

# Transfer Learning from Transformers to Fake News Challenge Stance Detection (FNC-1) Task

Valeriya Slovikovskaya, Giuseppe Attardi

University of Pisa

<https://www.unipi.it>

v.slovikovskaya@studenti.unipi.it, attardi@di.unipi.it

## Abstract

Transformer models, trained and publicly released over the last couple of years, have proved effective in many NLP tasks. We wished to test their usefulness in particular on the stance detection task. We performed experiments on the data from the Fake News Challenge Stage 1 (FNC-1). We were indeed able to improve the reported SotA on the challenge, by exploiting the generalization power of large language models based on Transformer architecture. Specifically (1) we improved the FNC-1 best performing model adding BERT sentence embedding of input sequences as a model feature, (2) we fine-tuned BERT, XLNet, and RoBERTa transformers on FNC-1 extended dataset and obtained state-of-the-art results on FNC-1 task.

**Keywords:** fake news detection, transfer learning, fine-tuning

## 1. Introduction

Two years ago, the Fake News Challenge Stage 1 (FNC-1) attracted the attention of several members of the linguistics research community: fifty teams participated and submitted their results and some of them like (Riedel et al., 2017) released their code and published research papers. Later, (Hanselowski et al., 2018) retrospectively examined the top submissions<sup>1</sup>, challenge task formulation, dataset characteristics as well as baseline model and evaluation metric proposed by organizers. Their contribution included revised FNC score, extended dataset and novel classification model that outperforms the winner’s system on class-wise F1 scores. Besides, they also determined upper bound for FNC-1 data classification conducting an annotation experiment that involved five human raters, who manually labelled 200 data instances. This detailed analysis of the task, together with the code released by the authors, are highly valuable as a baseline on which we can rely in order to obtain further improvements.

In the last two years since Fake News Challenge Cup, significant improvements have occurred in NLP technology, in particular with the development of large language models using contextualized word embeddings based on the Google Transformer architecture (Vaswani et al., 2017).

Applying transfer learning techniques to these general purpose pre-trained models helped to achieve improvements in a wide range of downstream tasks. Using these models as pre-training for downstream tasks allows lowering the amount of annotated data required for such tasks.

BERT (Devlin et al., 2018), GPT (Radford et al., 2019), XLNet (Yang et al., 2019), trained on the huge unlabeled datasets on GPUs and Cloud TPUs, attained significant accuracy increase on GLUE (Wang et al., 2018) and RACE (Lai et al., 2017) benchmarks<sup>2</sup>.

The public releases of pre-trained models led to intensi-

fied experimenting with transfer learning / fine-tuning for a broader range of linguistic tasks; this stimulated the development of APIs that unify and standardize the access to weights of different pre-trained models and made them accessible for the greatly extended community of professionals.

This paper is organized as follows. In the first part, we present Fake News Challenge Stage 1 (FNC-1) revised task: task formulation, dataset, evaluation metric, and the state-of-the-art model. In the second part, we introduce new semantic features that allow to improve the model performance. In the third part we present the further improved results obtained as a fruit of transfer learning from BERT, RoBERTa and XLNet transformers applied to the FNC-1 classification task.

## 2. FNC-1 Task

### 2.1. Task formulation

Let’s suppose we have a statement or a claim which truthfulness is verified or which source is trusted. How can we then detect liars about such claim in an upcoming news stream? We might try to determine whether the news text is related to the claim, and, if so, to find out whether it contradicts or supports it, i.e. its stance on the subject. The Fake News Challenge Stage 1 task is formulated in fact as a stance detection problem. The data set provided by organizers consists of pairs of a news article headline (as a claim), and a snippet of text taken either from the same or from another news article. Each pair in the training part of the data set is labeled with a tag attesting the relation between the headline and the snippet. The labels are “agree”, “disagree”, “discuss”, and “unrelated”. Table 1 presents a few examples of stance classification, taken from the training set of the challenge. The goal of the challenge is to classify the pairs in the test data set.

### 2.2. Dataset

#### 2.2.1. FNC-1 Data

The organizers of the Challenge constructed a news corpus with data from the Emergent dataset for stance classifica-

<sup>1</sup>including that of Hanselowski team Athene which ranked second

<sup>2</sup>GLUE leaderboard: <https://gluebenchmark.com/leaderboard/>, RACE leaderboard: [http://www.qizhexie.com/data/RACE\\_leaderboard.html](http://www.qizhexie.com/data/RACE_leaderboard.html)

Headline	Robert Plant Ripped up \$800M Led Zeppelin Reunion Contract
agree	Led Zeppelin’s Robert Plant turned down £500 MILLION to reform supergroup
disagree	No, Robert Plant did not rip up an \$800 million deal to get Led Zeppelin back together
discuss	Robert Plant reportedly tore up an \$800 million Led Zeppelin reunion deal
unrelated	Richard Branson’s Virgin Galactic is set to launch SpaceShipTwo today

Table 1: Headline and text snippets from document bodies with respective stances from the FNC-1 dataset

tion that contains 300 claims, and 2,595 associated article headlines.

