order of magnitude as the EWT (12,543 sentences vs. 31,213 headlines).

4 Experiments

Models are trained on either EWT, GSCproj, or both training sets. We evaluate all models on the GSCh test set according to token and headline accuracy, unless otherwise stated. We also consider training multi-domain models on the three auxiliary datasets as well as EWT and GSCproj, to explore the value added by training on more long-form text data. For the EWT-trained baseline taggers, we insert a final period to GSCh headlines before running inference so as to mitigate the mismatch between the training and test data.

4.1 Training Details

We train all models using the Adam optimizer with $\beta_1 = 0.99, \beta_2 = 0.999$. All BiGRU layers are 100-dimensional. For each architecture, we perform a random search with a budget of 10 runs over dropout rate in $[0.0, 0.4]$ and number of epochs in $[2, 6]$. The fixed dropout rate is applied to the hidden layers of the BiGRU, throughout the BERT encoder, and before the tag projection layer. We sample the exponent of the base learning rate (base 10) for non-contextual taggers uniformly at random from $[-5.5, -1.0]$ and for contextual taggers from $[-5.5, -4.0]$. Each model was trained using a single V100 GPU and training time varied from one to eight hours based on the size of the training set and number of training epochs.

We retrain each model three times with different random seeds, and select training parameters based on mean GSCproj validation token accuracy

across seeds (for models trained solely on EWT, models are selected based on the EWT validation set). When training jointly on only the EWT and auxiliary datasets, we select a model over hyper-parameters and decoder head based on token accuracy on the GSCproj validation set. We report test performance for the model found to be most accurate on the validation set. As is recommended by Søgaard et al. (2014), we use bootstrap sampling tests to test for statistical significance at the $p < 0.01$ level.⁵

4.2 Extrinsic Open IE Evaluation

To demonstrate the potential impact of explicit modeling of headlines, we consider the effect of POS tag quality on the downstream performance of an Open IE system. We experiment with a re-implementation of the state-of-the-art Open IE system PredPatt (White et al., 2016), which extracts propositions, tuples of predicate followed by n extracted arguments, using syntactic rules over a UD dependency parse of the source sentence. For the parser, we employ an arc-hybrid transition-based dependency parser using a bi-directional LSTM neural network (Kiperwasser and Goldberg, 2016), trained on data from sections of 2-21 of the WSJ Penn Treebank.⁶ The parser takes POS tags as input. In our experiments, we substitute the predicted POS tags from the models under consideration and evaluate their downstream impact on tuple extraction quality.

To evaluate and compare the systems, we perform a typological error analysis for a randomly sampled 100 sentence subset of GSCh where the output of the OpenIE systems on EWT and EWT+GSCproj-based predictions differ.

Setting	Train set	Non-contextual		BERT	
		Token	Headline	Token	Headline
projected from lead	EWT	90.16	53.22	92.08	60.77
	EWT	91.71*	59.03*	94.00*	68.39*
	GSCproj	92.14*	60.37*	93.56*	66.46*
shared	EWT + GSCproj	91.94*	59.38*	93.86*	67.76*
	EWT + GSCproj	92.26*	60.34*	93.93*	68.27*
multi-domain	EWT + Aux	92.02*	59.03*	92.69*	63.13
multi-domain	EWT + GSCproj + Aux	93.19* [†]	64.01* [†]	94.00*	68.10*

Table 1: Token and headline percent accuracy of each POS tagger on GSCh. *Aux* refers to all auxiliary datasets described in Section 2.3. *projected from lead* refers to performance attained by predicted tags projected from the lead sentence to headline. *shared* and *multi-domain* refer to whether the decoder is shared across registers, or decoder weights are specific to each register. The best performance for each type of encoder is in bold. Statistically significant performance over the EWT and GSCproj models are denoted by * and [†], respectively.

5 Results and Discussion

Table 1 displays the test performance for POS taggers on GSCh. The selected hyperparameters and validation performance of each model is given in Appendix B. Both non-contextual and contextual language models benefit from training on the silver-labeled GSCproj. Although the BERT EWT tagger achieves 95.40% accuracy on the EWT test set, which is on par with strong models such as Stanza (Qi et al., 2018), the performance degrades on GSCh, achieving only 92.08% accuracy. We found that this is in line with performance of an EWT-trained tagger applied to data from other domains: 94.34% token accuracy on the UD-converted English news treebank test set vs. 97.64% for a POS tagger trained on the Revised English News Text Treebank. This underscores the fact that the language and POS tag distribution in news headlines is substantially different from long-form text. However, by training on GSCproj alone, one can improve absolute accuracy by 1.48% at the token level and 5.69% at the headline level.

