

---

Machine Learning to Predict Sarcasm in Article headlines  
Siddarth Thota

**Abstract.**

This paper describes a machine learning model designed to assign sets of possible labels to article headlines. This was done by analyzing a dataset with article headlines from 2 sources: The Onion and The Huffpost. The Onion is a site known for generally sarcastic headlines, while the Huffpost has more genuine headlines. The data was then inputted into a Zero-Shot-Classification model, which labeled each headline as either with labels generally pertaining to "fake" or "real". This accuracy data is subsequently displayed with bar charts.

**Introduction**

Automation of headline categorization would no doubt save humans time and energy while creating an efficient approach. But can we trust machine learning algorithms to make these classifications? This work was created to assess the accuracy of a Zero-Shot-Classification model – one which is not trained in a task specific way – at classifying headlines into different sets of categories with the same general foundation.

**Methodology**


The dataset used was a collection of headlines compiled from two news sites: The Onion and The Huffpost. The dataset was chosen because of its dependability as The Onion is known for more satirical and fake titles, while the Huffpost displays realism. A for loop parses each line from the JSON file that contains the dataset. Each headline is then appended into the list input\_strs. From this list, data from the original file will be used throughout the program.

**Models**

A zero-shot classification model was used to predict the labels of the headlines. This neural network is known for its ability to classify data using general predictions although not being trained explicitly on said data. The 2d array contains pairs of labels (ex. ['funny', 'not funny']). A potential\_labels list initializes the potential assignments a headline could receive. Each label is given a score in terms of how well it would correspond to the particular headline. The label with the highest score for the particular headline is the one it is assigned to as a prediction.

(hypothetically a headline maybe predicted to be 'sarcastic' instead of 'not sarcastic' due to a higher numeric value being given to 'sarcastic').

In the main() an empty list accuracies = [] is defined. For each pair of labels (ex. ['genuine', 'not serious']), classification results are given, and compared with targets to calculate for accuracy. The results of the calculations are appended to the accuracies list. A bar plot is then created using the accuracy results.



## Results

The model was fairly accurate and consistent in terms of results. In a test, 3 label pairs were used: ['not sarcastic', 'sarcastic'], ['genuine', 'not serious'], ['real', 'fake']. When running `print(plt.bar(range(len(accuracies)), accuracies))`, it is shown that the pairs have a 54.46%, 52.36%, and 52.36% accuracy respectively when compared to the target/ given values.

## Conclusions

This model can be used to automate the task of headline categorization. It also shows the capabilities of pre-trained and general zero-shot classification models in categorization intensive tasks.

## Acknowledgements and References

Dataset:

<https://www.kaggle.com/datasets/rmisra/news-headlines-dataset-for-sarcasm-detection/data>

*Thank you to Efthimios G. for helping me in acquiring the data, helping me code, and reviewing this paper.*

