verifact.

FAKE NEWS DETECTION AND PREDICTING ITS DEGREE OF FACTUALITY

BY - Anish Sridharan, Purv Sogani, Swarup Laxmikant

what's going on?

a normal person -

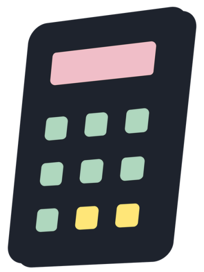

You see some news

See the same thing
on 10 different
platforms presenting it.

Don't know if it's true

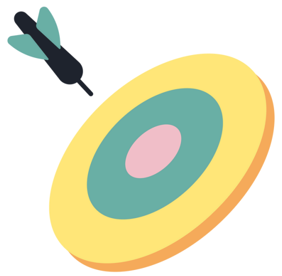

Take actions
Waste energy
Still Unsatisfied

bolo verifact.

remembering what needs to be done, so you don’t have to.

Here is your summary, sit back and chill.

News1: Fake
News2: Almost Fake
News3: Fair
News4: Likely to be true
News5: True

Proposed Solution

with rich and diverse data, this approach has potential beyond the walls of MLPR.

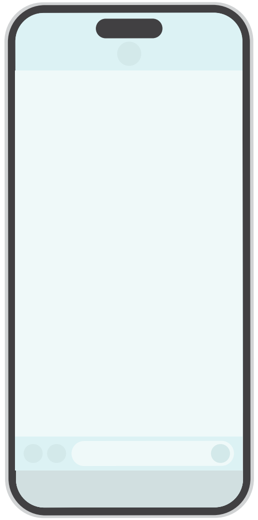

Applications

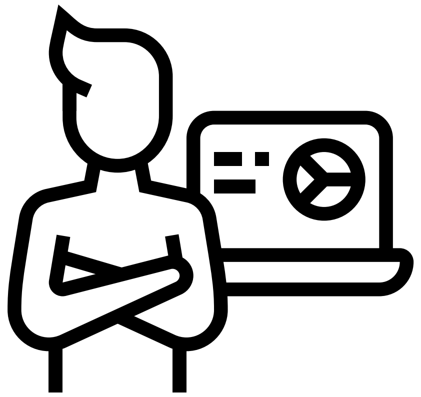

ChatBot

Web_Extensions

Lets talk about the Data

- Which Format?
- What is the goal?
- Looking for the best dataset for the model.
- Preprocessing
- Further Steps...

The problem - Datasets

The Need - Data

Labels according to the factness index

0: Half True
1: False
2- Mostly True
3- True
4: Barely True
5: Liar liar, pants on fire

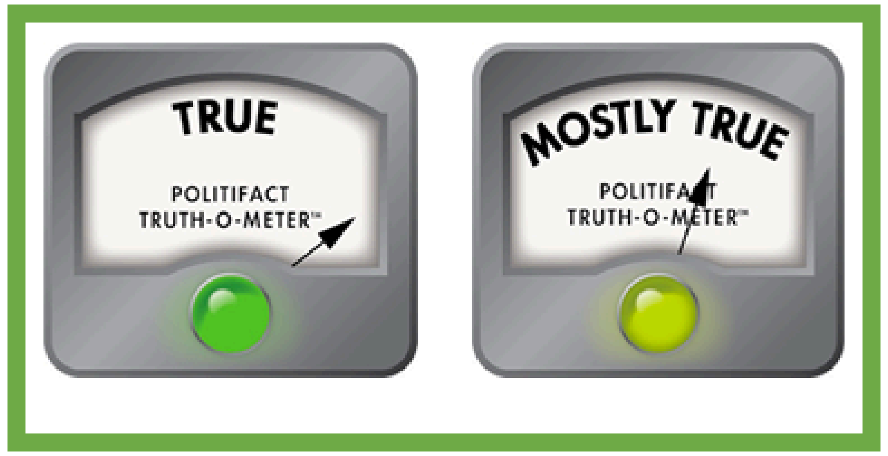

TRUE

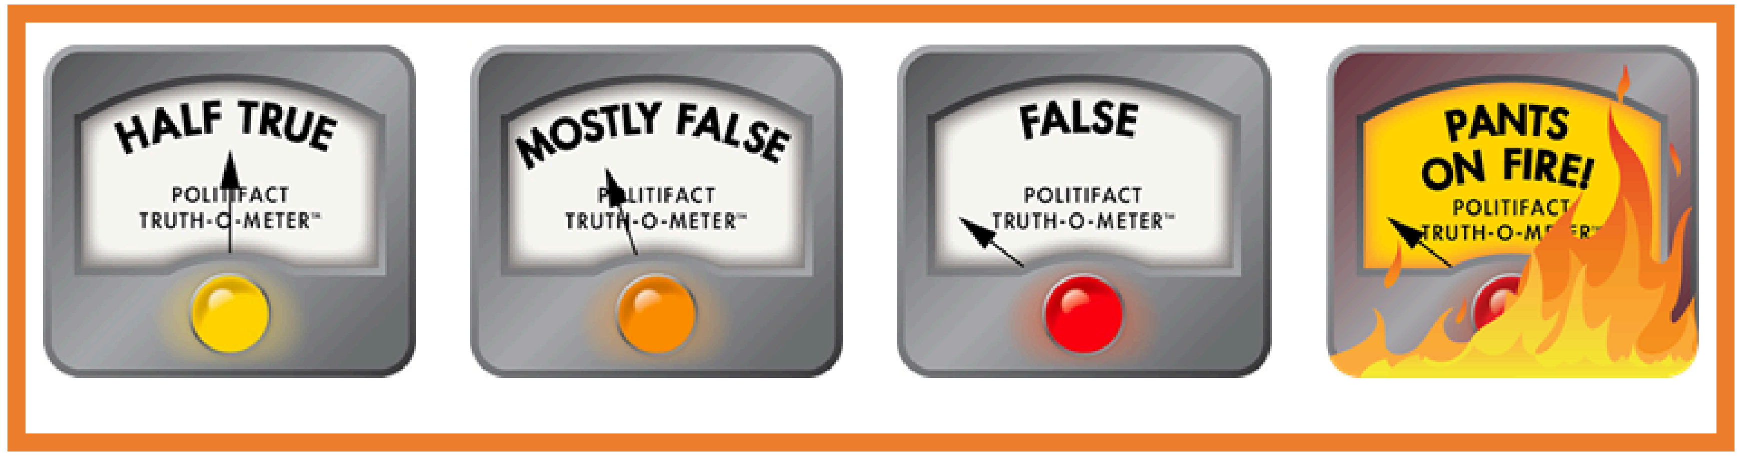

FALSE

Dataset - LIAR

"Liar, Liar Pants on Fire": A New Benchmark Dataset for Fake News Detection

- **Publically avalaible** for **free** (purposely for fake news detection)
- A **decade-long** of **12.8K** manually **labeled** short statements with **various contexts**
- Each statement is **evaluated** by a POLITIFACT.COM editor for its truthfulness.

A new benchmark dataset for fake news detection. (n.d.-a). https://aclanthology.org/P17-2067.pdf

