

Do CoNLL-2003 Named Entity Taggers Still Work Well in 2023?

Shuheng Liu and Alan Ritter

School of Interactive Computing

Georgia Institute of Technology

sliu775@gatech.edu, alan.ritter@cc.gatech.edu

Abstract

The CoNLL-2003 English named entity recognition (NER) dataset has been widely used to train and evaluate NER models for almost 20 years. However, it is unclear how well models that are trained on this 20-year-old data and developed over a period of decades using the same test set will perform when applied on modern data. In this paper, we evaluate the generalization of over 20 different models trained on CoNLL-2003, and show that NER models have very different generalization. Surprisingly, we find no evidence of performance degradation in pre-trained Transformers, such as RoBERTa and T5, even when fine-tuned using decades-old data. We investigate why some models generalize well to new data while others do not, and attempt to disentangle the effects of temporal drift and overfitting due to test reuse. Our analysis suggests that most deterioration is due to temporal mismatch between the pre-training corpora and the downstream test sets. We found that four factors are important for good generalization: model architecture, number of parameters, time period of the pre-training corpus, in addition to the amount of fine-tuning data. We suggest current evaluation methods have, in some sense, underestimated progress on NER over the past 20 years, as NER models have not only improved on the original CoNLL-2003 test set, but improved even more on modern data. Our datasets can be found at https://github.com/ShuhengL/ac12023_conllpp.

1 Introduction

The progress of natural language processing (NLP) is typically measured using performance metrics like accuracy or F_1 score on public test sets. For instance, the top line in Figure 1 shows the steady improvement of selected models on the CoNLL-2003 English named entity recognition (NER) test set (Tjong Kim Sang and De Meulder, 2003) over the course of 15 years (2005-2020) as measured by published F_1 scores.

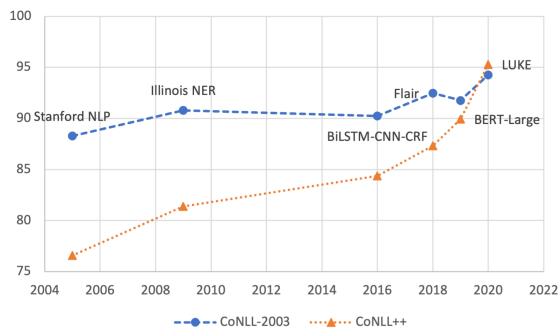


Figure 1: Progress on NER from 2005-2020, as measured using published F_1 scores on the CoNLL-2003 (data from 1996) and CoNLL++ (data from 2020) English NER test set. The gap between the two grows smaller as time passes by, showing improved generalization of models developed over time.

However, these scores are all calculated using the same publicly available test set, which raises several questions. One concern is how much of this progress is actually due to *adaptive overfitting*, i.e. over-estimating performance by reusing the same test set, as opposed to genuine improvement (Recht et al., 2019; Roelofs et al., 2019; Gorman and Bedrick, 2019). In addition, there is also the issue of *temporal drift* as training data ages, which can negatively impact performance on modern data (Rijhwani and Preotiuc-Pietro, 2020; Agarwal and Nenkova, 2022; Luu et al., 2022).

Performance degradation is a significant concern in applications that use NER, such as text deidentification (Morris et al., 2022), relation extraction (Zhong and Chen, 2021), linking entities to a knowledge base (De Cao et al., 2022), etc. However, continuously annotating, training, and evaluating new models on new data is not always possible. NER models that are trained on decades-old data and evaluated on heavily-used public development and test sets may struggle to perform well on modern data, which highlights the need to consider these factors when assessing performance.

Dataset	Time	# Tokens	LOC	MISC	ORG	PER	# Unique Tokens	Avg Sentence Length
CoNLL-2003	Dec. 1996	46,435	1668	702	1661	1617	9,489	12.67
CoNLL++	Dec. 2020	46,587	1128	697	1201	981	8,115	23.64

Table 1: Statistics of CoNLL-2003 test set and our CoNLL++. We report the publication time of the articles, the numbers of four different types of entities, the number of tokens, unique tokens and average number of tokens per sentence.¹

To understand how well NER works when models have been developed over 20 years using the same dataset, we created a new test set called CoNLL++. We closely modeled CoNLL++ after the CoNLL-2003 test set, using news articles from 2020 instead of 1996, as in the original dataset. We carefully controlled for other variables, making results on the two datasets as comparable as possible, with the exception of the time frame. An example of an annotated sentence from CoNLL++ is shown below:

AMBASSADOR O
TO O
THE O
UNITED I-ORG
NATIONS I-ORG
: O
LINDA I-PER
THOMAS-GREENFIELD I-PER

Using CoNLL++, we conduct an empirical study of more than 20 NER models that were trained on the original CoNLL-2003 training split. Our analysis shows that different models can have very different generalization when moving to modern data. Simply comparing the performance of models on the CoNLL-2003 test set does not tell the whole story of progress on NER over the past 20 years.

Similar to the findings of Recht et al. (2019) on the ImageNet dataset (Deng et al., 2009), we do not observe evidence of widespread overfitting on CoNLL-2003. On average, each point of F_1 improvement on the CoNLL-2003 test set translates to a larger improvement on CoNLL++ (see Figure 2), suggesting overall improvements on the original dataset between 2003-2020 are mostly *not* due to overfitting. Rather, most performance deterioration appears to be caused by temporal misalignment (Luu et al., 2022).

Suprisingly, for some models (e.g. RoBERTa and T5), we find no evidence of performance degradation at all, despite the fact they are fine-tuned on a 20-year-old public dataset. We conduct an extensive analysis, which suggests that model size,

architecture, amount of fine-tuning data, and pre-training corpus are all important factors for generalization in NER.

2 Annotating a New Test Set to Measure Generalization

Data Collection: The CoNLL-2003 shared task collected English data from the Reuters Corpus, including Reuters news articles published between Aug. 1996 and Aug. 1997. The test set was collected from December 1996 according to Tjong Kim Sang and De Meulder (2003). We find that almost all articles were published between Dec. 5th and 7th, 1996, except one article published on Nov. 29th and another on Dec. 16th. Our dataset follows this distribution to collect Reuters news articles published between December 5th and 7th, 2020, collected from the Common Crawl Foundation.² We tokenize the data with the same tokenizer used for the CoNLL-2003 shared task, and randomly select articles to match the total number of tokens in the original test set.

Annotation: We manually labeled this new dataset, which we refer to as CoNLL++, using the BRAT annotation interface (Stenetorp et al., 2012). Articles were distributed between two authors, where one author annotated 96.1% of the articles and the other annotated 50.0%. The first author’s annotation is used as the gold standard.³ During the annotation process, articles from the CoNLL-2003 test set were interleaved with new articles from

¹We notice that our dataset contains fewer entities than CoNLL-2003. This is mainly because there are a number of tabular data, with information such as results of sports events (e.g. 1. Jesper Ronnback (Sweden) 25.76 points), in CoNLL-2003. Such data greatly contribute to the number of entities. These tabular data also cause the average sentence length of CoNLL-2003 to be smaller than that of CoNLL++. By removing these data, we found that the average sentence length increased to 18.50, much more comparable to CoNLL++. Model performances reported in Figure 2 were also not affected by the removal of these tabular data. We include further analysis and explanation in the Appendix (§ A).

²<http://commoncrawl.org/>

³Articles only annotated by the second author were reviewed and then used as the gold standard.