We find that projecting predicted tags from the BERT tagger from the lead sentence onto the headline performs surprisingly well. However, this is not a realistic scenario at inference time as few articles have headlines that are subsequences of the lead sentence. We find that a multi-domain BERT tagger trained on five corpora achieves the same performance on GSCh as the POS tags projected from the lead sentence.⁷

⁵Using batches of 1056 headlines (20% of the data) sampled with replacement from GSCh and 2,000 replications.

⁶The parser achieves 94.47 % UAS and 93.50 % LAS on section 23 of WSJ, when using gold POS tags.

⁷Note that both the EWT+GSCproj and

Non-contextual performance Non-contextual taggers also benefit from training on GSCproj. In fact, a multi-domain non-contextual tagger that is trained on EWT and GSCproj along with the auxiliary datasets outperforms a BERT tagger that was only trained on EWT by over 1% absolute token accuracy. Although adding auxiliary long-form text datasets improves tagger accuracy over only training on EWT, the addition of the silver-labeled GSCproj training set leads to more than a 1% absolute improvement in token accuracy. One interesting difference between contextual and non-contextual taggers is that the domain-agnostic (*shared*) BERT tagger benefits from the concatenation of EWT and GSCproj training sets, while the non-contextual tagger does not. We posit that this is because the contextual token embeddings learned by the BERT model naturally discriminate between the headline and long-form text registers. In effect, the BERT tagger can learn a register-specific sense of each word simply due to distributional differences between the GSCh and EWT registers.

Performance on general headlines Table 2 shows the performance of trained BERT models on the NYT+GSC unconstrained headline evaluation set described in Section 2.3. Training on GSCproj alone improves the performance on the GSCh subset, but the model achieves similar token accuracy on the NYT subset. However, the multi-domain model trained on all available domains yields an absolute increase of 1.67% token accuracy on the NYT fold, even though none of the

EWT+GSCProj+Aux multi-domain BERT taggers also achieve statistically significantly better token accuracy than the GSCproj model, but only at the $p < 0.05$ level.

Train	Test	% Accuracy	
		Token	Head
EWT	NYT+GSCh	91.88	58.20
	NYT	92.14	60.26
	GSCh	91.59	56.46
GSCProj	NYT+GSCh	92.82	61.20
	NYT	92.03	56.77
	GSCh	93.59	65.68
EWT+ GSCProj+ Aux	NYT+GSCh	93.64	65.20
	NYT	93.81	66.38
	GSCh	93.43	63.84

Table 2: Performance of models on the general headline evaluation set. The "EWT+GSCProj+Aux" model is the multi-domain model, and evaluation set predictions are made using the GSCProj head. Performance is reported separately for the NYT and GSCh subsets, as well as their union.

training domains explicitly contain NYT headlines. This suggests that training on projected tags also improves headline POS tagging for headlines that are not strictly subsequences of the lead, or are drawn from a different wire, such as the New York Times.

5.1 Training with Gold Supervision

In the above experiments, we assume no gold in-domain supervision, only training on projected tags. To better understand how well a model performs with gold supervision, we also train a single-domain contextual tagger using the gold GSCh tags using 5-fold cross validation: training on 60%, validating on 20%, and testing on the remaining 20%. This tagger, trained on 3,149 gold headlines, achieves 93.55% token and 66.83% headline CV test accuracy, which is similar to that achieved by training on the 31,213 GSCproj training set (93.56%). A model CV-trained on 3,149 projected tag headlines instead achieves 92.73% token and 62.73% headline accuracy.