Analysis

Words - Text

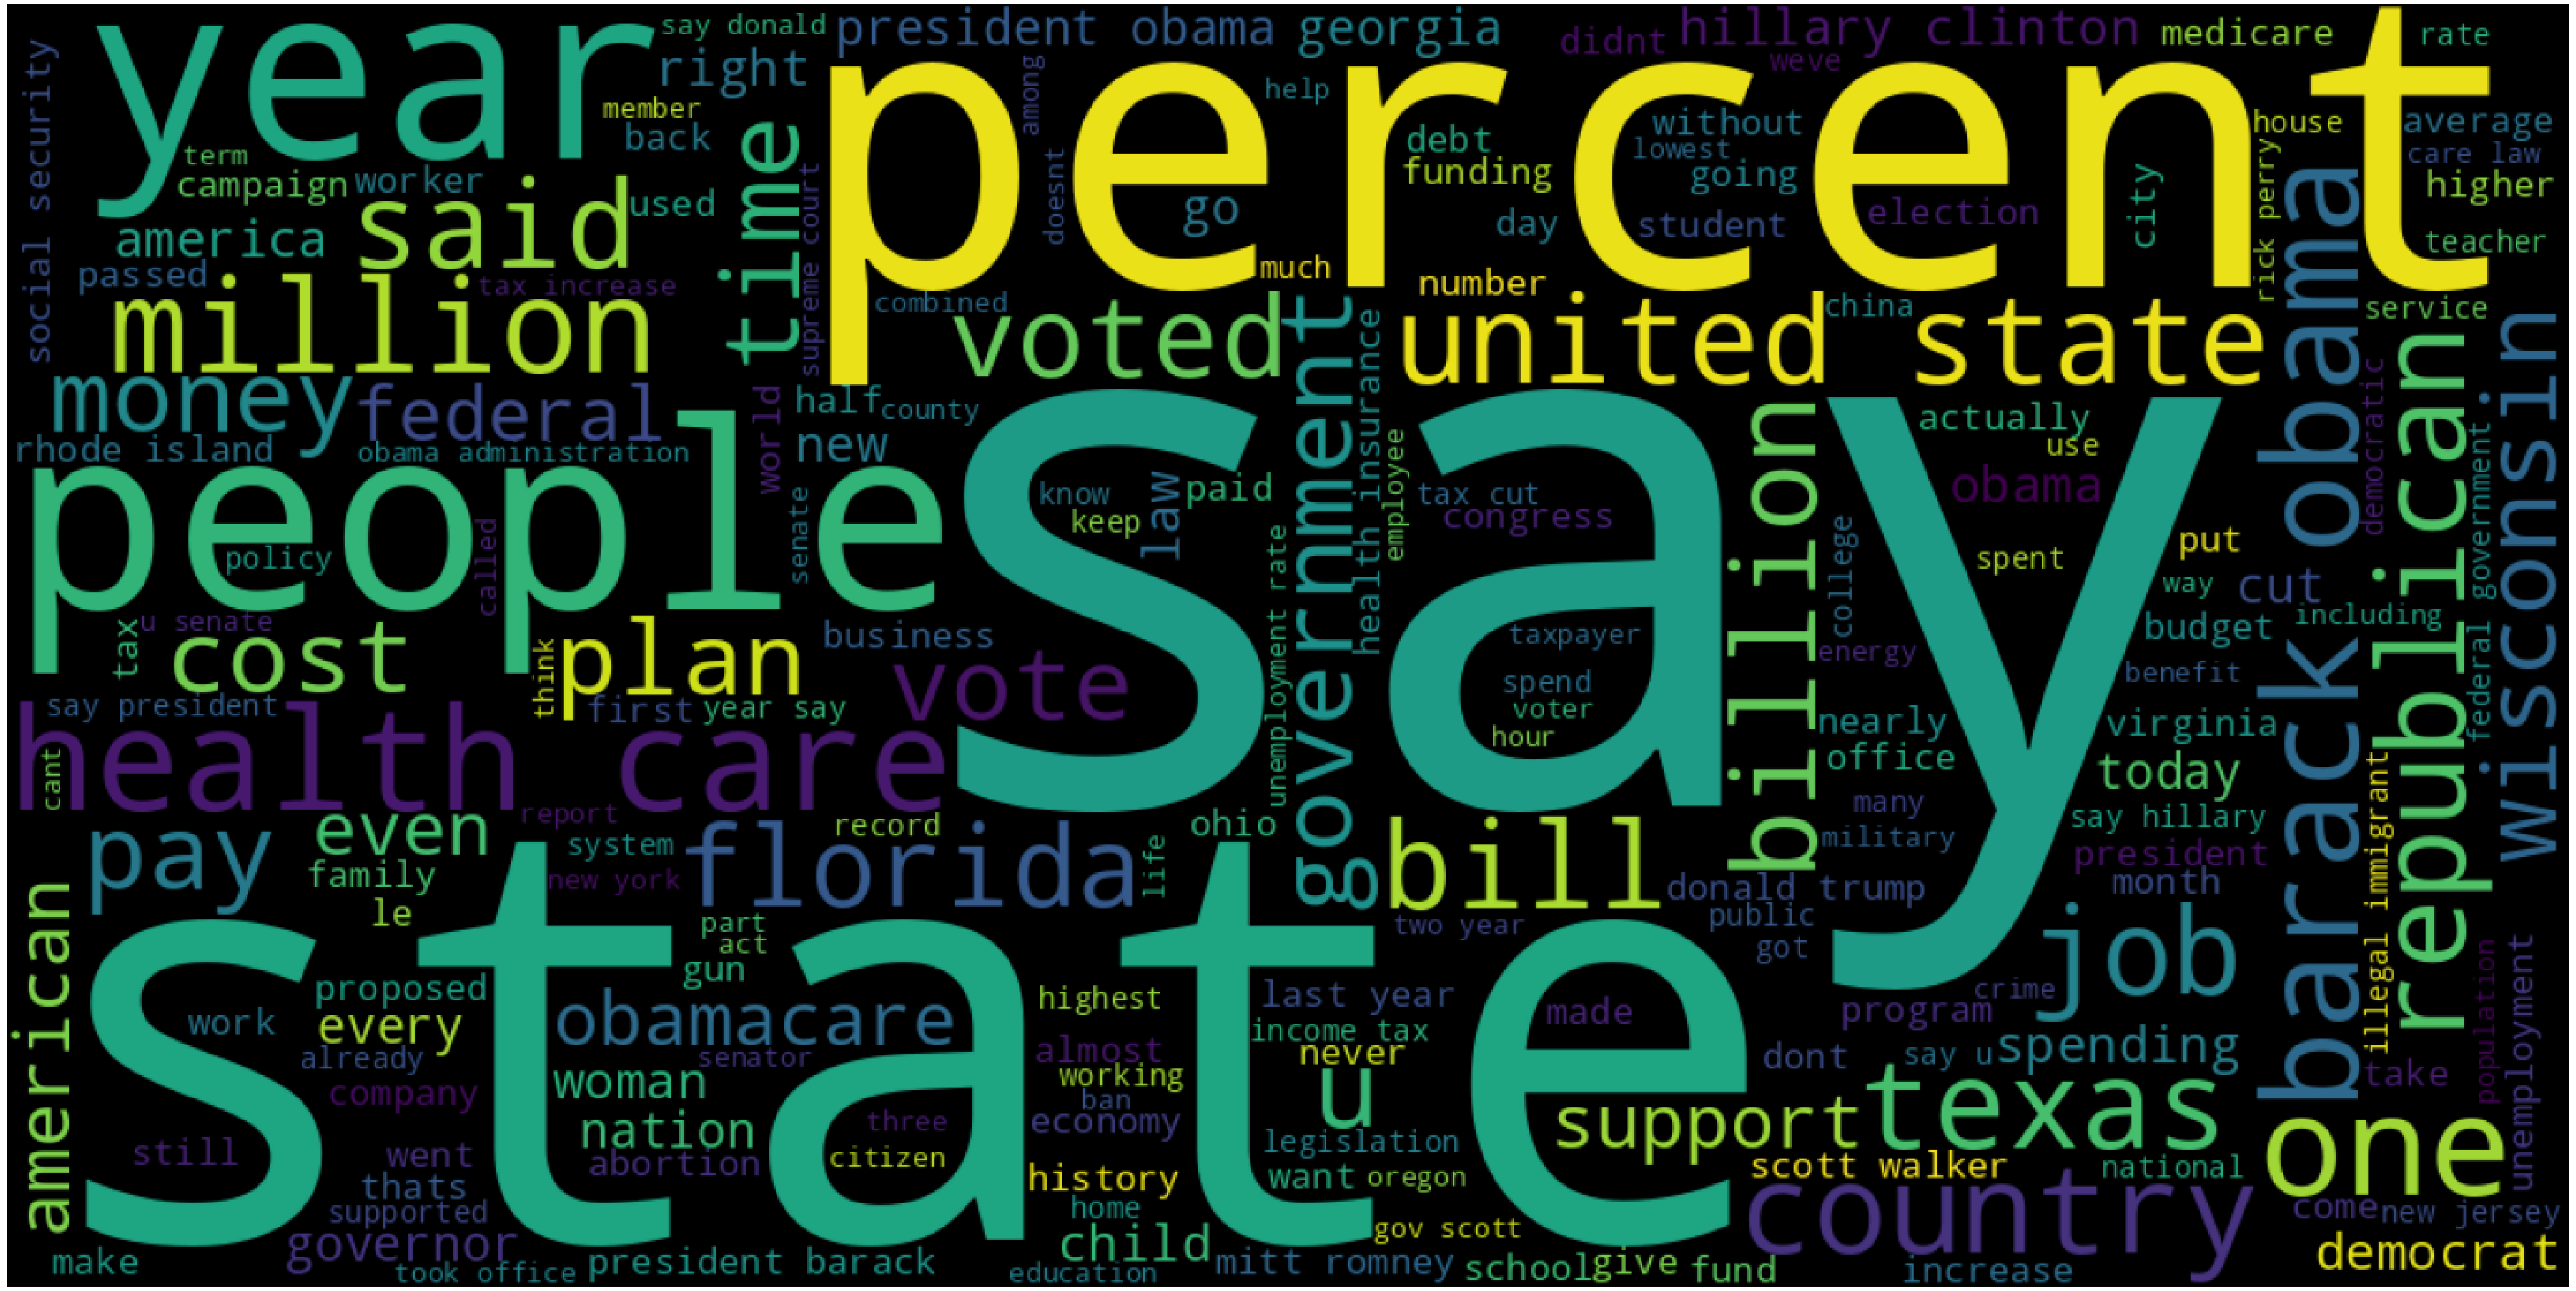

More Analysis