Name	Architecture	Reference	Corpus	Corpus Time	ΔF_1 (%)	Δ Rank
BiLSTM-CRF	GloVe+RNN+CRF	Lample et al. (2016) ⁵	WP	till 2014	-20.25	-7
BiLSTM-CNN	GloVe+RNN	Chiu and Nichols (2016) ⁶	WP	till 2014	-15.09	0
Stanford NLP	CRF	Finkel et al. (2005)	-	-	-13.25	0
SciBERT	BERT	Beltagy et al. (2019)	SS	till 2019*	-8.94	+2
BiLSTM-CNN-CRF	GloVe+RNN+CRF	Ma and Hovy (2016) ⁷	WP	till 2014	-6.52	-1
BiLSTM-CRF-ELMo	ELMo+RNN+CRF	Peters et al. (2018)	1B	till 2011*	-5.72	-7
Flair	GloVe+Flair+RNN+CRF	Akbik et al. (2018)	WP 1B	till 2014* till 2011*	-5.57	-7
Stanza	Flair+RNN	Qi et al. (2020)	1B	till 2011*	-5.19	-8
Pooled Flair	GloVe+Flair+RNN+CRF	Akbik et al. (2019)	WP 1B	till 2014* till 2011*	-4.65	-6
mBERT	BERT	Devlin et al. (2019)	WP	01/2001-2018*	-4.22	-1
GigaBERT	BERT	Lan et al. (2020a)	G5 WP OS	01/2009-12/2010 till 2019* 11/2018	-3.90	0
ALBERT _{Base}	BERT	Lan et al. (2020b)	BP	till 2018*	-3.61	+3
ALBERT _{XXL}	BERT	Lan et al. (2020b)	BP	till 2018*	-2.22	+4
BERT _{Large}	BERT	Devlin et al. (2019)	BC WP	till 2015* 01/2001-2018*	-2.01	+3
XLM-RoBERTa _{Base}	RoBERTa	Conneau et al. (2020)	CC	till 2020*	-0.90	+5
T5 _{Large}	Transformer	Raffel et al. (2020)	C4	04/2019	-0.59	+11
RoBERTa _{Large}	RoBERTa	Liu et al. (2019)	BP CN OW ST	till 2018* 09/2016-02/2019 till 2019* till 2018*	+0.64	+3
T5 _{3B}	Transformer	Raffel et al. (2020)	C4	04/2019	+0.67	+1
Longformer _{Base}	RoBERTa	Beltagy et al. (2020)	BP RP ST RN	till 2018* till 2019 till 2018* 12/2016-03/2019	+1.00	+5
news-RoBERTa _{Base}	RoBERTa	Gururangan et al. (2020)	RP RN	till 2019 12/2016-03/2019	+1.06	+5
LUKE _{Large}	RoBERTa+EASA [†]	Yamada et al. (2020)	RP WP	till 2019 till 12/2018	+1.10	0

Table 2: Details about the models selected, sorted by ΔF_1 . We list the models’ architectures and word embeddings, pre-training corpora, and the temporal coverage of the corpora. If the exact temporal coverage cannot be found, we report the time of publication of the corpus followed by *. For each model, we report the percentage change in F_1 and the change in ranking. Abbreviations: **BC** = BookCorpus (Zhu et al., 2015), **BP** = BERT Pre-training Corpus, **CC** = CC-100 (Conneau et al., 2020), **CN** = CC-News (Nagel, 2016), **C4** = Colossal Clean Crawled Corpus (Raffel et al., 2020), **G5** = Gigaword5 (Parker et al., 2011), **OS** = OSCAR (Suárez et al., 2019), **OW** = OpenWebText (Gokaslan et al., 2019), **RN** = REALNEWS (Zellers et al., 2019), **RP** = RoBERTa Pre-training Corpus, **SS** = Semantic Scholar (Cohan et al., 2019), **ST** = Stories (Trinh and Le, 2018), **WP** = Wikipedia, **1B** = 1B Benchmark (Chelba et al., 2014). [†]Entity-aware self-attention (Yamada et al., 2020).

2020, in order to measure how closely the annotators follow the style of the original dataset.

Inter-Rater Agreement: We find that the CoNLL++ annotations closely follow the style of the original dataset. When considering labels in the CoNLL-2003 test set as gold, our manual re-annotation achieves a 95.46 F_1 score.⁴ The second author’s annotation, when considering the first author’s as gold, receives a 96.23 F_1 score on overlapping articles. The token-level Cohen’s Kappa between the two authors is 97.42, which can be considered almost perfect agreement (Artstein and Poesio, 2008). Table 1 summarizes the statistics of

⁴For reference, the current state of the art for automated NER taggers is 94.60 (Wang et al., 2021).

the two datasets.

3 Experimental Setup

We select models with a variety of architectures and pre-training corpora and fine-tune these models to study how different factors affect generalization. None of the models used any pre-training data that temporally overlap with CoNLL++, eliminating the possibility of articles in CoNLL++ appearing in any pre-training corpus. A list of all models and their implementation details can be found in Table 2.

Scripts for fine-tuning Flair and ELMo are

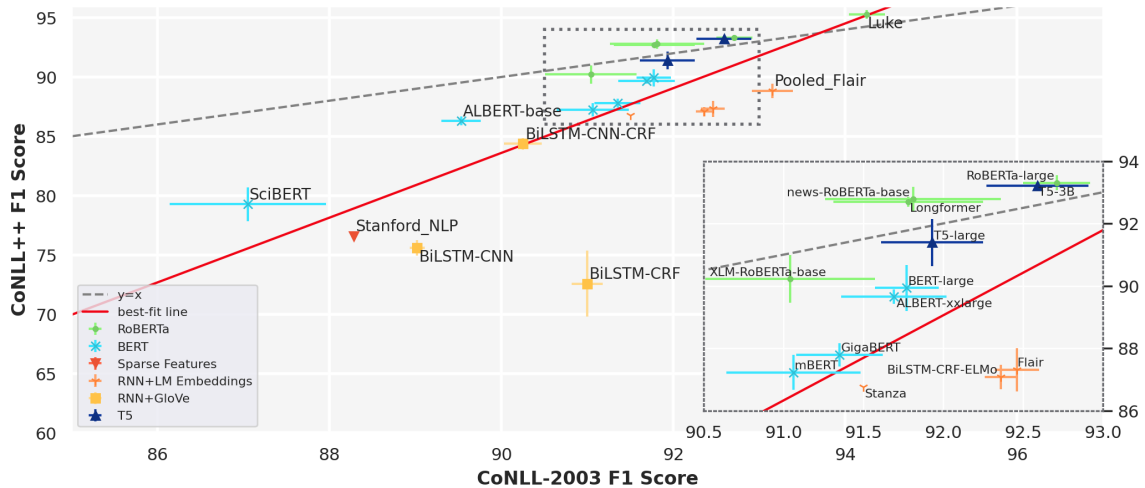


Figure 2: Plot of CoNLL++ F₁ scores against CoNLL-2003 F₁ scores. Each data point represents the average F₁ for each model, and the error bar represents one standard deviation. We observe that models show different level of generalization, while T5 and RoBERTa models generalize to CoNLL++. The solid best-fit line is steeper than the dashed $y = x$ ideal generalization line, providing evidence against adaptive overfitting (§ 5.1). This figure is best viewed in color.

adapted from Reiss et al. (2020).⁵ Other recurrent neural network (RNN) models are trained using various GitHub repositories (see footnotes 6, 7 and 8). We fine-tune the BERT and RoBERTa models with the HuggingFace transformers library (Wolf et al., 2020), except LUKE with AllenNLP (Gardner et al., 2018). T5 is fine-tuned to conditionally generate NER tags around entities (e.g. <per> Jane Doe </per>).

A hyperparameter search is conducted for each model. We follow the recommended search space for a model if available in its publication. {8, 16, 32} and {1e-5, 2e-5, 3e-5, 5e-5} are used for most searches for batch sizes and learning rates respectively. Appendix B provides more details on the hyperparameter search.

We train models on the CoNLL-2003 training set for 10 epochs, and use the dev set to select the best epoch and other hyperparameters for evaluation. Each model is evaluated five times with different random seeds on the CoNLL-2003 test set and on CoNLL++ to obtain the average F₁.