5.2 Error Analysis

The multi-domain EWT+GSCproj tagger improves over the EWT baseline by remedying errors related to the idiosyncracies of headlines. One such error is the frequent confusion of NOUN and PROPEN tags by the baseline model (Figure 5). This is to be expected as determiners are frequently omitted in headlines, removing a clear signal of whether a word is a PROPEN. Take this prediction from the baseline tagger:

- (1) Firefighter accused of setting fires
*PROPEN VERB ADP VERB NOUN

The second set of remedied errors are a product of the way that some verb phrases are formed in headlines, omitting the copula and using the t_0 VERB construction to signal future tense. These phenomena cause the baseline tagger to make particularly egregious blunders, such as tagging proper nouns as AUX when preceding a non-finite VERB:

- (2) Hideo Kojima working on a new game
PROPEN *AUX VERB ADP DET ADJ NOUN

or VERB if it precedes t_0 VERB:

- (3) Centrelink to review procedures
*VERB AUX VERB NOUN

We initially posited that this was due to the prior over label sequences learned by the CRF layer for the EWT register trumping the evidence. However, when we retrain the EWT model without a CRF layer it only achieved 91.90% accuracy on GSCh, performing worse than including a CRF layer.

The multi-domain tagger also remedies a smaller set of errors involving lexical ambiguity; for instance, confusing a nominal for a VERB:

- (4) Bosc faces spell on sidelines
PROPEN VERB *VERB ADP NOUN

Although the multi-domain model occasionally confuses ADP for SCONJ, after reviewing these examples, we find that many of the errors are due to noisy annotations where a word was incorrectly annotated as ADP instead of SCONJ. We also notice annotation inconsistency in constructions such as **(set|engaged|surprised|thrilled)** t_0 where the gold label should be ADJ, not VERB.

5.3 Extrinsic Evaluation

In the extrinsic Open IE-based evaluation, we performed an error analysis of tuples extracted by PredPatt using BERT EWT vs. EWT+GSCproj tags from 100 randomly selected headlines (which contained at least one tuple extracted from each and for which the extractions differed). Extracted tuples are labeled by two annotators who achieved an inter-annotator agreement rate of 90.0% before remedying discrepancies. Each extracted tuple is labeled for validity, and if invalid, the type of the salient error. 114 and 116 tuples are extracted by the baseline and multi-domain models, respectively.




Figure 5: Confusion matrix for the EWT baseline (left) and the difference between the EWT+GSCproj and EWT tagger confusion matrices (right). In the right plot, purple cells correspond to cases where the EWT+GSCproj model made a prediction more frequently than the EWT model, and vice-versa for green cells. Zero count cells are not annotated and the diagonal counts (correct predictions) are underlined.




Figure 6: Error types in OpenIE tuple extractions given the POS tags predicted by the EWT (inner donut) and EWT+GSCproj (outer donut) BERT taggers.

The types of salient error we focus on are: 1. an argument is mis-attached to the predicate; 2. the argument extraction is incomplete; 3. a core argument is missing; 4. the predicate is malformed; or 5. the predicate is incomplete (Table 3). Our typology is informed by previous work on Open IE evaluations (Schneider et al., 2017; Bhardwaj et al., 2019).

There are two striking differences between tuples extracted using the baseline EWT and EWT +

GSCproj POS tags. First, the multi-domain model results in far more precise extractions than the baseline, achieving a precision of 56.9% vs. 29.8%. Second, the majority of errors resulting from the baseline tagger are due to malformed predicates. Unlike error types such as “incomplete argument” or “core argument (attachment)”, these errors are particularly egregious, as the essential meaning of the extracted tuple is corrupted. Many instances of malformed predicates are due to part of the subject being improperly placed in the predicate, a byproduct of the baseline tagger’s tendency to identify words in second position or preceding “to” as AUX or VERB.

6 Related Work

“Headlines” has been identified as a unique register at least since Straumann (1935). Hallmarks of English headlines include the omission of articles and auxiliary verbs and using the infinitival form of verbs for future events. Subsequent work has found that, even within the English language, the syntax of news headlines varies as a function of publication (Mårdh, 1980; Siegal and Connolly, 1999), time period (Vanderbergen, 1981; Schneider, 2000;

Error Type	Example Extraction
Argument Misattachment	Joliet man accused of shooting [_{A1} 55 dogs] [_P released from] [_{A2} hospital]
Incomplete Argument	[_{A1} Schoolground knife] and weapon attacks have [_P doubled in] [_{A2} 10 years]
Missing Core Argument	[_{A1} Gene therapy] [_P boosts] hope for the deaf
Malformed Predicate	[_{A1} Newcomers] will need to learn [_P about military]
Bad Sub-predicate	[_P Affymetrix to acquire] [_{A1} Panomics]
	[_{A1} Maoists] [_P set] [_{A2} liquor unit] afire

Table 3: OpenIE error types. Extracted tuples encoded as $\langle [P \text{ Predicate}], [A_1 \text{ Argument 1}], [A_2 \text{ Argument 2}] \rangle$.