Words - Text_Tag

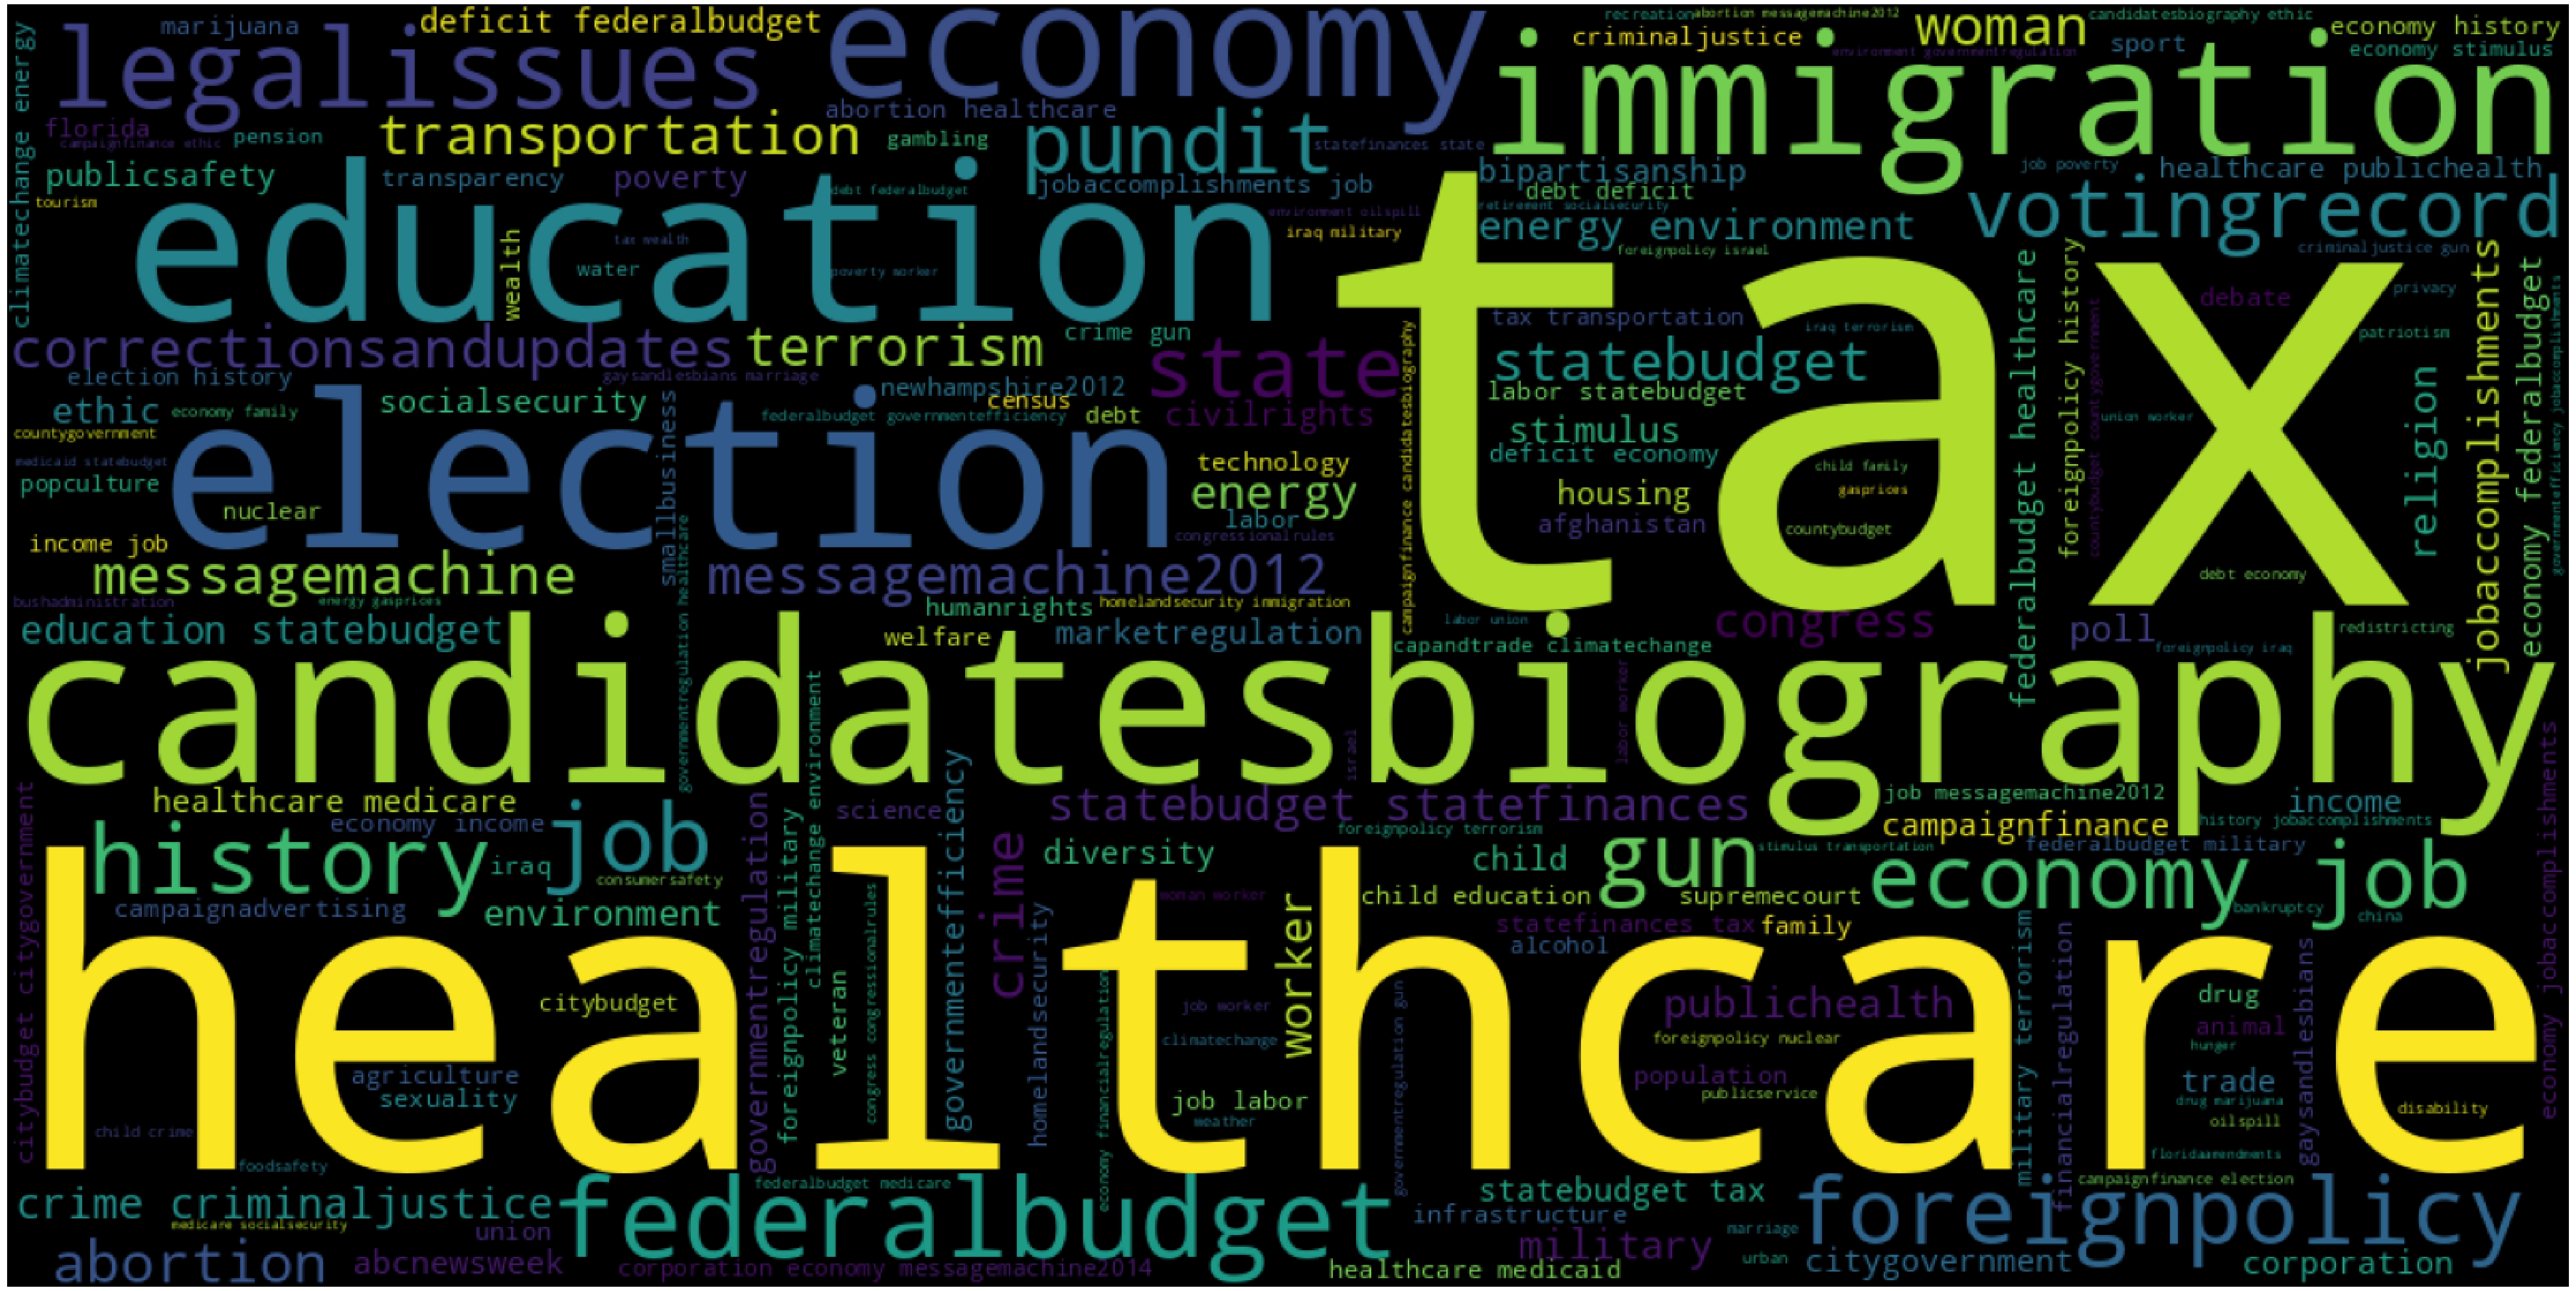

& More Analysis...