In Table 2, we report the percentage change of F₁, calculated as:

$$\Delta F_1 = \frac{F_1^{\text{CoNLL++}} - F_1^{\text{CoNLL-2003}}}{F_1^{\text{CoNLL-2003}}} \times 100$$

⁵Implementation from Reiss et al. (2021)

⁶Implementation from Jie (2020)

⁷Implementation from Kanakarajan (2019)

⁸Implementation from Reimers and Gurevych (2017)

where $F_1^{\text{CoNLL++}}$ and $F_1^{\text{CoNLL-2003}}$ are the F₁ scores on the CoNLL++ and CoNLL-2003 test sets respectively. The results are visualized in Figure 2. Raw F₁ scores are shown in Table 5 in the Appendix (§ C.1).

4 What Ingredients are Needed for Good Generalization?

As we can see in Figure 2 and Table 2, different models have very different generalization. Some models (e.g. RoBERTa-based models and T5_{3B}), have no performance drop on CoNLL++, whereas other models’ performances decrease significantly.

In the following sub-sections, we evaluate the impact of a number of factors on generalization. In §5, we attempt to disentangle to what extent the observed performance drops on CoNLL++ are caused by temporal deterioration, or adaptive overfitting.

4.1 Model Size

It has been shown that the size of pre-trained models affects their performance (Kaplan et al., 2020; Raffel et al., 2020). This inspired us to investigate the effect of model size on generalization. We compare the performance of BERT, RoBERTa, ALBERT and T5 models with different sizes on CoNLL++ and CoNLL-2003. The results are visualized in Figure 3. Details are available in Table 6 in the Appendix (§ C.2).

We observe, from Table 6, that larger models perform better on both test sets, but more impor-

tantly, as illustrated in Figure 3, performance degradation on CoNLL++ diminishes or even disappears as the model size grows. The only exception is the RoBERTa-based models, whose base-sized model already achieves comparable performance on CoNLL++. Figure 3 suggests that larger model sizes not only increase performance on a static test set, but also help models generalize better to new data.

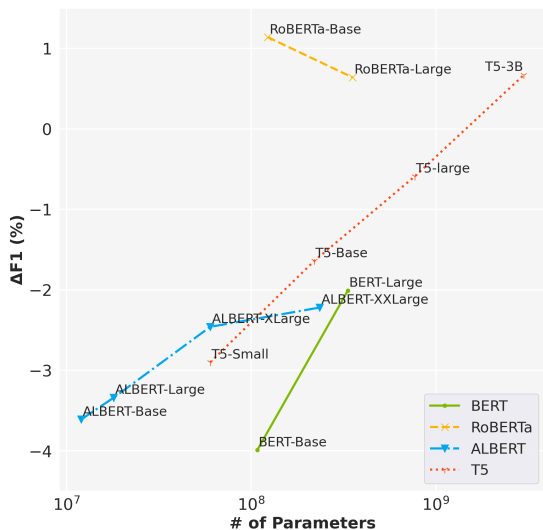


Figure 3: Plot of percentage change in F_1 scores (ΔF_1) against the number of parameters in log scale. All models except RoBERTa show an improvement in generalizability as the model size grows.

It is also informative to look at the individual trend within each model family. Whereas T5 models exhibit a linear relationship between the log number of parameters and ΔF_1 , the improvement of ΔF_1 for ALBERT models diminishes as the size grows larger. Additionally, models of similar sizes do not necessarily exhibit similar generalization. For example, $BERT_{Base}$ & $RoBERTa_{Base}$ ($\sim 100M$), $ALBERT_{xxLarge}$ & $T5_{Base}$ ($\sim 220M$) and $BERT_{Large}$ & $RoBERTa_{Large}$ ($\sim 300M$) all have similar sizes, but the performance changes within each pair are very different.

Both $RoBERTa_{Large}$ and $T5_{3B}$ achieve a performance increase of $\sim 0.6\%$, but the number of parameters of $T5_{3B}$ is ~ 10 times of that of $RoBERTa_{Large}$. This shows that the generalizability of model is also affected by factors other than the size of the model, but with the same architectures, larger models tend to generalize better.

4.2 Model Architecture

Based on the results from Table 2, we also observe that model architecture has a significant impact on generalizability. Most BERT, RoBERTa and T5 models have a small performance drop (less than 4% F_1) on CoNLL++. The performances of $RoBERTa_{Large}$, $news-RoBERTa_{Base}$, LUKE and $Longformer_{Base}$ improved slightly on CoNLL++. The fact that most Transformer-based models achieve higher rankings in CoNLL++ also confirms that pre-trained Transformers generalize better to new data.

BiLSTM models with Flair and ELMo embeddings, despite performing exceptionally on CoNLL-2003, show larger performance drops on CoNLL++ (5-6% F_1), and the performance of BiLSTM+GloVe models drops even more significantly (greater than 6% F_1). Such results show a clear trend that Transformer-based models generalize better to new data.

4.3 Number of Fine-Tuning Examples

The generalizability of a model may also be affected by the size of the fine-tuning dataset. We conduct experiments varying the number of CoNLL-2003 training examples used for fine-tuning from 10% to 100%. The fine-tuning is done with $RoBERTa_{Base}$ and Flair embeddings using the same experimental setup as in Section 3. We plot the percentage change in F_1 against the percentage of training examples in Fig 4.



Figure 4: Plot of change in F_1 scores (ΔF_1) against the percentage of CoNLL-2003 training data used for fine-tuning. Both $RoBERTa_{Base}$ and Flair show improved generalization as we use more training examples, although Flair shows a more pronounced improvement.

Both $RoBERTa_{Base}$ and Flair embeddings show improved generalization as we use more training

examples. However, this improvement is more pronounced for Flair than RoBERTa_{Base}. Even with 10% of the training data, RoBERTa_{Base} already shows a positive change in F₁ scores, and increasing the amount of training data to 100% only improves the change by an absolute value of 1%. In contrast, increasing the amount of training data from 10% to just 20% can already improve ΔF_1 by 2% for Flair.

The empirical evidence supports our hypothesis that having more training examples can improve the generalizability of the model, but such effect may vary across different models. RoBERTa-based models generalize well to new data even when only a small amount of fine-tuning data is available, whereas Flair benefits much more from having more fine-tuning data.

5 What Causes the Performance Drop Observed for Some Models?

Models in Table 2 show different levels of performance drop, or sometimes performance gain, on CoNLL++ compared to the CoNLL-2003 test set, and it is not entirely clear what causes this difference. We hypothesize two potential causes, namely adaptive overfitting (§ 5.1) and temporal drift (§ 5.2). In this section, we investigate each of these potential causes.

5.1 Adaptive Overfitting

We first investigate if the performance drop is caused by adaptive overfitting of models developed over the past 20 years. Roelofs et al. (2019) defined adaptive overfitting as the overfitting caused by reusing the same test set (*test reuse*). Recht et al. (2019) studied this phenomenon in the context of ImageNet by measuring to what extent can the improvement on the old test set translate to improvement on the new test set (*diminishing return*). We analyze both effects to conclude the presence of adaptive overfitting.

5.1.1 Diminishing Return

Following Recht et al. (2019), we measure the diminishing return on the CoNLL++ test set. Diminishing return measures if the improvement on CoNLL-2003 test set, gained by the continuous effort of developing NER taggers over 20 years, translates to smaller (*diminishing*) improvement on CoNLL++.

We fit a line to the data points in Figure 2, and then calculate its slope. A slope greater than 1

indicates that every unit of improvement on the CoNLL-2003 test set by the development of models translates to more than one unit of improvement on CoNLL++, i.e. there is no diminishing return. We measure the slope to be $2.729 > 1$, indicating that we have not found any diminishing return on CoNLL++, and therefore *no* adaptive overfitting caused by the model development over the past two decades.