Afful, 2014), and country (Ehineni, 2014). In this work, we primarily consider headlines from the GSC, drawn from over ten thousand English language news sites in 2012, and leave investigating syntactic variation as future work.

Most NLP work on news headlines has focused on the problem of headline generation (Banko et al., 2000; Rush et al., 2015; Takase et al., 2016; Tan et al., 2017; Takase and Okazaki, 2019) or summarization (Filippova and Altun, 2013). There has also been work on training headline classification models to label headlines by their expressed emotion (Kozareva et al., 2007; Oberländer et al., 2020), stance with respect to an issue (Ferreira and Vlachos, 2016), framing of/bias towards a political issue (Gangula et al., 2019; Liu et al., 2019), or the category or value of the news article (di Buono et al., 2017). However, we are not aware of work to develop headline-specific models for predicting traditional linguistic annotations such as POS tags or dependency trees, although there has been work on adapting machine translation models to headlines (Ono, 2003). As we show in this work, and has been observed in the past (Filippova and Altun, 2013), a model trained on non-headline-domain data will not necessarily perform well on headlines.

Our work is analogous to work on constructing POS taggers or syntactic parsers for less traditional domains such as tweets (Owoputi et al., 2013; Kong et al., 2014; Liu et al., 2018). Unlike this line of work, we did not have to craft a unique tag set for headlines, as headlines are not rife with the typographical errors and atypical constructions of tweets. At the same time, we do not presume a gold-annotated headline training set, but use projection to construct a silver-labeled training set. Data is annotated data purely for the purpose of evaluation.

7 Conclusion

This work is a first step towards developing stronger NLP tools for news headlines. We show that training a tagger on headlines with projected POS tags results in a far stronger model than taggers trained on gold-annotated long-form text. This suggests that more expensive syntactic annotations, such as dependency trees, may also be reliably projected onto headlines, obviating the need for gold dependency annotations when training a headline parser.

Although this work is focused on learning strong headline POS taggers, the projection technique described here can be adapted to train other strong headline sequence taggers; for example, training a headline chunker or named entity tagger on IOB tags projected from the lead sentence. Projection could potentially be applied to generate silver-labeled data for other domains such as simplified English (e.g., aligned sentences from simplified to original Wikipedia (Coster and Kauchak, 2011)) and other languages.

8 Acknowledgments

Many members of the Bloomberg AI group gave helpful feedback on early drafts of this work. In particular, this work benefited immensely from discussions with Amanda Stent, as well as aesthetic feedback and figure-tinkering from Ozan İrsoy.

References

- Isaac Afful. 2014. A diachronic study of the np structure in ghanaian newspaper editorials. *Journal of Advances in Linguistics*, 5:555–565.
- Lars Ahrenberg. 2015. Converting an english-swedish parallel treebank to universal dependencies. In *Third International Conference on Dependency Linguistics (DepLing 2015)*, Uppsala, Sweden, August 24–26, 2015, pages 10–19. Association for Computational Linguistics.