The claims were collected by journalists from a variety of sources such as rumour sites, e.g. Snopes.com, and Twitter accounts such as @Hoaxalizer. Their subjects include world and national U.S. news and technology stories. The journalist also summarises the article into a headline. In parallel to the article-level stance detection, a claim-level veracity judgement is reached as more articles associated with the claim are examined.

The dataset construction was conducted by Craig Silverman and his colleagues at the Tow Center for Digital Journalism at Columbia University<sup>3</sup> for Emergent Project (Silverman, 2015). Then it was used for independent research on rumor debunking and released<sup>4</sup> as Emergent dataset by (Ferreira and Vlachos, 2016)

Thus the Emergent Project dataset consists of 300 topics, each represented by a claim with 5–20 possibly related news articles. The labelling of a headline/article pair indicates whether the article is supporting, refuting, or just reporting the claim, respectively.

Other than for rumor debunking, the FNC-1 organizers match each document with every summarized headline, and then label the pair with one of the four stance labels: "agree", "disagree", "discuss", and "unrelated". To generate the "unrelated" class, headlines and articles belonging to different topics are randomly matched. Headline/article pairs of 200 claims/topics are reserved for training, the remaining headline/article pairs of 100 claims/topics for testing. Claims, headlines, and articles are therefore not shared between the two data splits. To prevent teams from using any unfair means by deriving the labels for the test set from the publicly available Emergent data, the organizers additionally created 266 instances. Table 2 shows the corpus size and label distribution.

### 2.2.2. ARC Data

To test the robustness of their models (i.e. how well they generalize to new datasets), (Hanselowski et al., 2018) introduced a novel test data for document-level stance detection based on the Argument Reasoning Comprehension (ARC) task proposed by (Habernal et al., 2017).

(Habernal et al., 2017) manually selected 188 debate topics with popular questions from the user debate section of the New York Times. For each topic, they collected user posts, which are highly ranked by other users, and created two claims representing two opposing views on the topic. Then, they asked crowd workers to decide whether a user post supports either of the two opposing claims or does not express a stance at all. This Argument Reasoning Comprehension (ARC) dataset consists of typical controversial topics from the news domain, such as immigration, schooling issues, or international affairs.

While this is similar to the FNC-1 dataset, there are significant differences, as a user post is typically a multi-sentence statement representing one viewpoint on the topic. In contrast, the news articles of FNC-1 are longer and usually provide more balanced and detailed perspective on the issue.

To allow using the ARC data for the FNC-1 stance detection setup (Hanselowski et al., 2018), consider each user post as an article and randomly select one of the two claims as the headline. They label the headline/article pair as "agree" if the claim has also been chosen by the workers, as "disagree" if the workers chose the opposite claim, and as "discuss" if the workers selected neither claim. Table 3 shows an example of the revised ARC corpus structure.

In order to generate the unrelated instances (Hanselowski et al., 2018), randomly match the user posts with claims, but avoid that a user post is assigned to a claim from the same topic. For training and testing, the authors split the corpus into 80% training/validation set and 20% testing set. Table 4 provides basic corpus statistics.

For all our experiments we use the combined FNC + ARC dataset made freely available by (Hanselowski et al., 2018) on GitHub.

### 2.3. Evaluation Metric

The FNC-1 dataset is highly unbalanced with respect to class distribution. The FNC-1 organizers proposed the hierarchical evaluation metric, FNC score, which first awards .25 points if an article is correctly classified as related or "unrelated" to a given headline. If it is related, .75 additional points are assigned if the model correctly classifies the headline/article pair as "agree", "disagree", or "discuss". The goal of this weighting schema is to balance out the large number of unrelated instances.

This metric was criticized by (Hanselowski et al., 2018) as the score that, albeit an hierarchical, fails to take into account the imbalanced class distribution of the three related classes "agree", "disagree", and "discuss".

(Hanselowski et al., 2018) reasoned that the models which perform well on the majority class and poorly on the minority classes are favored. Since it is not difficult to separate related from unrelated instances (the best systems reach about  $F1 = .99$  for the "unrelated" class), a classifier that just randomly predicts one of the three related classes would already achieve a high FNC score. A classifier that always predicts "discuss" for the related documents even reaches  $FNC = .833$ , which is even higher than the top-ranked FNC system.