5.1.2 Test Reuse

If the models are overfitting to the CoNLL-2003 test set due to test reuse, we should see not only a performance degradation on CoNLL++, but also a performance degradation on a test set taken from the same distribution as the CoNLL-2003 test set.

To obtain a new test set taken from the same distribution as the CoNLL-2003 test set, we resampled a new train/dev/test split from the CoNLL-2003 dataset, which we call CoNLL-2003'. Each split contains the same number of articles as its corresponding split in the CoNLL-2003 dataset. The "new" test set is thus certain to come from the same distribution as the original CoNLL-2003 test. We train and evaluate models on CoNLL-2003' with the same experimental setup as in Section 3, and report the results in Table 3.

Name	CoNLL++ ΔF_1 (%)	CoNLL-2003' ΔF_1 (%)
BiLSTM-CRF	-20.25	+2.53
BiLSTM-CNN	-15.09	+1.75
SciBERT	-8.94	-0.09
BiLSTM-CNN-CRF	-6.52	+2.95
BiLSTM-CRF-ELMo	-5.72	+1.58
Flair	-5.57	+0.76
Pooled Flair	-4.65	+1.60
mBERT	-4.22	+2.80
GigaBERT	-3.90	+1.75
ALBERT _{Base}	-3.61	+2.36
BERT _{Large}	-2.01	+0.47
XLm-RoBERTa _{Base}	-0.90	-1.52
T5 _{Large}	-0.59	+2.65
RoBERTa _{Large}	+0.64	+0.38
Longformer _{Base}	+1.00	+2.44
news-RoBERTa _{Base}	+1.06	+1.90
Luke	+1.10	+1.87

Table 3: Comparison between the performance change on CoNLL++ and CoNLL-2003' test sets. The table shows clearly that the performances of most models are not degrading because of test reuse. Detailed results can be found in Table 10 in Appendix C.5.

We only observe SciBERT and XLm-RoBERTa_{Large} models performing slightly worse on the CoNLL-2003' test set, while all other models appear to perform better. Most models suffering

from performance degradation on the CoNLL++ also perform better on the CoNLL-2003’ test set. This provides evidence that individual models are *not* overfitting to the CoNLL-2003 test set.

Based on our results above, the performance degradation on the CoNLL++ is likely *not* caused by overfitting on CoNLL-2003. Rather, it is more likely caused by temporal drift, which we discuss in the next section.

5.2 Temporal Drift

Temporal drift refers to the performance degradation of a model on the downstream task caused by the temporal difference between the train and test data. Prior work has shown that the performance on NER is affected by temporal drift. For example, [Rijhwani and Preotiuc-Pietro \(2020\)](#) showed that the performance of GloVe and Flair embeddings on NER degrades when the test data is more temporally distant from the train data of the downstream task. [Agarwal and Nenkova \(2022\)](#) also reported the same observation on GloVe embeddings.

In this section, we use the same term “temporal drift” but refer to the deterioration of *generalization* of models caused by the temporal difference between the pre-training corpus of their word embeddings and the test data of the downstream task. We hypothesize that generalization is largely affected by such temporal drift. We conduct experiments on Flair and ELMo, as well as on RoBERTa.

5.2.1 Temporal Drift in Flair and ELMo

We first investigate if bringing the pre-training corpora of Flair and ELMo closer to the test set can improve their generalizability. We notice that both embeddings were trained on 1B Benchmark ([Chelba et al., 2014](#)). This corpus was collected from WMT11 ([Callison-Burch et al., 2011](#)) English monolingual data, which is largely comprised of news data between 2007-2011. We hypothesize that pre-training these embeddings on a more recent corpus, e.g. REALNEWS corpus ([Zellers et al., 2019](#)) which contains news articles from 2016-2019, will improve their generalizability.

To control the experiment, we randomly sample 1 billion tokens of data from REALNEWS. We train Flair embeddings following the same procedure detailed in [Akbik et al. \(2018\)](#) to obtain both the forward and backward embeddings. Our embeddings achieve character level perplexity on the test set of 2.45 for the forward embeddings and 2.46 for the backward embeddings, comparable to 2.42

Name	CoNLL-2003	CoNLL++	ΔF_1 (%)
Flair	92.46 _{0.14}	87.31 _{0.69}	-5.57
Flair _{RN}	90.91 _{0.22}	88.46 _{0.69}	-2.69
Pooled Flair	93.15 _{0.24}	88.82 _{0.60}	-4.65
Pooled Flair _{RN}	92.98 _{0.14}	89.73 _{0.27}	-3.50
ELMo	92.36 _{0.10}	87.08 _{0.39}	-5.72
ELMo _{RN}	92.11 _{0.07}	90.79 _{0.50}	-1.43

Table 4: Percentage change in F_1 scores (ΔF_1) on CoNLL++ of Flair and ELMo embeddings when pre-trained on 1B Benchmark vs on REALNEWS corpus. Pre-training on REALNEWS, which is temporally closer to CoNLL++, improves the generalization of Flair and ELMo embeddings.

reported in [Akbik et al. \(2018\)](#). Similarly, we train ELMo embeddings following [Peters et al. \(2018\)](#), which achieves a perplexity of 40.07 on the test set, comparable to 39.7 reported. We use the same training scripts and hyperparameters as our experiments in Section 4.2 for Flair, Pooled Flair and ELMo. The newly trained models are dubbed as Flair_{RN}, Pooled Flair_{RN} and ELMo_{RN}.

It is clearly shown in Table 4 that having the training corpus for Flair and ELMo embeddings temporally closer to the CoNLL++ test set improves generalization. Notably, the performance gap for ELMo is reduced to -1.43%, better than that of BERT_{Large} (-2.01%). The improvements in generalization are attributed to the performance drops on the CoNLL-2003 test set and improvements on CoNLL++.

This provides evidence that the generalizability of the LSTM-based contextualized word embeddings is affected by temporal drift. However, even temporally closer data, these models still suffer from performance drops. This suggests that other ingredients, such as model architecture (§ 4.2), are still needed for a good generalization.

5.2.2 Temporal Drift in RoBERTa

Because pre-training a transformer model from scratch is expensive, we continue pre-training from the RoBERTa_{Base} checkpoint, leveraging the findings from [Gururangan et al. \(2020\)](#) that models learn to adapt to the distribution of the new corpora with continued pre-training.

We use the WMT20 English dataset ([Barrault et al., 2019](#)), consisting of English news data from 2007 to 2021. To avoid temporal overlap, we only use data from 2007 to 2019 as the pre-training corpora. The data are divided by year, and we pre-process the data such that the number of tokens per

year is the same. We train the RoBERTa_{Base} model for 3 epochs with the masked language modeling (MLM) objective. Checkpoints from each year are then fine-tuned on the CoNLL-2003 dataset, with the same experimental setup described in Section 3. We evaluate the models on the CoNLL-2003 test set and CoNLL++, and plot the results in Figure 5. Detailed performances are reported in Table 9 in the Appendix (§ C.4).

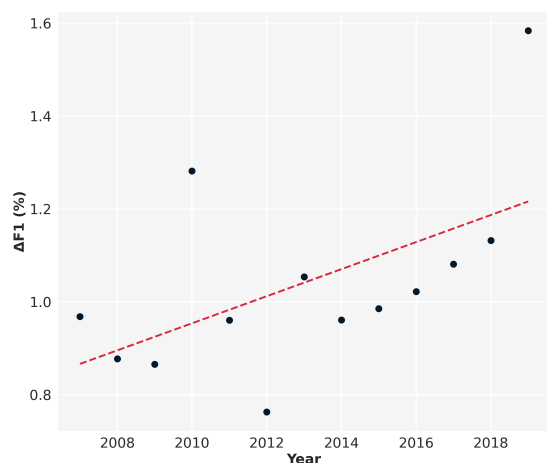


Figure 5: Plot of ΔF_1 scores against the year of data used for RoBERTa pre-training. The upward trend, indicated by the dashed best-fit line, shows that the generalization improves as the pre-training corpora used is temporally closer to CoNLL++.