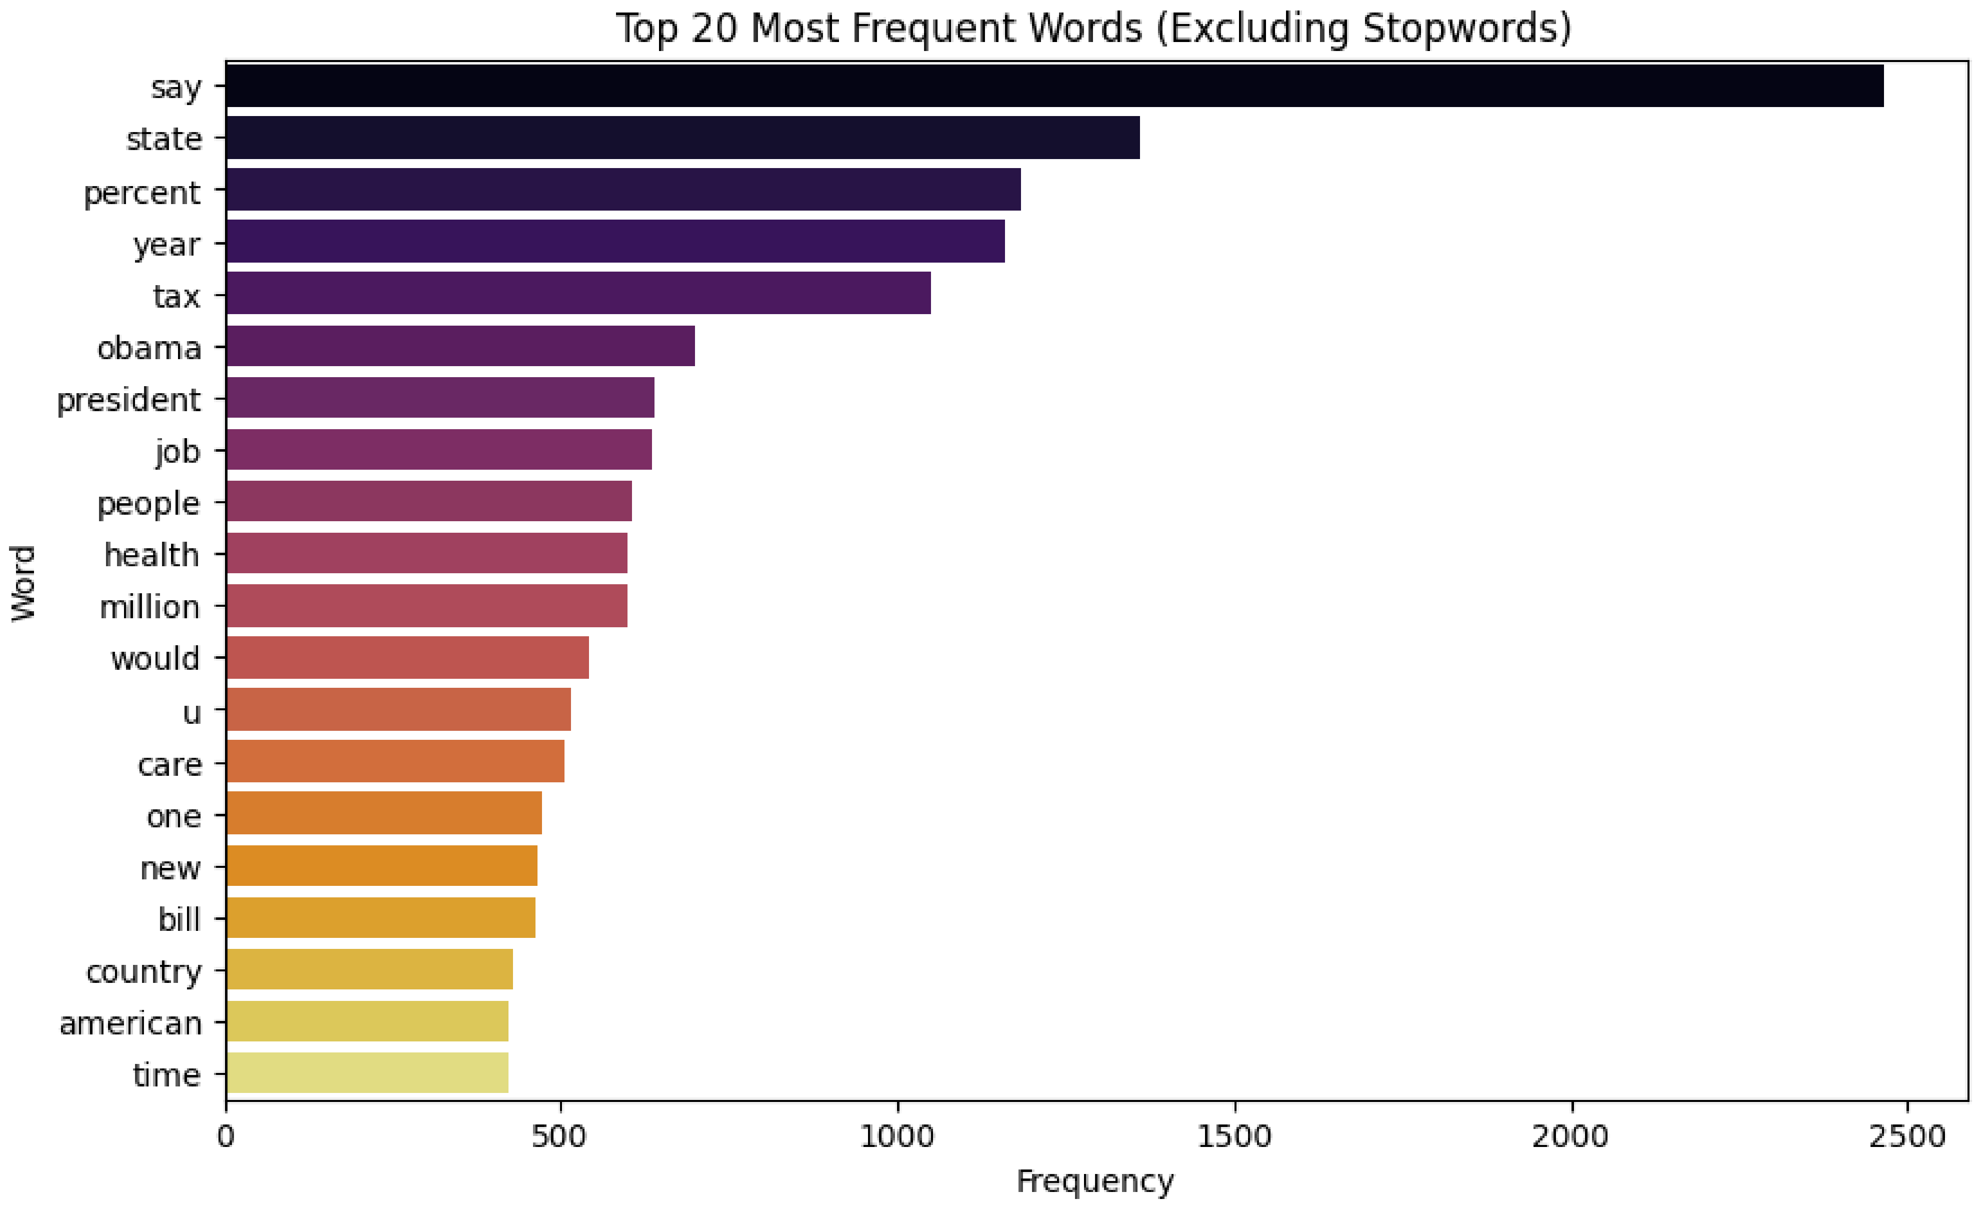

& a Little More Analysis...