Authors therefore argued that the FNC score metric is not appropriate for the task and proposed to rely on the class-

<sup>3</sup><http://towcenter.org/>

<sup>4</sup><https://github.com/willferreira/mscproject>

Dataset	headlines	documents	instances	agree	disagree	discuss	unrelated
FNC-1	2,587	2,587	75,385	7.4%	2.0%	17.7%	72.8%

Table 2: Corpus statistics and label distribution for the FNC-1 dataset

Example from the original ARC dataset		
Topic	Do same-sex colleges play an important role in education, or are they outdated?	
User post	Only 40 women’s colleges are left in the U.S. And, while there are a variety of opinions on their value, to the women who have attended ... them, they have been ... tremendously valuable. ...	
Claims	1. Same-sex colleges are outdated 2. Same-sex colleges are still relevant	
Label	Same-sex colleges are still relevant	
Generated instance in alignment with the FNC-1 problem setting		
Stance	Headline	Document
agree	Same-sex colleges are still relevant	Only 40 women’s colleges are left in the U.S. ...

Table 3: Example of the original ARC dataset and the generated instance to align with FNC-1 dataset

Dataset	headlines	documents	instances	agree	disagree	discuss	unrelated
ARC	4,448	4,448	17,792	8.9%	10.0%	6.1%	75.0%

Table 4: Corpus statistics and label distribution for the ARC dataset

wise and the macro-averaged F1 scores (F1m) not affected by the large size of the majority class. The naïve approach of perfectly classifying "unrelated" and always predicting "disagree" for the related classes would achieve only F1m = .444 and would reveal a poor ability of the model to distinguish between "agree", "disagree", and "discuss" categories.

## 2.4. Model Architecture

With respect to F1m score, the model of team Athene (Hanselowski et al., 2017) outperforms all other models. Its architecture, named "featMLP", is a multi-layer perceptron (MLP) inspired by the work of (Davis and Proctor, 2017). The model has six hidden and a softmax layers and incorporates multiple hand-engineered features. The one group of features includes unigrams, and the cosine similarity of word embeddings of nouns and verbs between headline and document tokens. The other group consists of topic models based on non-negative matrix factorization, latent Dirichlet allocation, and latent semantic indexing. featMLP also incorporates the FNC-1 baseline’s features suggested by the FNC-1 organizers such as word overlap, polarity words, refuting words, and co-occurrence feature that counts how many times word 1-/2-/4-grams, character 2-/4-/8-/16-grams, and stop words of the headline appear in the first 100, first 255 characters of the article, and how often they appear in the article overall. See (Hanselowski et al., 2018) and baseline (Galbraith et al., 2016) for major details. Depending on the feature type, they either form separate feature vectors for document and headline, or a joint feature vector.

## 2.5. Lexical Features

(Hanselowski et al., 2018) conduct the feature ablation test on the FNC-1 development set with 10-fold cross-validation to define the best feature set for their model. This feature set contains bag-of-word (BoW) and bag-of-character (BoC) features as well as topic modeling features. The BoW features are uni- and bi-grams with 5,000 tokens vocabulary for the headlines and articles. Based on a technique by (Das and Chen, 2007) a negation tag "NEG\_" is added as a prefix to every word following negation keywords (e.g. "not", "never", "no") until the next punctuation mark appears. The BoC features are three-grams with 5,000 tokens vocabulary. The set of winning Bow/BoC features also includes FNC-1 baseline’s co-occurrence feature.

As for topic modeling feature group, non-negative matrix factorization (NMF) (Lin, 2007), latent semantic indexing (LSI) (Deerwester et al., 1990), and latent Dirichlet allocation (LDA) (Blei et al., 2003) are used to create topic models and independent features out of them. For each topic model, (Hanselowski et al., 2018) extract 300 topics from the headlines and article texts, then compute the similarity for each (headline, topic) and (article, topic) pair and either concatenate the headline and article feature vectors (for NMF and LSI) or calculate the cosine distance between them as a single valued feature (for NMF and LDA).

Based on these features the model exploits similarity between the headline and the article in terms of lexical overlap without really capturing the semantics of the text.

## 2.6. Semantic Features

Stepping towards taking into account semantic characteristics of the text, (Hanselowski et al., 2018) experiment with "stackLSTM" model, which combines the best fea-

ture set found in the ablation test with a stacked long short-term memory (LSTM) network (Hermans and Schrauwen, 2013).

Headlines and first 100 tokens of article texts are vectorized using 50-dimensional GloVe word embeddings<sup>5</sup> (Pennington et al., 2014) and then are fed through two stacked LSTMs (with a hidden state size of 100 with a dropout of 0.2 each). The last hidden state of the second LSTM is concatenated with the feature set and fed into a 3-layer neural network with 600 neurons each. A dense layer with four neurons and softmax activation is put on top of the network in order to retrieve the class probabilities.