The results show a clear trend of performance degradation when the pre-training corpora is temporally distant from the test set. When the pre-training corpus is more recent, it becomes temporally closer to CoNLL++, leading to better CoNLL++ performance, and hence better generalization. In Figure 5, the ΔF_1 shows an upward trend with a correlation coefficients of 0.55, indicating a moderate positive correlation between generalization and the year of the pre-training corpora. This suggests that generalization is affected by the effect of temporal drift. This explains the better generalizability of models such as LUKE_{Large} and T5_{3B}, pre-trained on temporally closer data to the CoNLL++ test set, showing that temporal drift is the main driving factor for the different levels of generalization.

6 Related Work

How well pre-trained LMs adapt to data from future time periods has undergone extensive study. Temporal degradation has been found to be a challenge for many tasks, including language modeling (Lazaridou et al., 2021), NER (Augenstein

et al., 2017; Agarwal and Nenkova, 2022; Rijhwani and Preotiuc-Pietro, 2020; Ushio et al., 2022), QA (Dhingra et al., 2022), entity linking (Zaporojets et al., 2022), and others (Luu et al., 2022; Amba Hombaiah et al., 2021). All of this work has found that the performance of LMs degrades as the temporal distance between the training data and the test data increases, sometimes called “temporal misalignment” (Luu et al., 2022). In contrast to the prior work, we study performance deterioration on a dataset that has been heavily used to develop NER models over a period of 20 years, and conduct extensive experiments that aim to disentangle the effects of aging training sets from those due to heavy test reuse.

Most closely related to our work is Agarwal and Nenkova (2022), who analyzed a recently created Twitter NER dataset (Rijhwani and Preotiuc-Pietro, 2020) over the period 2014-2019, and found no performance deterioration when using RoBERTa-based representations. We build on this line of work by carefully measuring performance deterioration of models trained on the CoNLL-2003 dataset when evaluated on modern data. We analyze which factors are necessary for an NER model trained on a 20-year-old dataset to generalize well to modern data. Furthermore, the large 20-year gap helps us focus on not only temporal deterioration, but also if the extensive test reuse leads to adaptive overfitting. We present evidence in support of the hypothesis that most performance degradation is due to temporal drift and not adaptive overfitting.

Prior work has attempted to mitigate temporal degradation, mostly through continuously updating LMs with new data (Jang et al., 2022; Jin et al., 2022; Loureiro et al., 2022). Luu et al. (2022) explored this idea but found that temporal adaptation is not as effective as fine-tuning on the data from whose time period the dataset is drawn. In addition, catastrophic forgetting (Robins, 1995) can also be a problem when updating the LMs. Jin et al. (2022) found that applying knowledge distillation (Hinton et al., 2015) based approaches to continual learning can mitigate catastrophic forgetting, while improving the temporal generalization of LMs. Dhingra et al. (2022) proposed to train the LMs with an additional temporal objective by conditioning on the year of data, and found that this effectively mitigated catastrophic forgetting. Jang et al. (2022) created a lifelong benchmark for continuous training and evaluating LMs.

7 Conclusion and Future Directions

In this paper, we evaluate the generalization of NER models using CoNLL++, a CoNLL-style annotated NER test dataset with data from 2020. We conduct experiments on more than 20 models and find that models exhibit different generalizability.

Surprisingly, we find that generalizability is *not* affected by adaptive overfitting, but rather by temporal drift. To achieve better generalization, we need the combination of four factors: a modern transformer-based architecture, a large number of parameters, a large amount of fine-tuning data and a temporally closer pre-training corpus to the test set. We find that our progress on developing NER taggers is largely successful, showing not only good performance on individual test set but also good generalization on new data. This allows CoNLL-2003 taggers to still work in 2023.

Future research can focus on ways to mitigate temporal drift. Investigation on attributes of pre-training or fine-tuning corpora that causes temporal drift, such as change of entities mentioned, different usage of language, etc., can also shed light on the more specific impacts from temporal drift, thereby inspiring new and better ways to mitigate it.

We hope that our work provides insights on factors affecting generalization and how to mitigate the negative impact, and calls for more research on this everlasting problem of generalization in the NLP community.

8 Limitations

Our analysis on temporal drift (§ 5.2) was limited by the fact that the developer of many models in our study did not release the exact time period of the pre-training corpora used. Additionally, models such as BERT and RoBERTa were pre-trained on corpora that could be potentially be temporally close to the CoNLL++ test set.

In the section on test reuse (§ 5.1.2), due to a limited compute budget, we were only able to conduct this experiment on a single new train/dev/test split, so it is possible that the new split happens to be easier than the mean of the distribution. However, our experiments still provide additional evidence models are not overfitting the original CoNLL-2003 test set.

It is worth noting that when using older models trained on the CoNLL-2003 dataset, one additional reason for the performance degradation, especially

in real-world deployment, is that the data used to evaluate the models can be out-of-domain. In our experiments, we attempted to control the domain of the test data on which the models were evaluated to assess other factors for performance degradation. However, we acknowledge that in reality, model performance can be affected by factors such as the emerging text types (e.g. Twitter did not exist when CoNLL-2003 NER task was created), which leads to changes in domain, and therefore affects the generalizability of the models.

We acknowledge that having CoNLL++ will not resolve the problem of generalization to modern data. As new data keep emerging, there will always be the question of how well NER models generalize to that new data. We hope that our paper will encourage researchers in the NLP community to continuously annotate new test set to study this problem, so that we ensure the robustness and generalizability of models.

Acknowledgements

We would like to thank Fan Bai, Yang Chen, Chao Jiang, Junmo Kang and Mounica Maddela for providing feedback on earlier drafts of this paper, We would also like to appreciate the comments from the anonymous reviewers on how to improve the paper. This material is based upon work supported by the NSF (IIS-2052498) and IARPA via the BETTER program (2019-19051600004). The views, opinions, and/or findings expressed are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense, IARPA or the U.S. Government.

References

- Oshin Agarwal and Ani Nenkova. 2022. [Temporal effects on pre-trained models for language processing tasks](#). *Transactions of the Association for Computational Linguistics*, 10:904–921.
- Alan Akbik, Tanja Bergmann, and Roland Vollgraf. 2019. [Pooled contextualized embeddings for named entity recognition](#). 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 724–728, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. [Contextual string embeddings for sequence labeling](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages

- 1638–1649, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Spurthi Amba Hombaiyah, Tao Chen, Mingyang Zhang, Michael Bendersky, and Marc Najork. 2021. Dynamic language models for continuously evolving content. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 2514–2524.
- Ron Artstein and Massimo Poesio. 2008. [Survey article: Inter-coder agreement for computational linguistics](#). *Computational Linguistics*, 34(4):555–596.
- Isabelle Augenstein, Leon Derczynski, and Kalina Bontcheva. 2017. Generalisation in named entity recognition: A quantitative analysis. *Computer Speech & Language*, 44:61–83.
- Loïc Barrault, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. [Findings of the 2019 conference on machine translation \(WMT19\)](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, Florence, Italy. Association for Computational Linguistics.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The Long-Document Transformer. *ArXiv*.
- Chris Callison-Burch, Philipp Koehn, Christof Monz, and Omar F. Zaidan, editors. 2011. *Proceedings of the Sixth Workshop on Statistical Machine Translation*. Association for Computational Linguistics.
- Ciprian Chelba, Tomas Mikolov, Mike Schuster, Qi Ge, T. Brants, Phillip Todd Koehn, and Tony Robinson. 2014. One Billion Word Benchmark for Measuring Progress in Statistical Language Modeling. *ArXiv*.
- Jason P.C. Chiu and Eric Nichols. 2016. [Named entity recognition with bidirectional LSTM-CNNs](#). *Transactions of the Association for Computational Linguistics*, 4:357–370.
- Arman Cohan, Waleed Ammar, Madeleine van Zuylen, and Field Cady. 2019. [Structural scaffolds for citation intent classification in scientific publications](#). 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 3586–3596, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Nicola De Cao, Ledell Wu, Kashyap Papat, Mikel Artetxe, Naman Goyal, Mikhail Plekhanov, Luke Zettlemoyer, Nicola Cancedda, Sebastian Riedel, and Fabio Petroni. 2022. [Multilingual autoregressive entity linking](#). *Transactions of the Association for Computational Linguistics*, 10:274–290.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. [Imagenet: A large-scale hierarchical image database](#). In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255.
- 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.
- Bhuvan Dhingra, Jeremy R. Cole, Julian Martin Eisenschlos, Daniel Gillick, Jacob Eisenstein, and William W. Cohen. 2022. [Time-aware language models as temporal knowledge bases](#). *Transactions of the Association for Computational Linguistics*, 10:257–273.
- Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. [Incorporating non-local information into information extraction systems by Gibbs sampling](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pages 363–370, Ann Arbor, Michigan. Association for Computational Linguistics.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. [AllenNLP: A deep semantic natural language processing platform](#). In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia. Association for Computational Linguistics.
- Aaron Gokaslan, Vanya Cohen, Ellie Pavlick, and Stefanie Tellex. 2019. OpenWebText Corpus. <http://Skylion007.github.io/OpenWebTextCorpus>.

- Kyle Gorman and Steven Bedrick. 2019. [We need to talk about standard splits](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2786–2791, Florence, Italy. Association for Computational Linguistics.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. [Distilling the knowledge in a neural network](#).
- Joel Jang, Seonghyeon Ye, Changho Lee, Sohee Yang, Joongbo Shin, Janghoon Han, Gyeonghun Kim, and Minjoon Seo. 2022. [TemporalWiki: A lifelong benchmark for training and evaluating ever-evolving language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6237–6250, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Allan Jie. 2020. [Allanj/pytorch_neural_crf: Pytorch implementation of LSTM/Bert-CRF for named entity recognition](#).
- Xisen Jin, Dejiao Zhang, Henghui Zhu, Wei Xiao, Shang-Wen Li, Xiaokai Wei, Andrew Arnold, and Xiang Ren. 2022. [Lifelong pretraining: Continually adapting language models to emerging corpora](#). In *Proceedings of BigScience Episode #5 – Workshop on Challenges & Perspectives in Creating Large Language Models*, pages 1–16, virtual+Dublin. Association for Computational Linguistics.
- Kamal Raj Kanakarajan. 2019. [Kamalkraj/named-entity-recognition-with-bidirectional-LSTM-CNNs: Named-entity-recognition-with-bidirectional-LSTM-CNNs](#).
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling Laws for Neural Language Models](#). *CoRR*, abs/2001.08361.
- 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.
- Wuwei Lan, Yang Chen, Wei Xu, and Alan Ritter. 2020a. [An empirical study of pre-trained transformers for Arabic information extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4727–4734, Online. Association for Computational Linguistics.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020b. [Albert: A lite bert for self-supervised learning of language representations](#). In *International Conference on Learning Representations*.
- Angeliki Lazaridou, Adhi Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Tomas Kocisky, Sebastian Ruder, et al. 2021. [Mind the gap: Assessing temporal generalization in neural language models](#). *Advances in Neural Information Processing Systems*, 34:29348–29363.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [RoBERTa: A Robustly Optimized BERT Pretraining Approach](#). *CoRR*.
- Daniel Loureiro, Francesco Barbieri, Leonardo Neves, Luis Espinosa Anke, and Jose Camacho-collados. 2022. [TimeLMs: Diachronic language models from Twitter](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 251–260, Dublin, Ireland. Association for Computational Linguistics.
- Kelvin Luu, Daniel Khashabi, Suchin Gururangan, Karishma Mandyam, and Noah A. Smith. 2022. [Time waits for no one! analysis and challenges of temporal misalignment](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5944–5958, Seattle, United States. Association for Computational Linguistics.
- Xuezhe Ma and Eduard Hovy. 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)*, pages 1064–1074, Berlin, Germany. Association for Computational Linguistics.
- John Morris, Justin Chiu, Ramin Zabih, and Alexander Rush. 2022. [Unsupervised text deidentification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4777–4788, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Sebastian Nagel. 2016. [CC-News](#).
- Robert Parker, David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2011. [English Gigaword Fifth Edition LDC2011T07](#).
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke

- Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. 2019. Do ImageNet Classifiers Generalize to ImageNet? In *Proceedings of International Conference on Machine Learning*.
- Nils Reimers and Iryna Gurevych. 2017. [Reporting score distributions makes a difference: Performance study of LSTM-networks for sequence tagging](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 338–348, Copenhagen, Denmark. Association for Computational Linguistics.
- Fred Reiss, Karthik Muthuraman, Zachary Eichenberger, Brian Cutler, and Hong Xu. 2021. [CODAIT/identifying-incorrect-labels-in-CoNLL-2003: Research into identifying and correcting incorrect labels in the CoNLL-2003 corpus](#).
- Frederick Reiss, Hong Xu, Bryan Cutler, Karthik Muthuraman, and Zachary Eichenberger. 2020. [Identifying incorrect labels in the CoNLL-2003 corpus](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 215–226, Online. Association for Computational Linguistics.
- Shruti Rijhwani and Daniel Preotiuc-Pietro. 2020. [Temporally-informed analysis of named entity recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7605–7617, Online. Association for Computational Linguistics.
- Anthony V. Robins. 1995. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connect. Sci.*, 7:123–146.
- Rebecca Roelofs, Sara Fridovich-Keil, John Miller, Vaishaal Shankar, Moritz Hardt, Benjamin Recht, and Ludwig Schmidt. 2019. *A Meta-Analysis of Overfitting in Machine Learning*. Curran Associates Inc., Red Hook, NY, USA.
- Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun’ichi Tsujii. 2012. [brat: a web-based tool for NLP-assisted text annotation](#). In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 102–107, Avignon, France. Association for Computational Linguistics.
- Pedro Javier Ortiz Suárez, Benoît Sagot, and Laurent Romary. 2019. [Asynchronous pipelines for processing huge corpora on medium to low resource infrastructures](#). Proceedings of the Workshop on Challenges in the Management of Large Corpora (CMLC-7) 2019, Cardiff, 22nd July 2019, pages 9 – 16, Mannheim. Leibniz-Institut für Deutsche Sprache.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Trieu H. Trinh and Quoc V. Le. 2018. [A Simple Method for Commonsense Reasoning](#).
- Asahi Ushio, Francesco Barbieri, Vitor Sousa, Leonardo Neves, and Jose Camacho-Collados. 2022. [Named entity recognition in Twitter: A dataset and analysis on short-term temporal shifts](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 309–319, Online only. Association for Computational Linguistics.
- Xinyu Wang, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, and Kewei Tu. 2021. [Automated concatenation of embeddings for structured prediction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2643–2660, Online. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. [LUKE: Deep contextualized entity representations with entity-aware self-attention](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6442–6454, Online. Association for Computational Linguistics.

Klim Zaporozhets, Lucie-Aimée Kaffee, Johannes Deleu, Thomas Demeester, Chris Develder, and Isabelle Augenstein. 2022. [TempEL: Linking dynamically evolving and newly emerging entities](#). In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. *Defending Against Neural Fake News*. Curran Associates Inc., Red Hook, NY, USA.