- Joshua Bambrick, Minjie Xu, Andy Almonte, Igor Malioutov, Guim Perarnau, Vittorio Selo, and Iat Chong Chan. 2020. [NSTM: Real-time query-driven news overview composition at Bloomberg](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 350–361, Online. Association for Computational Linguistics.
- Michele Banko, Vibhu O Mittal, and Michael J Witbrock. 2000. Headline generation based on statistical translation. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, pages 318–325.
- Sangnie Bhardwaj, Samarth Aggarwal, and Mausam Mausam. 2019. [CaRB: A crowdsourced benchmark for open IE](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6262–6267, Hong Kong, China. Association for Computational Linguistics.
- Ann Bies, Justin Mott, and Colin Warner. 2015. [English news text treebank: Penn Treebank revised LDC2015T13](#).
- Maria Pia di Buono, Jan Šnajder, Bojana Dalbelo Bašić, Goran Glavaš, Martin Tutek, and Natasa Milic-Frayling. 2017. Predicting news values from headline text and emotions. In *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*, pages 1–6.
- Jan Buys and Jan A. Botha. 2016. [Cross-lingual morphological tagging for low-resource languages](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1954–1964, Berlin, Germany. Association for Computational Linguistics.
- William Coster and David Kauchak. 2011. Simple english wikipedia: a new text simplification task. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 665–669.
- Daniel Dahlmeier, Hwee Tou Ng, and Siew Mei Wu. 2013. Building a large annotated corpus of learner english: The nus corpus of learner english. In *Proceedings of the eighth workshop on innovative use of NLP for building educational applications*, pages 22–31.
- 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.
- Daniel Dor. 2003. On newspaper headlines as relevance optimizers. *Journal of pragmatics*, 35(5):695–721.
- Taiwo Oluwaseun Ehineni. 2014. A syntactic analysis of lexical and functional heads in nigerian english newspaper headlines. *International Journal of Linguistics*, 6(5):9.
- Ramy Eskander, Smaranda Muresan, and Michael Collins. 2020. Unsupervised cross-lingual part-of-speech tagging for truly low-resource scenarios. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4820–4831.
- William Ferreira and Andreas Vlachos. 2016. Emergent: a novel data-set for stance classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: Human language technologies*, pages 1163–1168.
- Katja Filippova and Yasemin Altun. 2013. Overcoming the lack of parallel data in sentence compression. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1481–1491.
- Victoria Fossum and Steven Abney. 2005. [Automatically inducing a part-of-speech tagger by projecting from multiple source languages across aligned corpora](#). In *Second International Joint Conference on Natural Language Processing: Full Papers*.
- Rama Rohit Reddy Gangula, Suma Reddy Duggenpudi, and Radhika Mamidi. 2019. Detecting political bias in news articles using headline attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 77–84.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. [Simple and accurate dependency parsing using bidirectional LSTM feature representations](#). *Transactions of the Association for Computational Linguistics*, 4:313–327.
- Lingpeng Kong, Nathan Schneider, Swabha Swayamdipta, Archna Bhatia, Chris Dyer, and Noah A Smith. 2014. A dependency parser for tweets. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1001–1012.
- Zornitsa Kozareva, Borja Navarro, Sonia Vázquez, and Andrés Montoyo. 2007. Ua-zbsa: a headline emotion classification through web information. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 334–337.
- Philippe Laban, Lucas Bandarkar, and Marti A Hearst. 2021. News headline grouping as a challenging nlu task. In *Proceedings of the 2021 Conference of the*

- North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3186–3198.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. [Neural architectures for named entity recognition](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California. Association for Computational Linguistics.
- Yeha Lee, Hun-young Jung, Woosang Song, and Jong-Hyeok Lee. 2010. Mining the blogosphere for top news stories identification. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 395–402.
- Siyi Liu, Lei Guo, Kate Mays, Margrit Betke, and Derry Tanti Wijaya. 2019. Detecting frames in news headlines and its application to analyzing news framing trends surrounding us gun violence. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 504–514.
- 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.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014a. [The Stanford CoreNLP natural language processing toolkit](#). In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, Baltimore, Maryland. Association for Computational Linguistics.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014b. [The Stanford CoreNLP natural language processing toolkit](#). In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60.
- Ingrid Mårdh. 1980. *Headlines: On the grammar of English front page headlines*, volume 58. Liberläromedel/Gleerup.
- Laura Ana Maria Oberländer, Evgeny Kim, and Roman Klinger. 2020. Goodnewseveryone: A corpus of news headlines annotated with emotions, semantic roles, and reader perception. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1554–1566.
- Kenji Ono. 2003. Translation of news headlines. In *Proceedings of MT Summit IX*.
- Olutobi Owoputi, Brendan O’Connor, Chris Dyer, Kevin Gimpel, Nathan Schneider, and Noah A Smith. 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.
- Nanyun Peng and Mark Dredze. 2017. [Multi-task domain adaptation for sequence tagging](#). In *Proceedings of the 2nd Workshop on Representation Learning for NLP*, pages 91–100, Vancouver, Canada. Association for Computational Linguistics.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Peng Qi, Timothy Dozat, Yuhao Zhang, and Christopher D. Manning. 2018. [Universal Dependency parsing from scratch](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 160–170, Brussels, Belgium. Association for Computational Linguistics.
- Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 379–389.
- Evan Sandhaus. 2008. The new york times annotated corpus. *Linguistic Data Consortium, Philadelphia*, 6(12):e26752.
- Kristina Schneider. 2000. The emergence and development of headlines in british newspapers. *English Media Texts, Past and Present: Language and Textual Structure*, 80:45.
- Rudolf Schneider, Tom Oberhauser, Tobias Klatt, Felix A. Gers, and Alexander Löser. 2017. [Analysing errors of open information extraction systems](#). In *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*, pages 11–18, Copenhagen, Denmark. Association for Computational Linguistics.
- Allan M Siegal and William G Connolly. 1999. *The New York Times manual of style and usage*. Three Rivers Press (CA).
- Natalia Silveira, Timothy Dozat, Marie-Catherine de Marneffe, Samuel Bowman, Miriam Connor, John Bauer, and Christopher D. Manning. 2014. A gold standard dependency corpus for English. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014)*.