The "stackLSTM" improves performance for the "disagree" class, which is the most difficult one to predict due to the low number of instances.

### 3. FNC-1 Top Model Improved

In our work on improving FNC-1 results we were looking for semantically sensitive model architectures and features. Therefore, we experimented with transfer learning from large pre-trained models reported as being able to capture semantics of the text.

There are two existing strategies for applying pre-trained language representations to downstream tasks: feature-based and fine-tuning. The feature-based approach uses task-specific architectures that include the pre-trained representations as additional features.

The fine-tuning does not require task specific model, instead it introduces into a general purpose pre-trained model a restricted set of task specific parameters and trains it on the downstream task dataset, in general for 2–3 epochs, fine-tuning all pre-trained parameters.

We apply both approaches to FNC-1 task, and report them both to be effective.

#### 3.1. InferSent-based Features

(Conneau et al., 2017) indicated the suitability of natural language inference (NLI) for transfer learning to other NLP tasks: NLI is a high-level understanding task that involves reasoning about the semantic relationships within sentences. Authors trained universal sentence representations using the supervised data of the Stanford Natural Language Inference (SNLI) dataset (Bowman et al., 2015) and made their encoder, based on a bi-directional LSTM architecture with max pooling, publicly available so that sentence embedding can be easily obtained and transferred as a feature to other text classification tasks.

Thus, for each headline/article body pair of FNC-ARC dataset, we calculated vector representation for every sentence of the article or headline using Facebook InferSent encoder (Conneau et al., 2017), then we averaged sentence vectors to obtain vector representation for whole article and whole headline, after that we concatenated the article and the headline vectors. We used the headline/article pair vectors as a feature on its own to directly feed into (Hanselowski et al., 2018) "featMLP" classifier and in addition, as a separate features, we calculated a cosine similarity score between headline and article vectors as

<sup>5</sup><http://nlp.stanford.edu/data/glove.twitter.27B.zip>

well as a maximum similarity score between headline and every sentence in the article.

These three features together slightly improve F1 macro and F1 -"agree", -"disagree" and -"discuss" scores of "featMLP" model as shown in Table 5. "Inf1" column reports the impact of InferSent embeddings for input sequence pair, "Inf3" shows the impact of this feature together with two similarity scores for two sequences in the input pair.

#### 3.2. BERT-based Features

Our next step was experimenting with transfer learning from pre-trained transformers.

(Devlin et al., 2018) introduced a new language representation model called BERT, which stands for **B**idirectional **E**ncoder **R**epresentations from **T**ransformers. BERT's model architecture is a multi-layer bidirectional Transformer encoder based on the original implementation described in (Vaswani et al., 2017) and released in the tensor2tensor library (Vaswani et al., 2018). BERT is designed to pre-train deep bidirectional representations from unlabeled text by jointly conditioning on both left and right context in all layers. BERT uses "masked language model" (MLM) pre-training objective, inspired by the Cloze task (Taylor, 1953). Masked language model randomly masks some of the tokens from the input, and the objective is to predict the masked word based on its context. In addition to the masked language model, BERT also use a "next sentence prediction" (NSP) task that jointly pre-trains text-pair representations.

For the pre-training corpus (Devlin et al., 2018) use the BooksCorpus (800M words) (Zhu et al., 2015) and English Wikipedia (2,500M words).<sup>6</sup>

Self attention mechanism in the Transformer allows BERT to model many downstream tasks – whether they involve single text or text pairs. For each task, the steps are: (1) simply plug in the task specific inputs and outputs into BERT and (2) fine-tune all the parameters end-to-end.

Authors showed that fine-tuned BERT outperforms all systems on all GLUE tasks by a substantial margin: BERT-base obtains a 4.5% average accuracy improvement over the prior state of the art, and BERT-large achieves 7.0%.

(Devlin et al., 2018) report fine-tuning to be relatively inexpensive compared to pre-training: at most 1 hour on a single Cloud TPU, or a few hours on a GPU, starting from the same pre-trained model. Authors pre-train BERT model of two sizes: BERT-base with 12 transformer blocks, 12 attention heads, 768 neurons in hidden layers, and 110 million parameters, and BERT-large with 24 transformer blocks, 16 attention heads, 1024 neuron in hidden layers, and 340 million parameters. Both models are made publicly available and their success raised the popularity of libraries like HuggingFace's Transformers (Wolf et al., 2019) that facilitate access to pre-trained models and optimize their integration into NLP pipelines.

<sup>6</sup>Authors note that it is critical to use a document-level corpus rather than a shuffled sentence-level corpus such as the Billion Word Benchmark (Chelba et al., 2013) in order to extract long contiguous sequences.