Zexuan Zhong and Danqi Chen. 2021. [A frustratingly easy approach for entity and relation extraction](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 50–61, Online. Association for Computational Linguistics.

Y. Zhu, R. Kiros, R. Zemel, R. Salakhutdinov, R. Urtasun, A. Torralba, and S. Fidler. 2015. [Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books](#). In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 19–27, Los Alamitos, CA, USA. IEEE Computer Society.

Appendix

A Tabular Data in CoNLL-2003

We found a significant amount of documents in the CoNLL-2003 test set that list the outcomes of various sports events, which contributes to the larger proportion of named entities in Table 1. These documents appear as though they may have been intended for display on news tickers.⁹ We present an example below.

SOCCKER SHOWCASE-BETTING ON
REAL MADRID V BARCELONA .
MADRID 1996-12-06
William Hill betting on Saturday 's
Spanish first division match between Real
Madrid and Barcelona :
To win : 6-5 Real Madrid ; 7-4 Barcelona
Draw : 9-4
Correct score :
Real Madrid to win Barcelona to win
1-0 13-2 1-0 15-2
2-0 9-1 2-0 12-1
2-1 8-1 2-1 10-1
3-0 20-1 3-0 28-1
3-1 16-1 3-1 22-1
3-2 25-1 3-2 25-1
4-0 50-1 4-0 100-1
4-1 40-1 4-1 80-1
4-2 50-1 4-2 80-1
Draw :
0-0 8-1
1-1 11-2
2-2 14-1
3-3 50-1
Double result :
half-time full-time
5-2 Real Madrid Real Madrid
14-1 Real Madrid Draw
28-1 Real Madrid Barcelona
5-1 Draw Real Madrid
4-1 Draw Draw
11-2 Draw Barcelona
25-1 Barcelona Real Madrid
14-1 Barcelona Draw
4-1 Barcelona Barcelona
First goalscorer of match :
Real Madrid Barcelona
9-2 Davor Suker 9-2 Ronaldo

⁹https://en.wikipedia.org/wiki/News_ticker

5-1 Pedrag Mijatovic 7-1 Luis Figo
 7-1 Raul Gonzalez 7-1 Juan Pizzi
 12-1 Fernando Redondo 9-1 Giovanni
 14-1 Victor Sanchez 12-1 Guillermo
 Amor
 16-1 Jose Amavisca 14-1 Roger Garcia
 16-1 Manolo Sanchis 14-1 Gheorghe
 Popescu
 16-1 Roberto Carlos 16-1
 JosepGuardiola
 20-1 Fernando Hierro 20-1 Ivan de
 laPena
 20-1 Luis Milla 25-1 Luis
 Enrique
 33-1 Fernando Sanz 25-1
 AbelardoFernandez
 40-1 Carlos Secretario 28-1 Sergi Barjuan
 40-1 Rafael Alkorta 33-1 Albert
 Ferrer
 40-1 Chendo Porlan 33-1 Miguel Nadal
 40-1
 Laurent Blanc

Being one of the 231 articles (0.43%) in the CoNLL-2003 test set, this article contains 59 (1.04%) named entities, including 23 ORG (1.38%), 34 PER (2.10%), 1 LOC and 1 MISC. We counted that there are in total 71 (30.7%) such files which contribute to 872 ORG (52.5%), 889 PER (55.0%), 657 LOC (39.4%) and 159 MISC (22.6%).

Additionally, as each line is considered to be a sentence in CoNLL-2003 dataset (separated by an empty line in the original format), and as items by spaces are considered to be tokens, this also demonstrates why the average token per sentence is much lower in CoNLL-2003 than in CoNLL++. The tabular data contains much shorter sentences in plethora, which significantly lowers the average token per sentence.

B Hyperparameter Search

In this section, we include the details on how we conducted the hyperparameter search for the transformer-based models. We trained most models with different sets of hyperparameters for 10 epochs and save the checkpoints that achieved the highest dev F_1 score. For each model, we compare performance on the dev set of checkpoints trained with different hyperparameters and select the set of hyperparameters with the best performance.

We tuned the learning rate and batch sizes for all models. If the instructions on how to tune the hyperparameters for a model are stated in its publication, we followed the instructions as closely as possible. Otherwise, we would tune the model using a default set of hyperparameters, where $\text{learning_rate} = \{1e-5, 2e-5, 3e-5, 5e-5\}$ and $\text{batch_size} = \{8, 16, 32\}$. Here we only list models for which we did not use the default set of hyperparameters.

- ALBERT:
 - $\text{learning_rate} = \{1e-5, 2e-5, 3e-5, 5e-5\}$
 - $\text{batch_size} = \{16, 32, 48, 128\}$
- GigaBERT:
 - $\text{learning_rate} = \{1e-5, 2e-5, 5e-5, 1e-4\}$
 - $\text{batch_size} = \{4, 8, 16, 32\}$
- Longformer:
 - $\text{learning_rate} = \{1e-5, 2e-5, 3e-5, 5e-5\}$
 - $\text{batch_size} = \{16, 32\}$
 - $\text{total_num_epoch} = 15$
- news_roberta_base:
 - $\text{learning_rate} = \{1e-5, 2e-5, 3e-5\}$
 - $\text{batch_size} = \{16, 32\}$
- XLM-RoBERTa:
 - $\text{learning_rate} = \{1e-5, 2e-5, 3e-5, 5e-5\}$
 - $\text{batch_size} = \{16, 32\}$
- T5:
 - $\text{learning_rate} = \{2e-5, 3e-5, 5e-5, 1e-4\}$
 - $\text{batch_size} = \{4, 8\}$

C Detailed Results

In this section, we include all the performance statistics.

C.1 CoNLL-2003 vs CoNLL++

Table 5 shows the performance statistics of all models on the CoNLL++ and CoNLL-2003 test sets.

C.2 Model Size

Table 6 includes the results from Section 4.1, showing the performance statistics of BERT-based, ALBERT-based, RoBERTa-based, and T5-based models with various sizes on the CoNLL-2003 and CoNLL++ test set. One side note is that our results also confirms the previous findings that the performance on a downstream task has a positive correlation with the model size.

Name	CoNLL-2003	CoNLL++	ΔF_1 (%)
BiLSTM-CRF	91.00 _{0.18}	72.57 _{2.78}	-20.25
BiLSTM-CNN	89.02 _{0.09}	75.59 _{0.66}	-15.09
Stanford NLP	88.28	76.58	-13.25
SciBERT	87.05 _{0.91}	79.27 _{1.43}	-8.94
BiLSTM-CNN-CRF	90.25 _{0.22}	84.37 _{0.49}	-6.52
BiLSTM-CRF-ELMo	92.36 _{0.10}	87.08 _{0.39}	-5.72
Flair	92.46 _{0.14}	87.31 _{0.69}	-5.57
Stanza	91.50	86.75	-5.19
Pooled Flair	93.15 _{0.24}	88.82 _{0.60}	-4.65
mBERT	91.06 _{0.42}	87.22 _{0.56}	-4.22
GigaBERT	91.35 _{0.27}	87.79 _{0.37}	-3.90
ALBERT _{Base}	89.53 _{0.23}	86.30 _{0.39}	-3.61
BERT _{Large}	91.77 _{0.20}	89.93 _{0.74}	-2.01
XLM-RoBERTa _{Base}	91.04 _{0.53}	90.22 _{0.77}	-0.90
T5 _{Large}	91.93 _{0.32}	91.39 _{0.75}	-0.59
RoBERTa _{Large}	92.71 _{0.21}	93.30 _{0.24}	+0.64
Longformer _{Base}	91.78 _{0.47}	92.70 _{0.16}	+1.00
news-RoBERTa _{Base}	91.81 _{0.55}	92.78 _{0.40}	+1.06
Luke	94.25 _{0.21}	95.29 _{0.37}	+1.10

Table 5: Detailed performances of the models on the CoNLL-2003 test set and the CoNLL++ test set, ranked by the ΔF_1 . The performances are F_1 scores calculated by taking the average over five runs and the standard deviations are presented in subscripts. The best results are highlighted in bold.