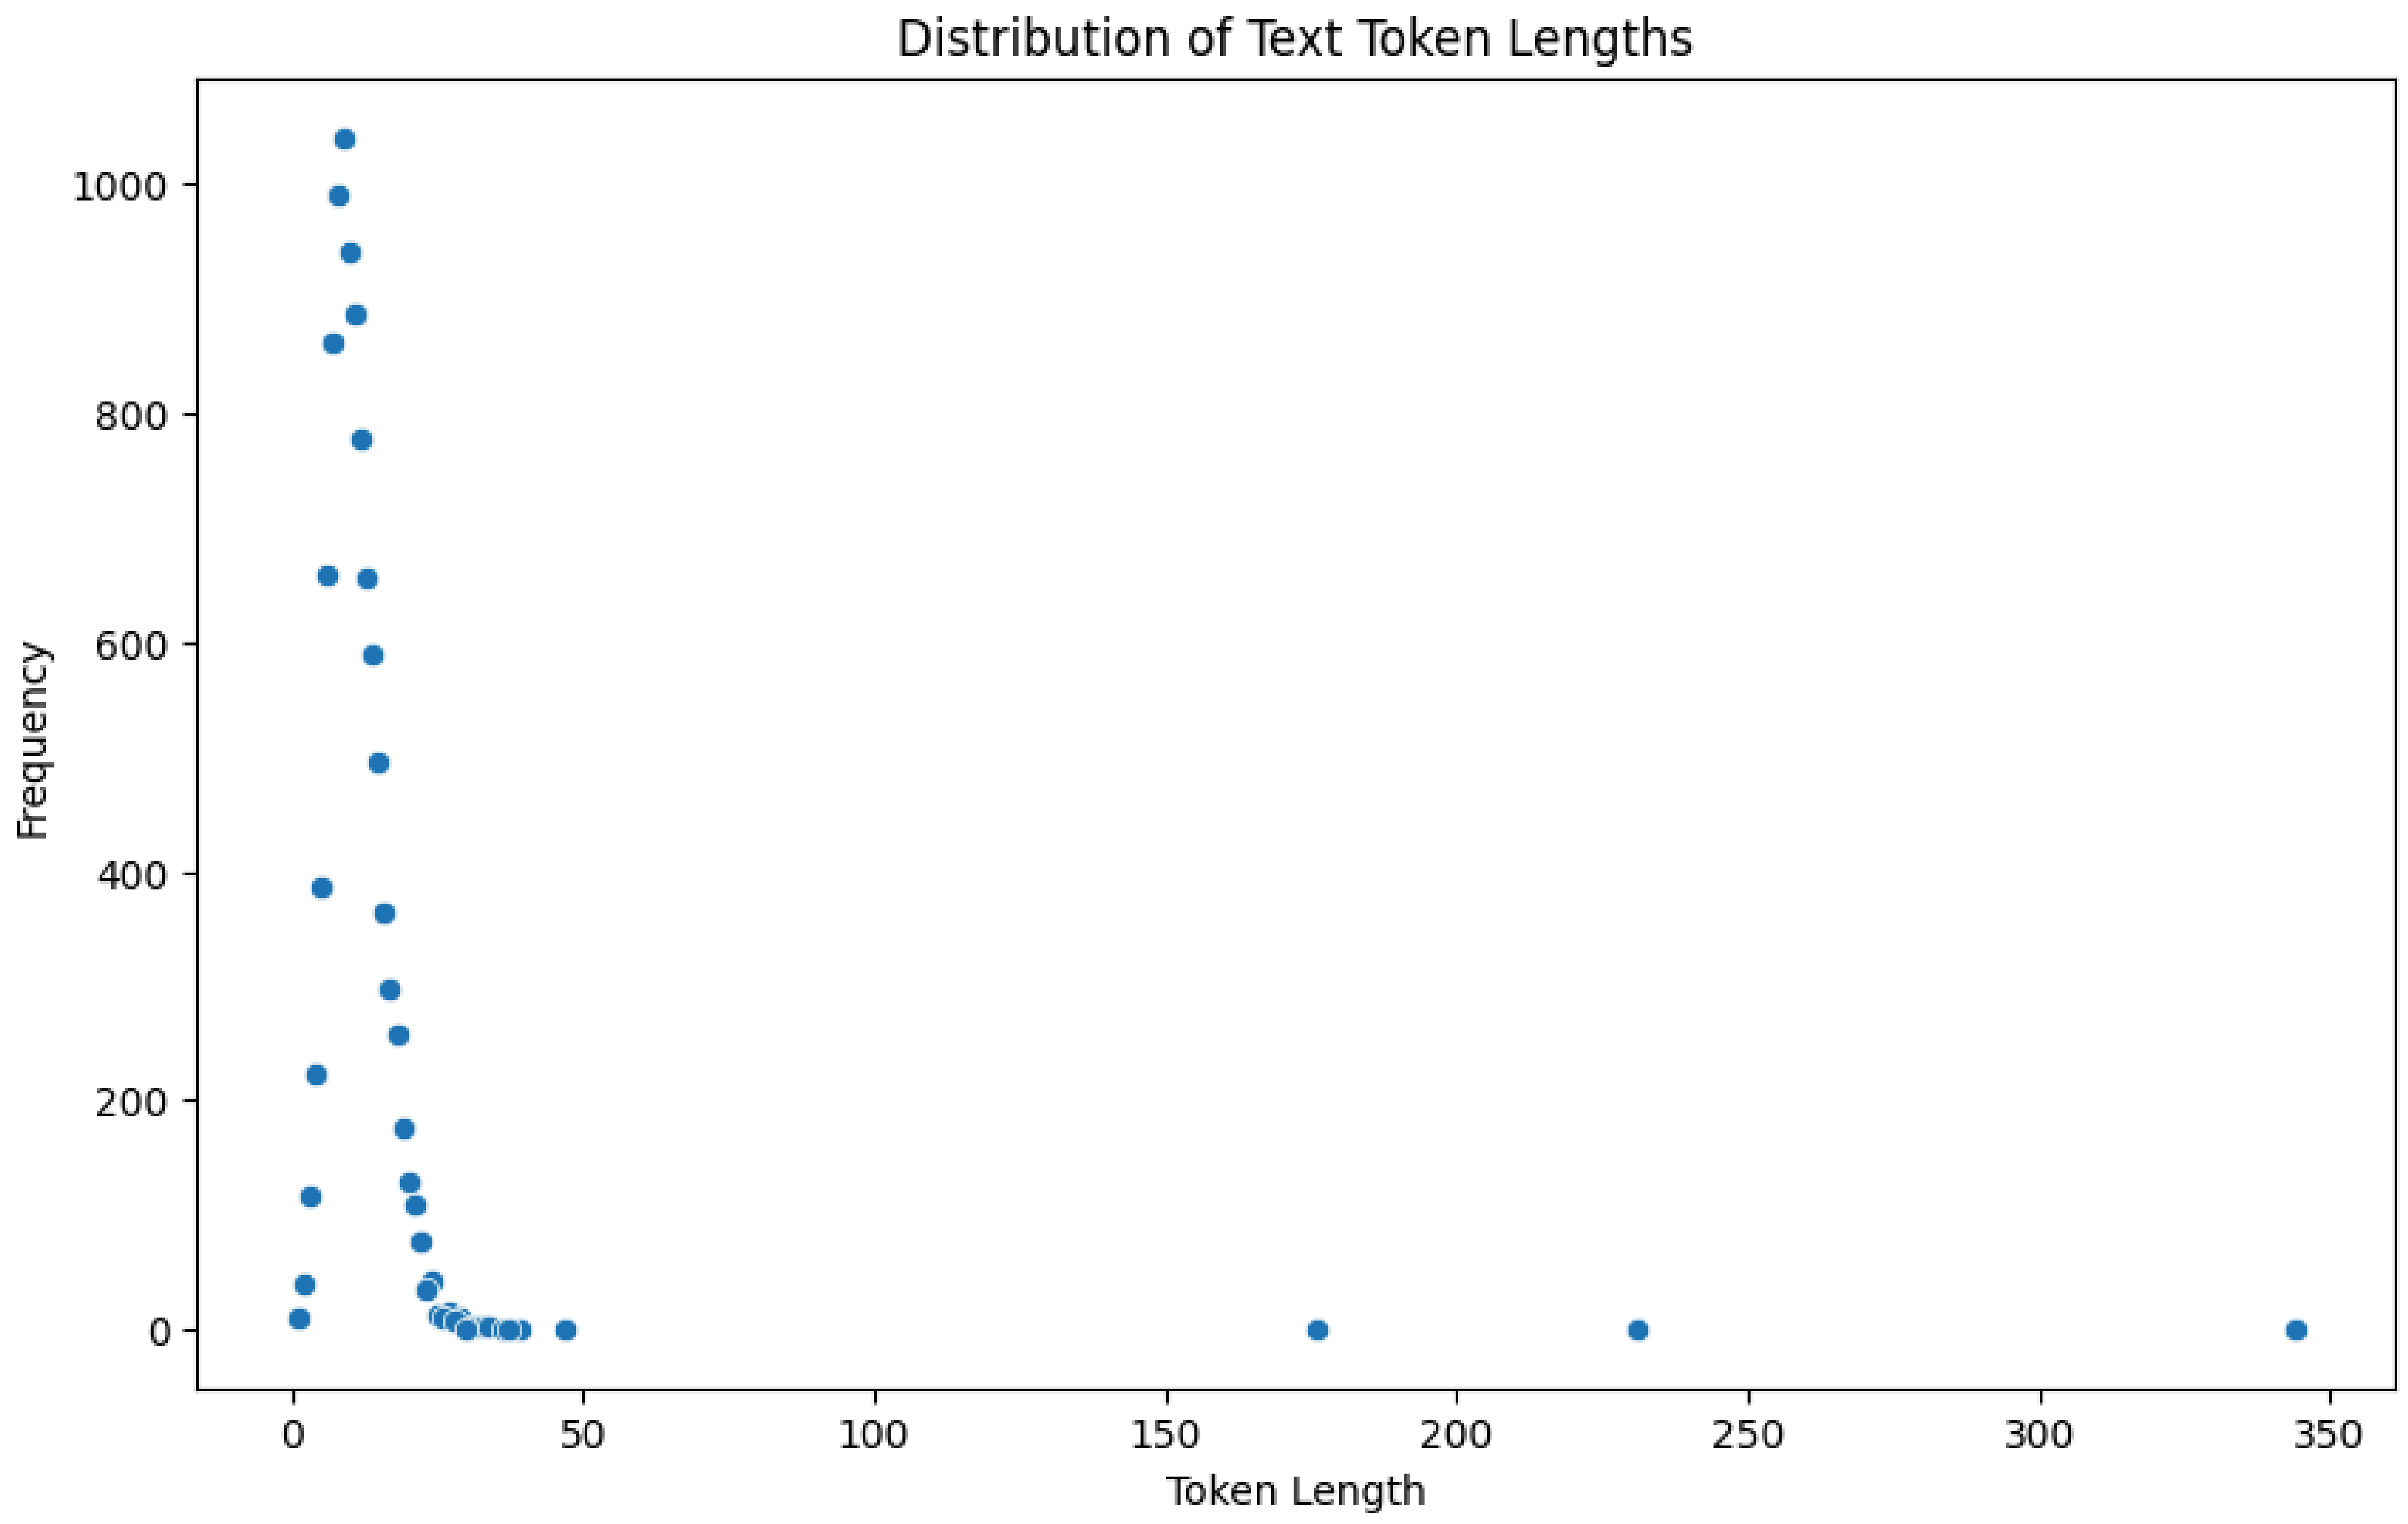

existing work -

Models	Valid.	Test
Majority	0.204	0.208
SVMs	0.258	0.255
Logistic Regress0ion	0.257	0.247
Bi-LSTMs	0.223	0.233
CNNs	0.260	0.270
Hybrid CNNs		
Text + Subject	0.263	0.235
Text + Speaker	**0.277**	0.248
Text + Job	0.270	0.258
Text + State	0.246	0.256
Text + Party	0.259	0.248
Text + Context	0.251	0.243
Text + History	0.246	0.241
Text + All	0.247	**0.274**

Table 2: The evaluation results on the LIAR dataset. The top section: text-only models. The bottom: text + meta-data hybrid models.

The dataset's authors had done six-class classification with various techniques. The image on the left captures the same.

Source:
https://paperswithcode.com/paper/liar-liar-pants-on-fire-a-new-benchmark

existing work -

Dataset	Modality	Size	Labels	Type	URL
Fake news	Text	20,800	Unreliable, reliable	News articles	https://www.kaggle.com/c/fake-news/data.
Weibo [27]	Text & image	40k tweets	Rumor, Non-rumor	Social media data	https://drive.google.com/file/d/14VQ7EWPiFeGzxp3XC2DeEHi-BEisDINn/view
Twitter15 [28]	Propagation trees	1,381 propagation trees, 276,663 users	Unverified, true, false, non-rumor	Social media data	https://www.dropbox.com/s/7ewzdrbelpmrnxu/rumdetect2017.zip?dl=0
Twitter16 [28]	Propagation trees	1,181 propagation trees, 173,487 users	Unverified, true, false, non-rumor	Social media data	https://www.dropbox.com/s/7ewzdrbelpmrnxu/rumdetect2017.zip?dl=0
LIAR [29]	Text	12.8K	Pants on fire, false, barely true, half-true, mostly true, and true	Political statements	https://paperswithcode.com/dataset/liar
PHEME [30]	Text	5800 tweets	Rumor, Non-rumor	Social media data	https://figshare.com/articles/dataset/PHEME_dataset_of_rumours_and_non-rumours/4010619
FNC-1	Text	75K	Agrees, disagrees, discusses, unrelated	News articles	https://github.com/FakeNewsChallenge/fnc-1
FakeNewsNet [31]	Text	5K	Fake, real	News articles, social media data	https://github.com/KaiDMML/FakeNewsNet
News Aggregator	Text	422,937	Real	News articles	https://www.kaggle.com/uciml/news-aggregator-dataset
Bend the truth [32]	Text	900	Fake, real	News articles	https://github.com/MaazAmjad/Datasets-for-Urdu-news.git
FacebookHoax [33]	Text	15,500	Hoax, non-hoax	scientific news	https://github.com/gabll/some-like-it-hoax/tree/master/dataset
Twitter [34]	Text and Image	992	Rumor, non-rumor	Fact-checked claims	https://github.com/MKLab-ITI/image-verification-corpus/tree/master/mediaeval2015
KaggleFN	Text	13K	Fake	News articles	https://www.kaggle.com/mrisdal/fake-news
FakevsSatire [35]	Text	486	Fake, satire	Political news	https://github.com/jgolbeck/fakenews

Comparing options in datasets. Ours had the most number of labels.