Score, %	Base	Inf1	Inf3	BERT1	BERT3	BERT +Inf3	BERT only
Accuracy	87.18	86.93	87.05	87.57	87.75	<b>87.94</b>	81.45
FNC score	78.59	77.82	78.16	79.03	79.42	<b>79.72</b>	68.65
F1 macro	61.61	61.63	62.02	63.32	<b>63.96</b>	63.91	56.27
F1 agree	50.26	50.66	50.69	53.77	53.15	53.58	43.16
F1 disagree	27.22	27.81	28.85	31.43	<b>34.04</b>	31.22	31.75
F1 discuss	72.43	71.70	72.20	71.21	71.61	<b>73.88</b>	58.55
F1 unrelated	96.54	96.35	96.36	96.89	<b>97.05</b>	96.95	91.61

Table 5: Performance of featMLP model without and with InferSent- and BERT- based features

In our experiments, we follow a feature-based approach first to transfer learning and use BERT sentence embeddings to create the new features for "featMLP" model. As with InferSent sentence embeddings, we introduce three separate features: we calculate BERT sentence embeddings for each sentence of an article and a headline, then we average sentence embeddings to obtain vector representation for the whole article and the whole headline, after that we concatenate the article and the headline vectors. We use the headline/article pair vectors as a feature on its own, and, in addition, as two separate features, we calculate the cosine similarity score between a headline and article vectors, and the maximum similarity score between headline and each sentence of the article.

To obtain BERT sentence embeddings we leverage the "bert-as-service" application (Xiao, 2018)<sup>7</sup>. "bert-as-service" is build on pre-trained 12/24-layer BERT models released by Google AI, it uses BERT as a sentence encoder and hosts it as a service via ZeroMQ. The application includes build-in HTTP server and a dashboard; it provides asynchronous encoding and multicasting.

The performance of the "featMLP" model enriched with BERT-based features is shown in Table 5. "BERT1" column shows the impact of BERT vector representation of an input, and "BERT3" shows the impact of this feature together with two similarity scores between BERT embeddings of two sequences in the input pair. The column "BERT3 + Inf3" reports InferSent and BERT cumulative effect. The most significant F1-score increase, that of 7% is attained for most difficult "disagree" class.

To demonstrate the power of BERT embeddings we conduct a kind of ablation test: we remove all Bow/BoC and topic modelling features and use only headline/article pair BERT vectors (without similarity scores), this way we obtaine the results shown in the "BERT1 only" column of Table 5: it can be seen that the "disagree" F1 score improvement is relied mostly on the BERT embedding feature.

These results pushed us further on the way of experimenting with pre-trained transformers. Following (Devlin et al., 2018) we fine-tuned on FNC-1 task three pre-trained models: BERT, XLNet and RoBERTa, and as expected, received substantially improved F1 scores with respect to (Hanselowski et al., 2018) "featMLP" model performance.

<sup>7</sup>with default settings listed here: <https://github.com/hanxiao/bert-as-service>

## 4. Transfer Learning: from Transformers to FNC-1 Task

### 4.1. XLNet

BERT predicts all masked positions independently, meaning that during the training, it does not learn to handle dependencies between predicted masked tokens. This reduces the number of dependencies BERT learns at once, making the learning signal weaker than it could be.

Another problem with BERT is that the [MASK] token – which is at the center of training BERT – never appears when fine-tuning BERT on downstream tasks. That means that the [MASK] token is a source of train-test skew while fine-tuning.

XLNet (Yang et al., 2019) incorporates bidirectional context while avoiding the [MASK] tokens and independent predictions. It does this by introducing "permutation language modeling": instead of predicting the tokens in sequential order, it predicts tokens in some random order.<sup>8</sup>

Aside from using permutation language modeling, XLNet improves upon BERT by using the Transformer XL as its base architecture.

Both BERT and XLNet can take the pair of text sequences as an input. To enable the model to distinguish between words in two different segments, BERT learns a segment embedding. In contrast, XLNet learns an embedding that represents whether two words are from the same segment. This embedding is used during attention computation between any two words.

XLNet integrates the segment recurrence mechanism and relative encoding scheme of TransformerXL (Dai et al., 2019) into pre-training and enables the model to reuse hidden states from previous segments.<sup>9</sup>

XLNet outperforms BERT and achieves state-of-the-art performance across 20 tasks including text classification, question answering, natural language inference, duplicate sentence (question) detection, document ranking, corefer-

<sup>8</sup>For example, the sequence order is  $x_1, x_2, x_3, x_4$ . So for this 4 (N) tokens sentence, there are 24 (N!) permutations. If we want to predict the  $x_3$ , so there are 4 patterns with  $x_3$  in the 24 permutations:  $x_3$  is at the 1st position, 2nd position, 3rd position, and 4th position:  $[x_3, xx, xx, xx]$ ,  $[xx, x_3, xx, xx]$ ,  $[xx, xx, x_3, xx]$ ,  $[xx, xx, xx, x_3]$ . Here we set the position of  $x_3$  as t-th position and t-1 tokens are the context words for predicting  $x_3$ . Intuitively, the model will learn to gather information from all positions on both sides.