- Anders Søgaard, Anders Johannsen, Barbara Plank, Dirk Hovy, and Héctor Martínez Alonso. 2014. What’s in a p-value in nlp? In *Proceedings of the eighteenth conference on computational natural language learning*, pages 1–10.
- Carlo Strapparava and Rada Mihalcea. 2007. Semeval-2007 task 14: Affective text. In *Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007)*, pages 70–74.
- Heinrich Straumann. 1935. *Newspaper headlines: A study in linguistic method*. London, Allen.
- Oscar Täckström, Dipanjan Das, Slav Petrov, Ryan McDonald, and Joakim Nivre. 2013. Token and type constraints for cross-lingual part-of-speech tagging. *Transactions of the Association for Computational Linguistics*, 1:1–12.
- Sho Takase and Naoaki Okazaki. 2019. Positional encoding to control output sequence length. 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 3999–4004.
- Sho Takase, Jun Suzuki, Naoaki Okazaki, Tsutomu Hirao, and Masaaki Nagata. 2016. Neural headline generation on abstract meaning representation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1054–1059.
- Jiwei Tan, Xiaojun Wan, and Jianguo Xiao. 2017. From neural sentence summarization to headline generation: a coarse-to-fine approach. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 4109–4115.
- AM Simon Vanderbergen. 1981. *The Grammar of Headlines in the Times: 1870-1970*, volume 95. AWLSK.
- Aaron Steven White, Drew Reisinger, Keisuke Sakaguchi, Tim Vieira, Sheng Zhang, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2016. [Universal Decompositional Semantics on Universal Dependencies](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1713–1723, Austin, Texas. Association for Computational Linguistics.
- Rachel Wities, Vered Shwartz, Gabriel Stanovsky, Meni Adler, Ori Shapira, Shyam Upadhyay, Dan Roth, Eugenio Martínez-Cámara, Iryna Gurevych, and Ido Dagan. 2017. A consolidated open knowledge representation for multiple texts. In *Proceedings of the 2nd workshop on linking models of lexical, sentential and discourse-level semantics*, pages 12–24.
- David Yarowsky, Grace Ngai, and Richard Wicentowski. 2001. [Inducing multilingual text analysis tools via robust projection across aligned corpora](#). In *Proceedings of the First International Conference on Human Language Technology Research*.
- Amir Zeldes. 2017. [The GUM corpus: Creating multilayer resources in the classroom](#). *Language Resources and Evaluation*, 51(3):581–612.
- Daniel Zeman, Martin Popel, Milan Straka, Jan Hajic, Joakim Nivre, Filip Ginter, Juhani Luotolahti, Sampo Pyysalo, Slav Petrov, Martin Potthast, Francis Tyers, Elena Badmaeva, Memduh Gokirmak, Anna Nedoluzhko, Silvie Cinkova, Jan Hajic jr., Jaroslava Hlavacova, Václava Kettnerová, Zdenka Uresova, Jenna Kanerva, Stina Ojala, Anna Misislä, Christopher D. Manning, Sebastian Schuster, Siva Reddy, Dima Taji, Nizar Habash, Herman Leung, Marie-Catherine de Marneffe, Manuela Sanguinetti, Maria Simi, Hiroshi Kanayama, Valeria dePaiva, Kira Droganova, Héctor Martínez Alonso, Çağrı Çöltekin, Umut Sulubacak, Hans Uszkoreit, Vivien Macketanz, Aljoscha Burchardt, Kim Harris, Katrin Marheinecke, Georg Rehm, Tolga Kayadelen, Mohammed Attia, Ali Elkahky, Zhuoran Yu, Emily Pitler, Saran Lertpradit, Michael Mandl, Jesse Kirchner, Hector Fernandez Alcalde, Jana Strnadová, Esha Banerjee, Ruli Manurung, Antonio Stella, Atsuko Shimada, Sookyoung Kwak, Gustavo Mendonca, Tatiana Lando, Rattima Nitisaroj, and Josie Li. 2017. [Conll 2017 shared task: Multilingual parsing from raw text to universal dependencies](#). In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 1–19, Vancouver, Canada. Association for Computational Linguistics.