Source: A Comprehensive Review of Fake News Detection with Deep Learning. Mridha et al., 2021.

https://ieeexplore.ieee.org/abstract/document/9620068

existing work -

Method	Advantages	Disadvantages	References
TF-IDF	The TF-IDF model includes information on both the more significant and less important words.	Slow for large vocabularies. Does not capture position in text, semantics, co-occurrences in different documents, etc.	[67, 4, 68, 49, 46, 52, 63, 69, 17, 70, 71, 72]
Bag-of-words	The ease of implementation.	It ignores the ordering of the words in a given document. Ignores the semantic relations among words	[68, 46, 73, 74, 70, 75]
Word2Vec [76]	Maintains the semantic meaning of various words in a text. The context information is preserved. The size of the embedding vector is very small	Inability to deal with unfamiliar words. There are no common representations at the sub-word level.	[40, 42, 41, 74, 77, 46, 78, 79, 73, 60, 80, 55, 51, 81]
Doc2Vec [82]	A numeric representation of a document, regardless of its length. Faster than Word2vec.	The benefit of using doc2vec is diminished for shorter documents	[83, 84]
GloVe [85]	GloVe, unlike Word2vec, does not rely solely on local statistics (Words local context information)	In order to obtain word vectors, global statistics (word co-occurrence) are used.	[40, 50, 41, 86, 60, 51, 87, 80, 88]
BERT [89]	Identify and capture contextual meaning in a sentence or text	Compute-intensive at inference time	[79, 90, 75, 91, 92, 93, 81, 94, 95]

A table comparing word embedding methods. We eventually used MiniLM for its small size, and yet having the strengths of BERT.

Source: A Comprehensive Review of Fake News Detection with Deep Learning. Mridha et al., 2021.
https://ieeexplore.ieee.org/abstract/document/9620068

existing work -

In case of 6 classification, this method achieved an accuracy of 23.6% which is improved a lot in the next method.

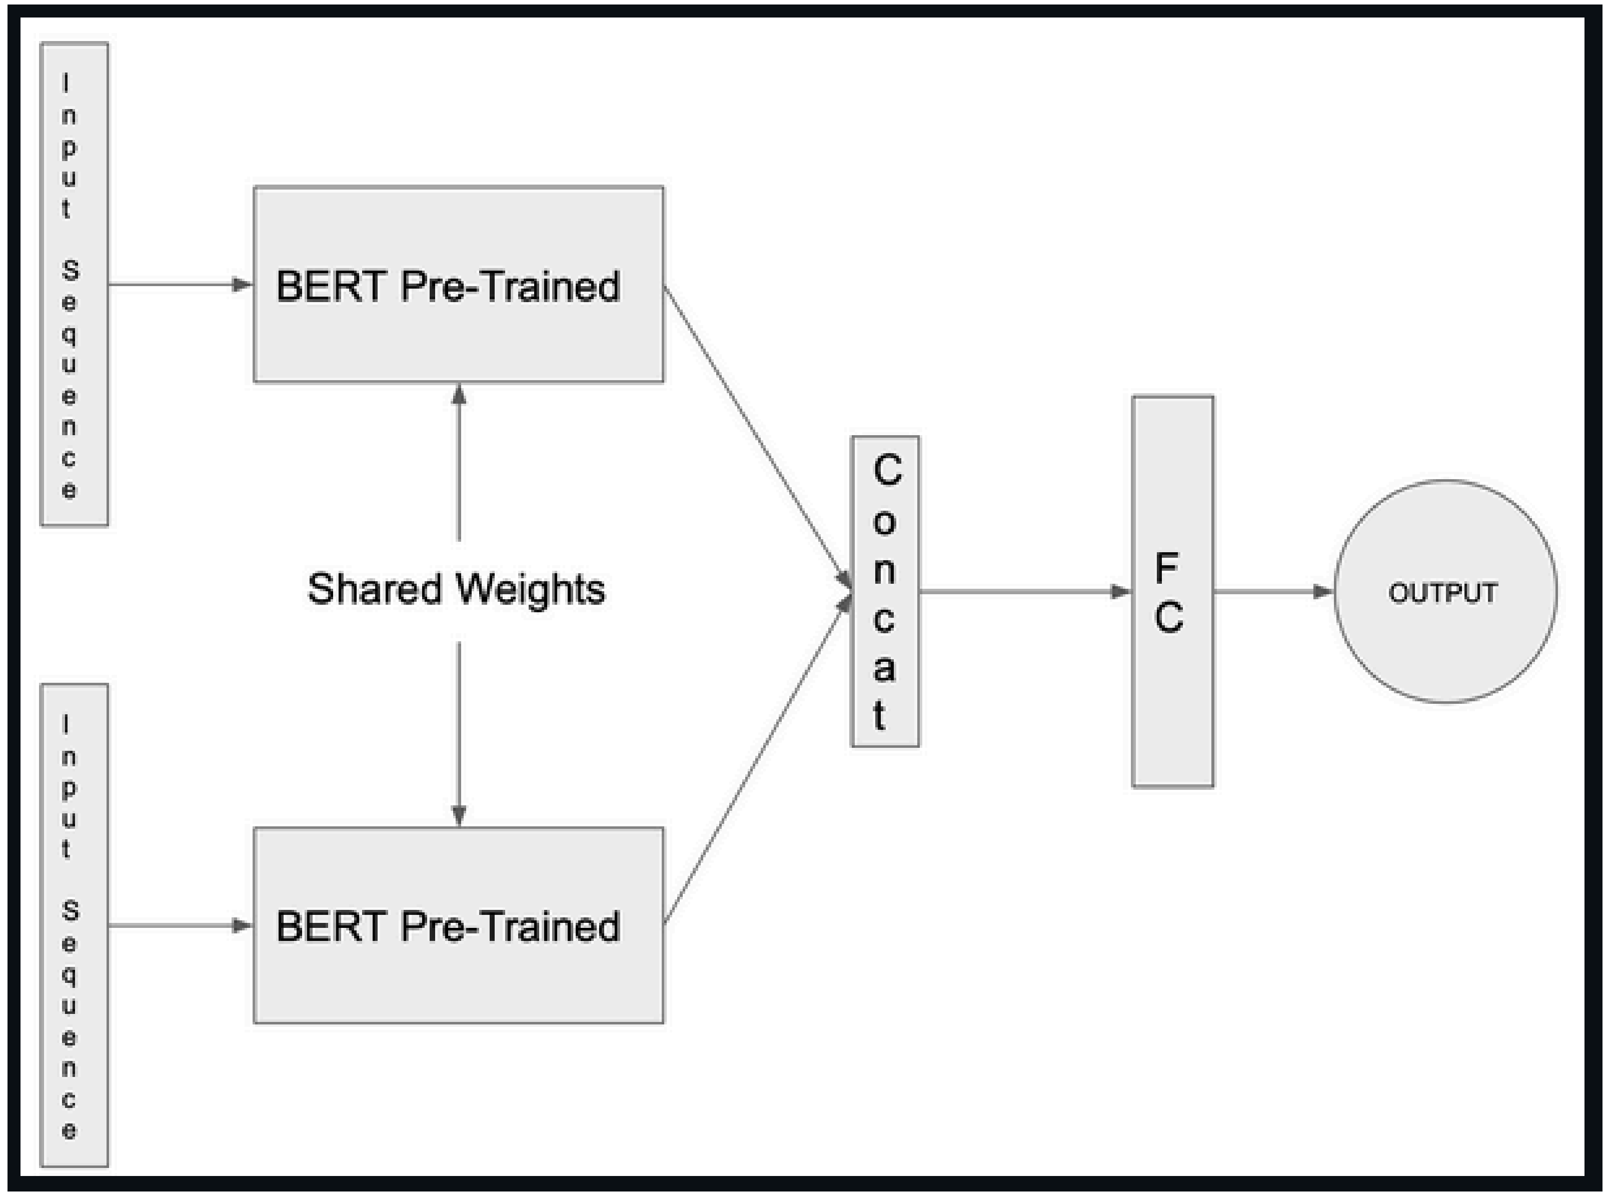

Quite different from the binary classification, there was an improvement in accuracy in the case of 6 class classification to 32.8%.

Out of 7 available classification implementations on PapersWithCode [1], the most upvoted repository [2] was the only one to provide solid numbers in its results. Images on the left are from its README.

The repository's code uses a Siamese network of BERT pre-trained models. Top uses two models, bottom uses three.

Sources:

[1] https://paperswithcode.com/paper/liar-liar-pants-on-fire-a-new-benchmark

[2] https://github.com/manideep2510/siamese-BERT-fake-news-detection-LIAR

support vector machines

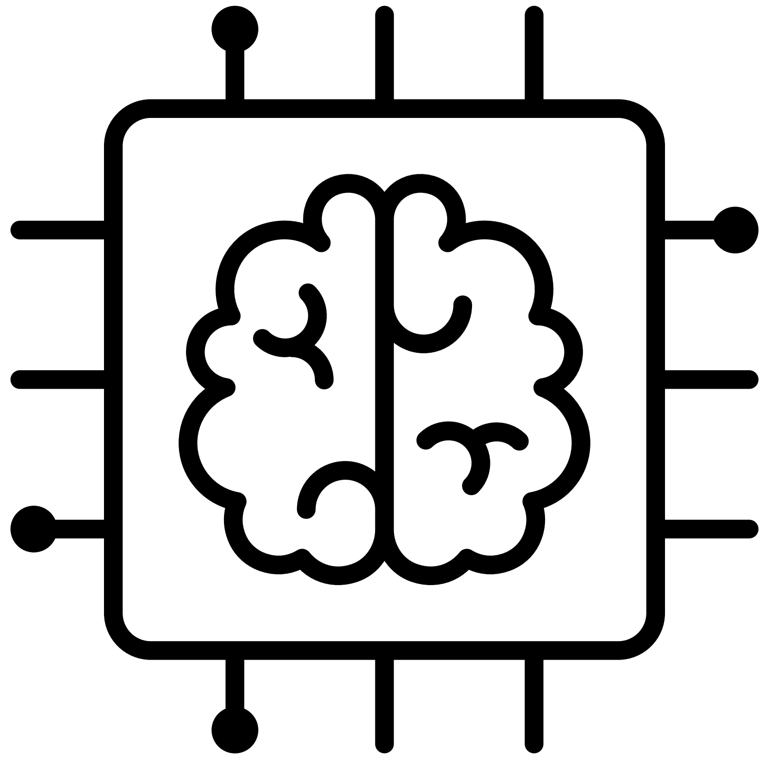

- Support Vector Machine (SVM) is a machine learning algorithm that finds the best hyperplane to separate different classes in a dataset, maximizing the margin between them.
- Our baseline model employs a Support Vector Machine, processing unigram features extracted through TF-IDF vectorization of the BERT Summaries.
- it served as an excellent base classical machine learning model used for multi -class classification.