<sup>9</sup>XLNet is named after TransformerXL

ence resolution.

## 4.2. RoBERTa

(Liu et al., 2019) found that BERT was significantly undertrained and proposed an improved recipe for training BERT models, which they call RoBERTa (Robustly optimized BERT approach), that can match or exceed the performance of all of the post-BERT methods. The recipe includes: (1) training the model longer, with bigger batches, over more data; (2) removing the "next sentence prediction" objective; (3) training on longer sequences; and (4) dynamically changing the masking pattern applied to the training data.

To train RoBERTa authors use five English-language corpora of varying sizes and domains, totaling over 160GB of uncompressed text.<sup>10</sup>

They also demonstrate that removing the "next sentence prediction" loss together with segment-pair input format matches or slightly improves downstream task performance. RoBERTa achieves state-of-the-art results on all 9 of the GLUE task development sets. It consistently outperforms both BERT and XLNet.

## 4.3. Transformers Fine-tuned on FNC-1 Task

For our experiments on fine-tuning transformers on FNC-1 task we used the Simple Transformers (Rajapakse, 2019) wrapper around Hugging Face Transformers library (Wolf et al., 2019).

Score, %	featMLP	BERT	XLNet	RoBERTa
Accuracy	87.18	91.32	92.10	93.19
FNC score	78.59	86.16	87.90	89.17
F1 macro	61.61	72.82	76.06	78.12
F1 agree	50.26	64.70	68.61	70.71
F1 disagree	27.22	47.84	54.85	58.06
F1 discuss	72.43	80.07	82.10	84.57
F1 unrelated	96.54	98.68	98.65	99.16

Table 6: Performance of featMLP, BERT, XLNet, and RoBERTa fine-tuned models

The development of Transformers, a library for natural language processing with transfer learning models, was motivated by the need to share state-of-the-art pre-trained models and to increase the availability of published research code. With this library low-resource users can reuse pre-trained models without having to train them from scratch.

<sup>10</sup>These corpora include (1) BOOK CORPUS (Zhu et al., 2015) plus English Wikipedia, the original data used to train BERT (16GB); (2) CC-NEWS, which authors collected from the English portion of the CommonCrawl News dataset (Nagel, 2019), the data contains 63 million English news articles crawled between September 2016 and February 2019 (76GB after filtering); (3) OPEN WEB TEXT (Gokaslan and Cohen, 2019), an open-source recreation of the WebText corpus described in (Radford et al., 2019), the text is web content extracted from URLs shared on Reddit with at least three upvotes (38GB); (4) STORIES, a dataset introduced in (Trinh and Le, 2018) containing a subset of CommonCrawl data filtered to match the story-like style of Winograd schemas (31GB).

These models are accessed through a simple and unified API that follows a classic NLP pipeline: setting up configuration, processing data with a tokenizer and encoder, and using a model either for training (adaptation in particular) or inference. The model implementations provided in the library are tested to ensure they match the original author implementations' performances on various benchmarks. A list of architectures for which reference implementations and pre-trained weights are currently provided in Transformers includes BERT (Devlin et al., 2018), DistilBERT (Sanh et al., 2019), RoBERTa (Liu et al., 2019), XLNet (Yang et al., 2019), GPT and GPT2 (Radford et al., 2019).

The Simple Transformers (Rajapakse, 2019) library is built on top of the Hugging Face Transformers. The idea behind it was to make it as simple as possible, abstracting a lot of the implementation details.

Thus, with Simple Transformers on the shoulders of Hugging Face Transformers we could access pre-trained BERT, XLNet and RoBERTa in unified way without a lot of pre-processing coding.

We modified only limited number of parameters, setting maximum sequence length to be equal to 512 tokens<sup>11</sup>. We used the base versions of all transformers ("bert-base-uncased", "xlnet-base-cased", "roberta-base"), those with 12 transformer blocks. We trained every transformers for 5 epochs with batch size of 4, with learning rate of 3e-5 for BERT and 1e-5 for XLNet and RoBERTa.<sup>12</sup>

Our results, reported in Tables 6, 7, and 8 show that all three models fine-tuned on the Fake News Challenge task outperform (Hanselowski et al., 2018) "featMLP" model with significant margins varying from 8 to 20% for "related" classes. The XLNet and RoBERTa demonstrate better performance than BERT, that corresponds to our expectations; RoBERTa outperforms both BERT and XLNet, that is in line with the results reported in (Liu et al., 2019).