A Annotation Guidelines

The annotation UI is shown in Figure 7. See below for full instructions for the annotation task. When in doubt, annotators consulted the UD POS tagging guidelines at <https://universaldependencies.org/u/pos/> and similar sentences in the EWT.

Annotation Guidelines for POS Tagging Headlines

For each task, you will be shown a headline along with POS tags from the Universal Dependencies (UD) 2.0 schema. These tags are automatically assigned by a trained model.

Each word in the headline has its assigned POS tag displayed underneath. Click on any word to correct its tag. POS tags that you changed will be marked with a red border. If all of the words are assigned the correct tag, be sure to select the "All POS tags are correct" checkbox, otherwise you will not be able to click Submit and proceed to the next task.

You should assign POS tags according to the directions given in the UD 2.0 guidelines. Be sure to read these guidelines thoroughly before beginning annotation. When in doubt, you can search for similar sentences in the English Web Treebank here: http://bionlp-www.utu.fi/dep_search/ (select "English (UDv2.0)" and note their parts of speech).

Tricky Cases

There are a few classes of mistakes our model makes consistently. There are also a few classes of examples that may be difficult for humans to decide which is the correct POS tag. Below are the most common cases we encountered.

NOUN vs. PROPN Our model will often incorrectly label NOUN tokens as PROPN, because headlines often do not precede nouns with a determiner. These are easy errors for humans to fix, but just be aware that they are common. Example:

- (5) Haye hopes fight is only postponed

"Haye" should be tagged as a PROPN in this headline, not NOUN.

Compound Nouns From the UD 2.0 guidelines: "A noun modifying another noun to form a compound noun is given the tag NOUN not ADJ." Example:

- (6) But air traffic controllers in Baghdad have no record of the flights, which supposedly took off between July 2004 and July 2005.

In this sentence, "air", "traffic", and "controllers" should all be labeled as NOUN. When in doubt, defer to what the typical dictionary POS tag for a word would be.

Multiword Proper Nouns All tokens constituting a multiword name that overall functions like a proper noun should be tagged as PROPN, even if each constituent word would typically be given another POS tag. Examples:

- (7) United Airlines
PROPN PROPN
- (8) The Yellow Pages
PROPN PROPN PROPN

These are typically names of organizations or people. This also goes for titles preceding a PROPN. Examples:

- (9) US Rep. Lamborn
PROPN PROPN PROPN
- (10) State Senator Stewart Greenleaf
PROPN PROPN PROPN PROPN
- (11) Former President Bush
ADJ PROPN PROPN

In the above, "Former" is tagged as an ADJ, since it is not considered part of the title. One major exception is the possessive marker, "'s", which is tagged as PART, a participle, even if it is part of an organization's name. Examples:

- (12) Pakistan People 's Party
PROPN PROPN PART PROPN
- (13) Garden Grove seniors invited to Mother
PROPN PROPN NOUN VERB ADP PROPN
's Day event
PART PROPN NOUN

Some corner cases:

Tokens that consist of all digits should be labeled as NUM (e.g., the 7 in Windows 7). Acronyms of proper nouns should be labeled as PROPN: USA, NATO, HBO. When in doubt, refer to similar sentences in the English Web Treebank.

Currencies Tagging currencies can be tricky. If the currency denotes the denomination of a sum of money, it should be marked as a SYM:



Figure 7: Example of the annotation UI initially (left) and after selecting a POS tag to correct (right). The instructions are hidden.