Name	# Parameters	CoNLL-2003	CoNLL++	ΔF_1 (%)
BERT _{Base}	108M	91.38 _{0.33}	87.73 _{0.51}	-3.99
BERT _{Large}	334M	91.77 _{0.20}	89.93 _{0.74}	-2.01
RoBERTa _{Base}	123M	92.08 _{0.22}	93.13 _{0.31}	+1.14
RoBERTa _{Large}	354M	92.71 _{0.21}	93.30 _{0.24}	+0.64
ALBERT _{Base}	12M	89.53 _{0.23}	86.30 _{0.39}	-3.61
ALBERT _{Large}	18M	90.46 _{0.21}	87.44 _{0.47}	-3.34
ALBERT _{XL}	60M	90.80 _{0.17}	88.57 _{1.03}	-2.46
ALBERT _{XXL}	235M	91.69 _{0.33}	89.65 _{0.23}	-2.22
T5 _{Small}	60M	88.94 _{0.32}	86.36 _{0.08}	-2.90
T5 _{Base}	220M	91.55 _{0.27}	90.05 _{0.45}	-1.64
T5 _{Large}	770M	91.93 _{0.32}	91.39 _{0.75}	-0.59
T5 _{3B}	3B	92.59 _{0.32}	93.21 _{0.09}	+0.67

Table 6: Performances of the models of different sizes on the CoNLL-2003 test set and the CoNLL++ test set. The performances are F_1 scores calculated by taking the average over five runs and the standard deviations are presented in subscripts.

C.3 Number of Fine-Tuning Examples

Table 7 and Table 8 show the results from Section 4.3 of the RoBERTa-based and Flair-based models on the two test sets respectively when varying the number of examples fine-tuned on.

C.4 Temporal Drift

Table 9 show the results from Section 5.2. The "Year" column indicates the time period from which the data used for continued pre-training on RoBERTa_{Base} was used.

C.5 Test Reuse

Table 10 shows the results from Section 5.1.2.

Training Example	CoNLL-2003	CoNLL++	ΔF_1 (%)
10%	88.28 _{0.38}	88.49 _{0.67}	+0.24
20%	90.23 _{0.30}	91.08 _{0.47}	+0.94
30%	90.81 _{0.21}	91.36 _{0.40}	+0.61
40%	91.10 _{0.12}	91.64 _{0.48}	+0.59
50%	91.42 _{0.15}	91.76 _{0.49}	+0.37
60%	91.45 _{0.27}	91.93 _{0.34}	+0.52
70%	91.82 _{0.10}	92.25 _{0.34}	+0.47
80%	91.98 _{0.15}	92.97 _{0.46}	+1.07
90%	92.04 _{0.20}	92.94 _{0.50}	+0.98

Table 7: Performances of RoBERTa_{Base} on the CoNLL-2003 test set and the CoNLL++ test set when varying the percentage of training examples used. The performances are F_1 scores calculated by taking the average over five runs and the standard deviations are presented in subscripts.

Training Example	CoNLL-2003	CoNLL++	ΔF_1 (%)
10%	86.90 _{0.15}	79.11 _{0.53}	-8.96
20%	88.42 _{0.45}	82.26 _{0.67}	-6.96
30%	89.04 _{0.24}	83.17 _{0.71}	-6.59
40%	89.74 _{0.11}	83.98 _{0.49}	-6.43
50%	90.15 _{0.13}	84.47 _{0.26}	-6.30
60%	90.40 _{0.28}	84.64 _{0.74}	-6.38
70%	90.62 _{0.16}	85.08 _{0.83}	-6.11
80%	90.68 _{0.16}	85.39 _{0.62}	-5.83
90%	90.84 _{0.17}	85.44 _{0.46}	-5.94

Table 8: Performances of Flair on the CoNLL-2003 test set and the CoNLL++ test set when varying the percentage of training examples used.

Year	CoNLL-2003	CoNLL++	ΔF_1 (%)
2007	91.96 _{0.44}	92.85 _{0.31}	+0.97
2008	91.88 _{0.09}	92.69 _{0.17}	+0.88
2009	92.24 _{0.17}	93.10 _{0.11}	+0.87
2010	91.92 _{0.25}	93.10 _{0.41}	+1.28
2011	92.07 _{0.35}	92.95 _{0.15}	+0.96
2012	92.07 _{0.34}	92.77 _{0.33}	+0.76
2013	91.87 _{0.23}	92.84 _{0.27}	+1.05
2014	92.01 _{0.32}	92.89 _{0.21}	+0.96
2015	91.95 _{0.29}	92.92 _{0.63}	+0.99
2016	91.98 _{0.23}	92.92 _{0.26}	+1.02
2017	91.93 _{0.13}	92.93 _{0.18}	+1.08
2018	91.89 _{0.38}	92.93 _{0.44}	+1.13
2019	91.80 _{0.29}	93.25 _{0.44}	+1.58

Table 9: Performances of different checkpoints obtained by continued pre-training RoBERTa_{Base} with data from different years on the CoNLL-2003 test set and CoNLL++.

Name	CoNLL-2003	CoNLL-2003'	ΔF_1 (%)
BiLSTM-CRF	91.00 _{0.18}	93.30 _{0.19}	+2.53
BiLSTM-CNN	89.02 _{0.09}	90.58 _{0.57}	+1.75
SciBERT	87.05 _{0.91}	86.97 _{0.84}	-0.09
BiLSTM-CNN-CRF	90.25 _{0.22}	92.61 _{0.29}	+2.95
BiLSTM-CRF-ELMo	92.36 _{0.10}	93.82 _{0.07}	+1.58
Flair	92.46 _{0.14}	93.16 _{0.13}	+0.76
Pooled Flair	93.15 _{0.24}	94.64 _{0.09}	+1.60
mBERT	91.06 _{0.42}	93.61 _{0.33}	+2.80
GigaBERT	91.35 _{0.27}	92.95 _{0.84}	+1.75
ALBERT _{Base}	89.53 _{0.23}	91.64 _{0.17}	+2.36
BERT _{Large}	91.77 _{0.20}	92.20 _{0.85}	+0.47
XLM-RoBERTa _{Base}	91.04 _{0.53}	89.66 _{0.42}	-1.52
T5 _{Large}	91.93 _{0.32}	94.37 _{0.32}	+2.65
RoBERTa _{Large}	92.71 _{0.21}	93.06 _{0.63}	+0.38
Longformer _{Base}	91.78 _{0.47}	94.02 _{0.48}	+2.44
news-RoBERTa _{Base}	91.81 _{0.55}	93.55 _{0.23}	+1.90
Luke	94.25 _{0.21}	96.01 _{0.15}	+1.87

Table 10: Detailed F_1 scores of the models on the CoNLL-2003 and the CoNLL-2003' test set.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 8
- A2. Did you discuss any potential risks of your work?
Not applicable. Left blank.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract at the beginning of the paper, and introduction in Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
Table 2
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. Left blank.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Our task is on named entity recognition, for which having actual names in the dataset is more desirable
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Not applicable. Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Table 1

C Did you run computational experiments?

Section 4 and 5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4.1 does discuss the number of parameters in some of the models used. However, number of parameters in other models, total computational budget and computing infrastructure are largely irrelevant to our paper

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Experimental setup in Section 3, and discussion on hyperparameter in Appendix B
- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
Section 3, 4 and 5
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Table 2
- D** **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
Section 2
- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
The annotations were completed by two authors of the paper. No additional annotators were hired.
- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
Not applicable. Left blank.
- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
Not applicable. Left blank.
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Not applicable. Left blank.
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Not applicable. Left blank.