We also performed cross domain validation experiments fine-tuning transformers on FNC-1 dataset and evaluating them on ARC dataset and vice versa. The Table 9 shows that for these test RoBERTa can demonstrate worse performance in comparison with BERT and XLNet models.

## 5. Conclusion

We substantially improved Fake News Challenge Stage 1 results, putting the task in the line of Natural Language Processing tasks that are reported to benefit from transfer learning from pre-trained transformers.

## 6. Acknowledgments

The experiments were carried out on a server with 4 GPUs Nvidia Tesla Pascal P100, partially funded by the University of Pisa through the grant Grandi Attrezzature 2016.

## 7. Bibliographical References

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent Dirichlet Allocation. *J. Mach. Learn. Res.*, 3:993–1022.

<sup>11</sup>The maximum possible value to set given the parameters of pre-trained models

<sup>12</sup>Every time fine-tuning was taking several hours on Nvidia Tesla P100 GPU with 16GB of memory.

	agree	disagree	discuss	unrelated
agree	<b>1104 1534 1597 1592</b>	152 328 297 265	738 619 474 377	243 24 50 32
disagree	306 151 114 129	<b>230 432 517 560</b>	253 136 150 150	280 18 35 21
discuss	617 486 482 487	175 239 205 210	<b>3337 3732 3894 4009</b>	514 222 262 132
unrelated	129 66 44 29	64 70 50 34	244 156 125 107	<b>20586 20759 20676 20838</b>

Table 7: Confusion matrix: every cell shows featMLP, BERT, XLNet, RoBERTa result correspondingly

	Precision, %	Recall, %	F1-score, %	Support
0	<b>53 69 71 71</b>	<b>48 61 66 70</b>	<b>50 65 69 71</b>	2237 2505 2418 2266
1	<b>38 40 48 52</b>	<b>24 59 63 65</b>	<b>29 48 55 58</b>	1069 737 816 860
2	<b>73 80 84 86</b>	<b>74 80 80 83</b>	<b>73 80 82 85</b>	4643 4679 4843 4838
3	<b>95 99 98 99</b>	<b>98 99 99 99</b>	<b>97 99 99 99</b>	21023 21051 20895 21008

Table 8: Precision, recall, F1-score for featMLP, BERT, XLNet, RoBERTa. Every cell contains 4 results corresponding to four models in order

FNC - ARC				
F1 score, %	featMLP	BERT	XLNet	RoBERTa
agree	25.71	39.55	43.82	45.61
disagree	9.67	28.39	28.31	24.91
discuss	15.42	18.58	08.30	18.23
unrelated	89.83	93.48	92.53	93.91
ARC - FNC				
F1 score, %	featMLP	BERT	XLNet	RoBERTa
agree	27.57	39.03	41.01	0.73
disagree	10.43	8.41	05.44	16.10
discuss	8.03	5.66	14.07	10.40
unrelated	85.43	95.72	95.59	93.57

Table 9: Class-wise F1 scores for "featMLP" model (with BoW/BoC and topic modeling features) and fine-tuned transformers in cross-domain, FNC - ARC and ARC - FNC, experiments

Conneau, A., Kiela, D., Schwenk, H., Barrault, L., and Bordes, A. (2017). Supervised Learning of Universal Sentence Representations from Natural Language Inference Data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.

Dai, Z., Yang, Z., Yang, Y., Carbonell, J. G., Le, Q. V., and Salakhutdinov, R. (2019). Transformer-XL: Attentive Language Models Beyond a Fixed-Length Context. *CoRR*, abs/1901.02860.

Das, S. and Chen, M. Y. (2007). Yahoo! for Amazon: Sentiment Extraction from Small Talk on the Web. *Management Science*, 53(9):1375–1388.

Davis, R. and Proctor, C. (2017). Fake News, Real Consequences: Recruiting Neural Networks for the Fight Against Fake News. *web.stanford.edu*.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., and Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41(6):391–407.

Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2018).

BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR*, abs/1810.04805.

Habernal, I., Wachsmuth, H., Gurevych, I., and Stein, B. (2017). The Argument Reasoning Comprehension Task. *CoRR*, abs/1708.01425.

Hanselowski, A., PVS, A., Schiller, B., and Caspelherr, F. (2017). Description of the system developed by team Athene in the FNC-1. [https://github.com/hanselowski/athene\\_system/blob/master/system\\_description\\_athene.pdf](https://github.com/hanselowski/athene_system/blob/master/system_description_athene.pdf).

Hanselowski, A., PVS, A., Schiller, B., Caspelherr, F., Chaudhuri, D., Meyer, C. M., and Gurevych, I. (2018). A Retrospective Analysis of the Fake News Challenge Stance-Detection Task. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1859–1874, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Hermans, M. and Schrauwen, B. (2013). Training and Analysing Deep Recurrent Neural Networks. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems 26*, pages 190–198. Curran Associates, Inc.