(14) Rs. 500 in Pakistan
SYM NUM ADP PROP

(15) \$ 500 in USA
SYM NUM ADP PROP

(16) kr 500 in Denmark
SYM NUM ADP PROP

A currency that is referred to generically should be labeled as a NOUN:

(17) Yuan was little changed against the dollar
NOUN AUX ADV VERB ADP DET NOUN

(18) these countries could put the dollar
DET NOUN AUX VERB DET NOUN
under intense pressure
ADP ADJ NOUN

(19) the Japanese government holds dollar
DET ADJ NOUN VERB NOUN
reserves of approximately \$ 1
NOUN ADP ADJ SYM NUM
trillion
NUM

Hyphenated Compounds and Multiword Expressions Multiword expressions are phrases made up of at least two words whose meaning cannot be directly inferred from its constituent words. A canonical example is the phrase "kick the bucket", which means "to die", and can take the place of a verb in a sentence. Each word in a multiword expression should be assigned the POS tag that is typical for its usage. Example:

(20) let the cat out of the bag
VERB DET NOUN ADP ADP DET NOUN

Similarly, hyphenated expressions like "drive-through" are split by punctuation into separate tokens ("drive", "-", "through"). POS tags should be

assigned to each constituent token independently. Example:

(21) drive - through
VERB PUNCT ADP

When in doubt, defer to the typical POS tag for a token rather than the context.

Passive Constructions Be aware of passive constructions when the subject of the verb is omitted.

(22) Tourism Industry is least affected

(23) Grand - standing is ignored .

In these cases, "affected" and "ignored" should both be labeled as VERB, although one might be tempted to label them as ADJ.

Copula is AUX Copulas (is/are/some inflection of to be) are always labeled as AUX, even when used as the main verb in a sentence. Example:

(24) What are catch shares ?

"are" should be tagged as AUX here, not VERB.

Tagging "than"/"for" SCONJ vs. ADP Words like "than" and "for" may be particularly difficult to tag. When these words are followed by a noun phrase, they are likely ADP. When they are followed by a clause, they are likely SCONJ. Examples:

(25) Meiring is more than an
PROPN AUX ADJ ADP DET
embarrassment .
NOUN PUNCT

- (26) The Kashmir issue generates far more
 DET PROP N NOUN VERB ADV ADJ
 terrorism , and even the threat
 NOUN PUNCT CCONJ ADV DET NOUN
 of nuclear war , than Iraq
 ADP ADJ NOUN PUNCT SCONJ PROP N
 ever did .
 ADV VERB PUNCT
- (27) Now the Electoral Commission is
 ADV DET PROP N PROP N AUX
 refusing to punish people for mere
 VERB PART VERB NOUN ADP ADJ
 past Baath Party membership .
 ADJ PROP N PROP N NOUN PUNCT

B Model Validation Performance

Here we report the selected hyperparameters and validation performance for models presented in Table 1. Note that validation performance is over the silver-labeled GSCproj dataset, which means that it is likely an overestimate of the performance on the gold validation set. The one exception to this is the EWT baseline tagger, where validation performance is over the EWT validation set. Table 4 gives the selected hyperparameters for each of these models along with their validation performance. For the EWT+Aux tagger, validation performance is reported for the EWT decoder for both non-contextual and BERT taggers, as this decoder achieved the highest token accuracy on GSCproj.

Encoder Type	Model	LR	DR	Epochs	% Token Acc
Non-contextual	EWT	3.49e-03	0.0438	2	93.52 (± 0.0549)
	GSCproj	3.12e-04	0.265	5	93.70 (± 0.0618)
	EWT+GSCproj[shared]	1.29e-03	0.315	5	93.83 (± 0.0462)
	EWT+GSCproj[multi]	2.95e-04	0.0763	5	93.38 (± 0.128)
	EWT+Aux[multi]	6.43e-05	0.118	6	90.44 (± 0.120)
	EWT+GSCproj+Aux[multi]	8.71e-05	0.327	6	93.78 (± 0.0224)
BERT	EWT	1.77e-05	0.262	6	95.58 (± 0.0519)
	GSCproj	5.22e-06	0.0519	6	95.66 (± 0.0523)
	EWT+GSCproj[shared]	4.27e-06	0.191	5	95.91 (± 0.0403)
	EWT+GSCproj[multi]	6.71e-06	0.276	5	95.93 (± 0.0494)
	EWT+Aux[multi]	3.57e-06	0.244	6	93.35 (± 0.0303)
	EWT+GSCproj+Aux[multi]	1.73e-05	0.164	4	95.72 (± 0.0616)

Table 4: Hyperparameters (*LR*: base learning rate, *DR*: dropout rate, and number of training *Epochs*) and validation mean % token accuracy (\pm standard deviation) for models presented in Table 1. Standard deviation and mean token accuracy are computed over three training runs with different random seeds.