Performance-

	precision	recall	f1-score	support
0	0.26	0.47	0.34	250
1	0.27	0.34	0.30	267
2	0.27	0.41	0.32	249
3	0.30	0.08	0.12	211
4	0.29	0.09	0.14	214
5	0.00	0.00	0.00	92
accuracy			0.27	1283
macro avg	0.23	0.23	0.20	1283
weighted avg	0.26	0.27	0.23	1283

Performance-

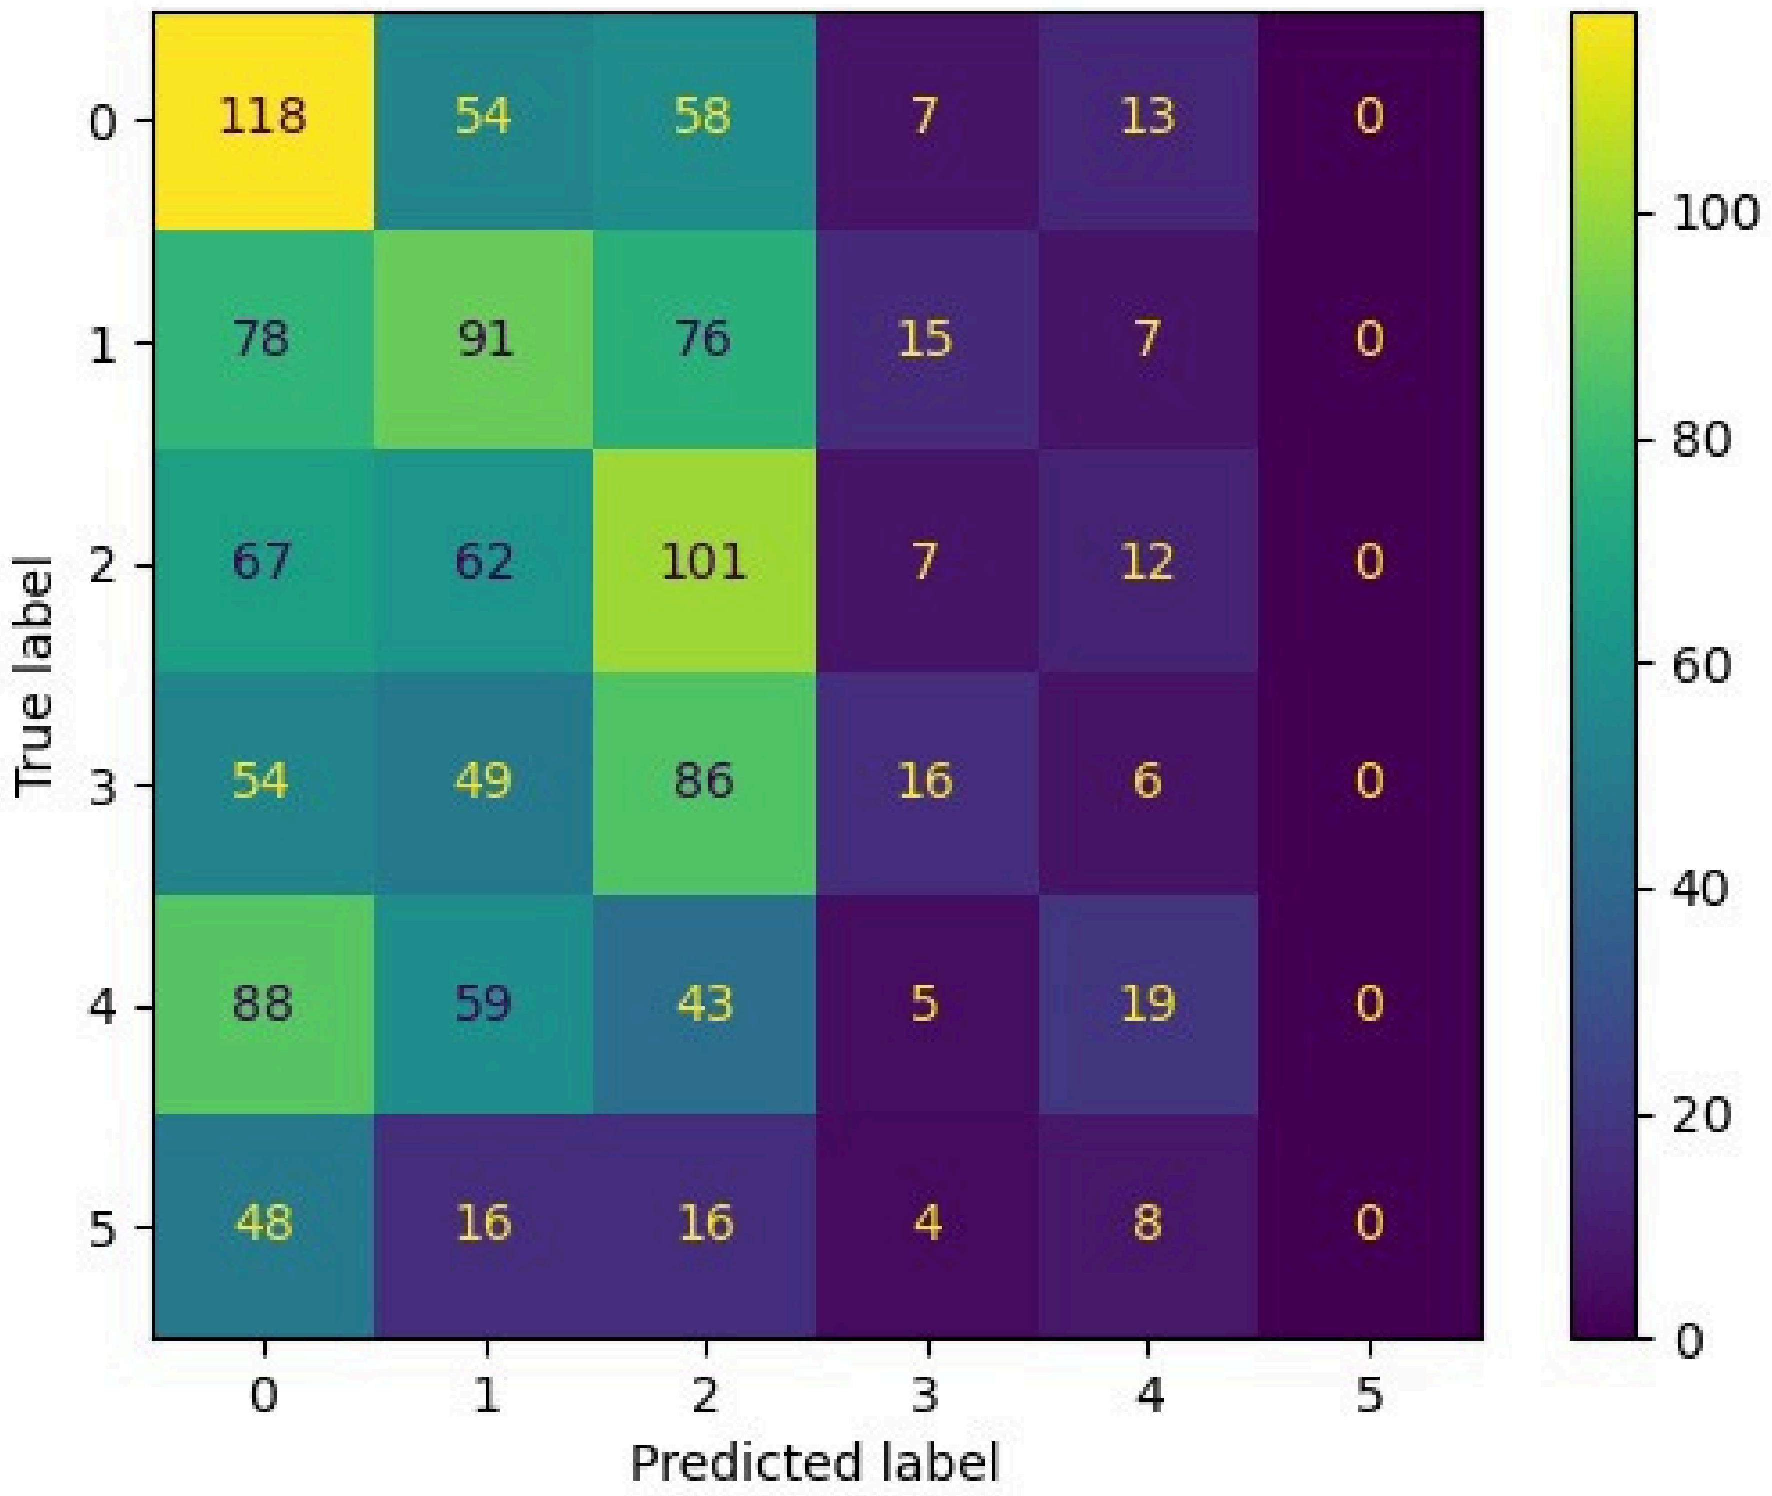

RNN

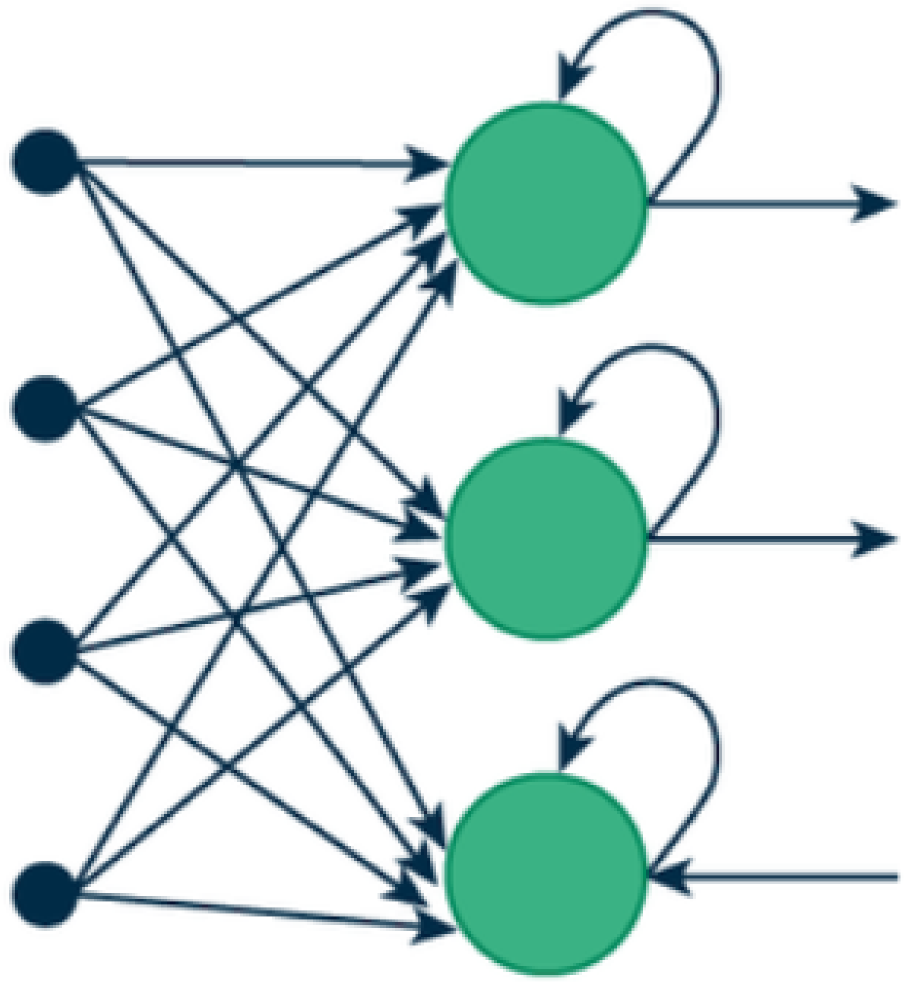

(a) Recurrent Neural Network

What is RNN?

RNN, or Recurrent Neural Network, is a type of artificial neural network designed to handle sequential data by retaining memory of past inputs. It's commonly used in tasks like natural language processing and time series prediction.

Why RNN?

RNNs are ideal for fake news detection because they can understand the sequential nature of text, capturing nuances and language patterns crucial for identifying deceptive content.

LSTM

- LSTM model understands and analyzes both short text and longer articles akin to human reading capabilities.
- It identifies key keywords, comprehends sentiment, and grasps contextual meaning.
- Trained on a dataset of labeled news articles (real vs. fake), learning to differentiate between truthful and deceptive content.
- Detects subtle cues like word choice, phrasing, and context that humans might miss.
- Can swiftly evaluate new articles and accurately classify them as real or fake news.
- Automated fake news detection aids in filtering out misinformation, offering users fact-based content.

Performance- SVM & RNN + LSTM

SVM

```
Validation Accuracy for Support Vector Machine: 0.26246105919003115