Lin, C.-J. (2007). Projected Gradient Methods for Non-Negative Matrix Factorization. *Neural computation*, 19:2756–79.

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). RoBERTa: A Robustly Optimized BERT Pre-training Approach. *CoRR*, abs/1907.11692.

Nagel, S. (2019). Cc-news. <http://chuachinhon.pythonanywhere.com/>.

Pennington, J., Socher, R., and Manning, C. D. (2014). GloVe: Global Vectors for Word Representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. (2019). Language Models are Unsupervised Multitask Learners. preprint on webpage at <https://paperswithcode.com/paper/language-models-are-unsupervised-multitask>.

- Rajapakse, T. (2019). Simple Transformers. <https://github.com/ThilinaRajapakse/simpletransformers>.
- Riedel, B., Augenstein, I., Spithourakis, G. P., and Riedel, S. (2017). A simple but tough-to-beat baseline for the Fake News Challenge stance detection task. *CoRR*, abs/1707.03264.
- Sanh, V., Debut, L., Chaumond, J., and Wolf, T. (2019). DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter.
- Silverman, C. (2015). Emergent: A real-time rumor tracker. <http://www.emergent.info/>.
- Taylor, W. L. (1953). "Cloze Procedure": A New Tool for Measuring Readability. *Journalism Bulletin*, 30(4):415–433.
- Trinh, T. H. and Le, Q. V. (2018). A Simple Method for Commonsense Reasoning. *CoRR*, abs/1806.02847.
- Vaswani, A., Bengio, S., Brevdo, E., Chollet, F., Gomez, A. N., Gouws, S., Jones, L., Kaiser, L., Kalchbrenner, N., Parmar, N., Sepassi, R., Shazeer, N., and Uszkoreit, J. (2018). Tensor2Tensor for Neural Machine Translation. *CoRR*, abs/1803.07416.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention Is All You Need. *CoRR*, abs/1706.03762.
- Wang, A., andF Julian Michael, A. S., Hill, F., Levy, O., and Bowman, S. R. (2018). GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. *CoRR*, abs/1804.07461.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., and Brew, J. (2019). HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *ArXiv*, abs/1910.03771.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J. G., Salakhutdinov, R., and Le, Q. V. (2019). XLNet: Generalized Autoregressive Pretraining for Language Understanding. *CoRR*, abs/1906.08237.
- Zhu, Y., Kiros, R., Zemel, R. S., Salakhutdinov, R., Urtasun, R., Torralba, A., and Fidler, S. (2015). Aligning Books and Movies: Towards Story-like Visual Explanations by Watching Movies and Reading Books. *CoRR*, abs/1506.06724.
- Galbraith, B., Iqbal, H., van Veen, H., Rao, D., Thorne, J., and Pan, Y. (2016). Baseline FNC implementation. <https://github.com/FakeNewsChallenge/fnc-1-baseline>.
- Gokaslan, A. and Cohen, V. (2019). OpenWeb-Text Corpus. <http://Skylion007.github.io/OpenWebTextCorpus>.
- Hanselowski, A., PVS, A., Schiller, B., Caspelherr, F., Chaudhuri, D., Meyer, C. M., and Gurevych, I. (2018). Repository of the COLING 2018 paper: A Retrospective Analysis of the Fake News Challenge Stance Detection Task. [https://github.com/UKPLab/coling2018\\_fake-news-challenge](https://github.com/UKPLab/coling2018_fake-news-challenge).
- Lai, G., Xie, Q., Liu, H., Yang, Y., and Hovy, E. (2017). RACE: Large-scale ReAding Comprehension Dataset From Examinations. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.
- Rajapakse, T. (2019). Simple Transformers. <https://github.com/ThilinaRajapakse/simpletransformers>.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., and Brew, J. (2019). HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *ArXiv*, abs/1910.03771.
- Xiao, H. (2018). bert-as-service. <https://github.com/hanxiao/bert-as-service>.

## 8. Language Resource References

- Bowman, S. R., Angeli, G., Potts, C., and Manning, C. D. (2015). A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Chelba, C., Mikolov, T., Schuster, M., Ge, Q., Brants, T., and Koehn, P. (2013). One Billion Word Benchmark for Measuring Progress in Statistical Language Modeling. *CoRR*, abs/1312.3005.
- Conneau, A., Kiela, D., Schwenk, H., Barrault, L., and Bordes, A. (2017). InferSent. <https://github.com/facebookresearch/InferSent>.
- Ferreira, W. and Vlachos, A. (2016). Emergent: a novel data-set for stance classification. In *HLT-NAACL*.