Multimodality in fake news detection using neural networks and support vector machines

José David Nieto Vitola
jdnietov@unal.edu.co

Nicolás Campuzano Angulo
ncampuzanoa@unal.edu.co

Description

After researching the state-of-the-art methods available for fake news detection, we implemented improvements on two proposed solutions: first, the convolutional neural network FDNNet, using Facebook's FastText as a word embedding, and then tokenization for a TF-IDF heuristic that would feed a support vector machines. Both models were then merged into a multimodal model.

Dataset

We used Kaggle's ISOT Fake News Dataset. This dataset contains two types of articles, fake and real news. Real news were obtained by crawling articles from Routers.com, while the fake news were collected from different resources, flagged by Politifact and Wikipedia.

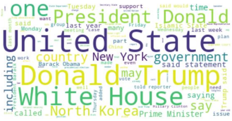

21.147 real news

23.481 fake news

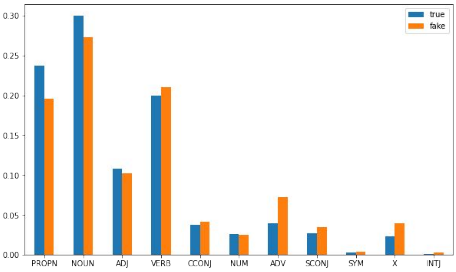

POS tag distribution across the dataset (with stopwords).

FDNNet + FastText

The deep neural network trained in this project was based on FDNNet's convolutional architecture. The original paper used GloVe as a word embedding; here, we use FastText because, unlike its predecessors like GloVe and Word2Vec FastText creates embeddings with character-level information of words, which makes it, for instance, immune to out-of-vocabulary terms.

Counterintuitively, FDNNet architecture obtains a good performance with less embedding size given by FastText vector representation. It allows us to create better and lightweight models to detect fake news.

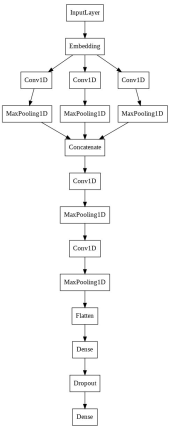

TABLE II
F-SCORES ON THE FDNNET + FASTTEXT EXPERIMENTS

Embedding size / Optimizer	100	200	300	400	500
Adam	0.939	0.897	0.870	0.843	0.808
Adamax	**0.939**	0.897	0.870	0.849	0.809
SGD	0.889	0.766	0.760	0.765	0.748
Adadelta	**0.939**	0.899	0.869	0.847	0.807

SVM

There were two alternative models based on SVM:

1. A model based on POS tag counting. We tagged the ISOT dataset using spaCy's built-in tokenizer and tagger, which is superior to average tokenizers because it is pretrained with named entity recognition. This model worked best including stopwords in the text.

2. A model based on BoW (bag-of-words) and TF-IDF (term frequency-inverse document frequency) metrics. This model worked best removing stopwords from the text.

TABLE V
F-SCORES FOR BOW AND TF-IDF MODELS

Heuristic / C	BoW	TF-IDF
0.125	**0.9961**	0.9942
0.25	0.9960	0.9952
0.5	0.9958	0.9962
1	0.9957	0.9968
2	0.9958	0.9971
4	0.9958	0.9972
8	0.9960	**0.9973**

Multimodality - Results

To combine the best two models obtained, we use two different logical gates. Both the AND an the OR gates give us better results than the standalone FDNNet architecture. However, to apply the logical gates we transform the FDNNet prediction from float values to binaries, with a threshold of 0.5. That means every value under 0.5 is mapped to 0. The resulting f-score metric in both gates have a good effect on improving FDNNet and regularizing the f-score in the TF-IDF model.

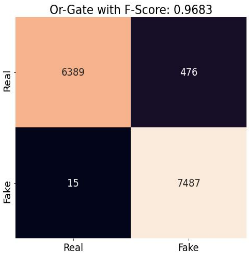

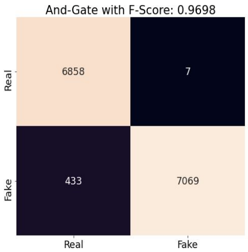

Explanation video: http://youtu.be/ahTxzK4dR48

Machine Learning 2020-I

UNIVERSIDAD NACIONAL DE COLOMBIA