```

RNN + LSTM

```
161/161 ━━━━━━━━━━━━━━━━━━━━ 10s 41ms/step - accuracy: 0.9930 - loss: 0.0210 - val_accuracy: 0.2274 - val_loss: 7.8848
41/41 ━━━━━━━━━━━━━━━━━━━━ 1s 14ms/step
Test Accuracy for RNN + LSTM: 0.2244738893219018
```

The Final Architecture- Bi-LSTM

```
41/41 ━━━━━━━━━━━━━━━━━━━━ 1s 26ms/step
              precision    recall  f1-score   support

           0       0.26      0.41      0.32       250
           1       0.25      0.41      0.31       267
           2       0.24      0.27      0.26       249
           3       0.30      0.12      0.17       211
           4       0.26      0.09      0.13       214
           5       0.32      0.09      0.14        92

    accuracy                           0.26      1283
   macro avg       0.27      0.23      0.22      1283
weighted avg       0.27      0.26      0.24      1283
```

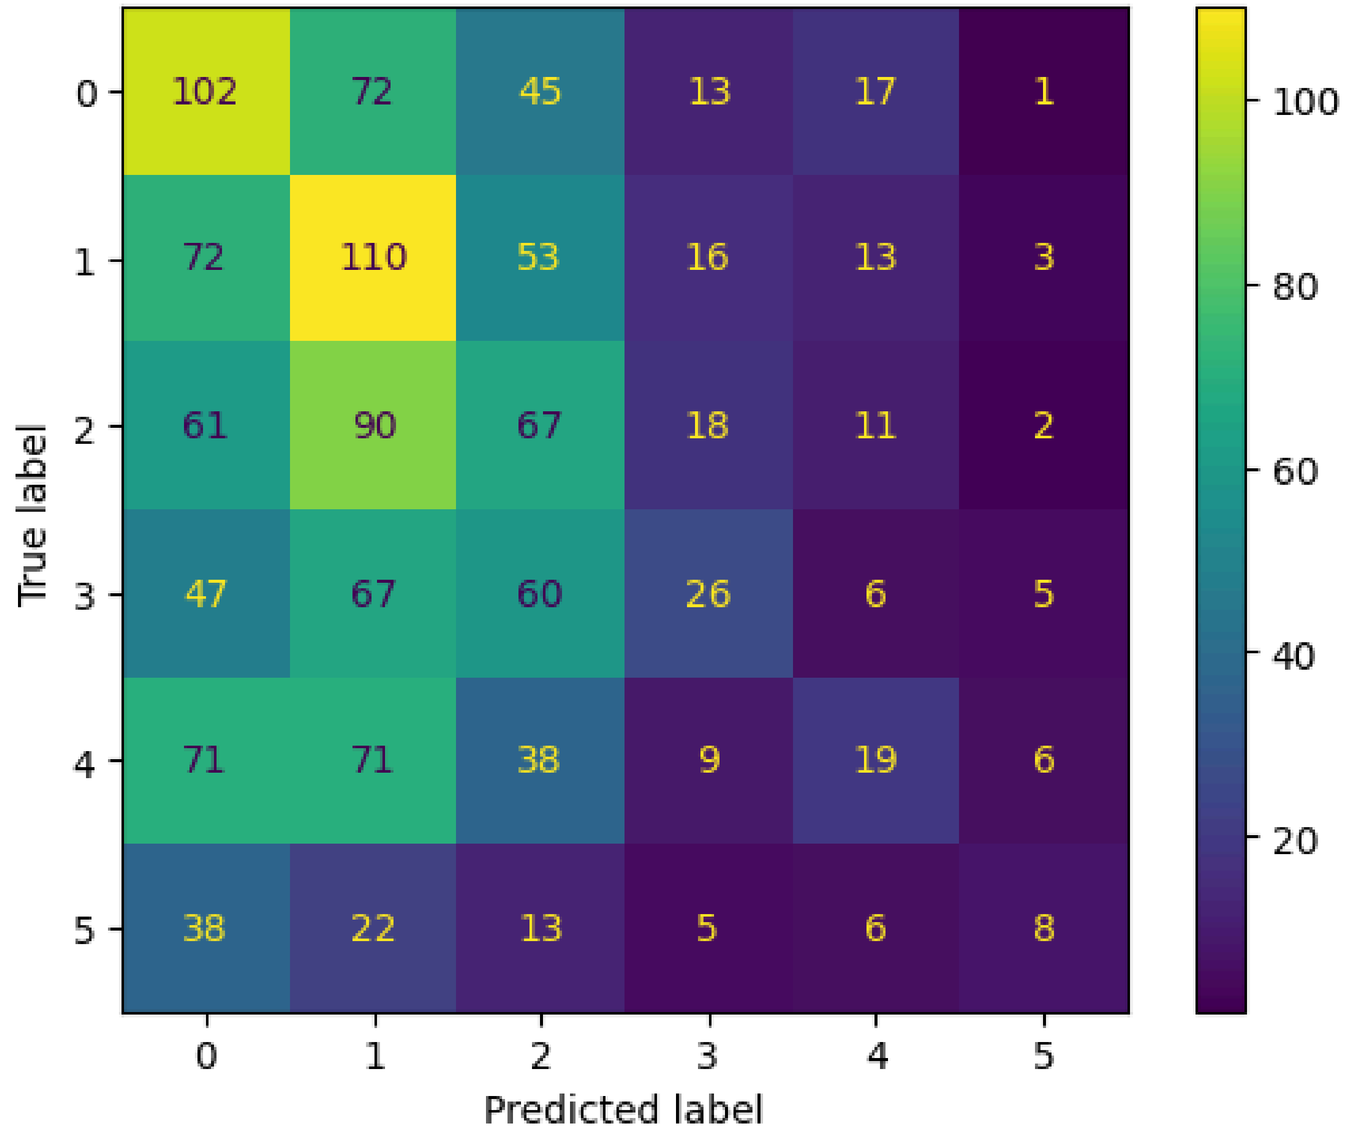


```
accuracy: 0.2896 - loss: 1.6991
```

challenges:

1. **Overfitting:** While our training accuracy at times was reaching 99%, our validation and testing accuracy kept wandering around 25-27%. This is a known problem in the literature, and is being worked on by the community to solve.
2. **Resource Constraints:** Training and deploying Bi-LSTM models require significant computational resources. Managing these resources efficiently, especially for larger datasets, can be a big stone to move.
3. **Learning:** We still require to innovate on better architectures, for example possibly implementing attention mechanisms in the future.

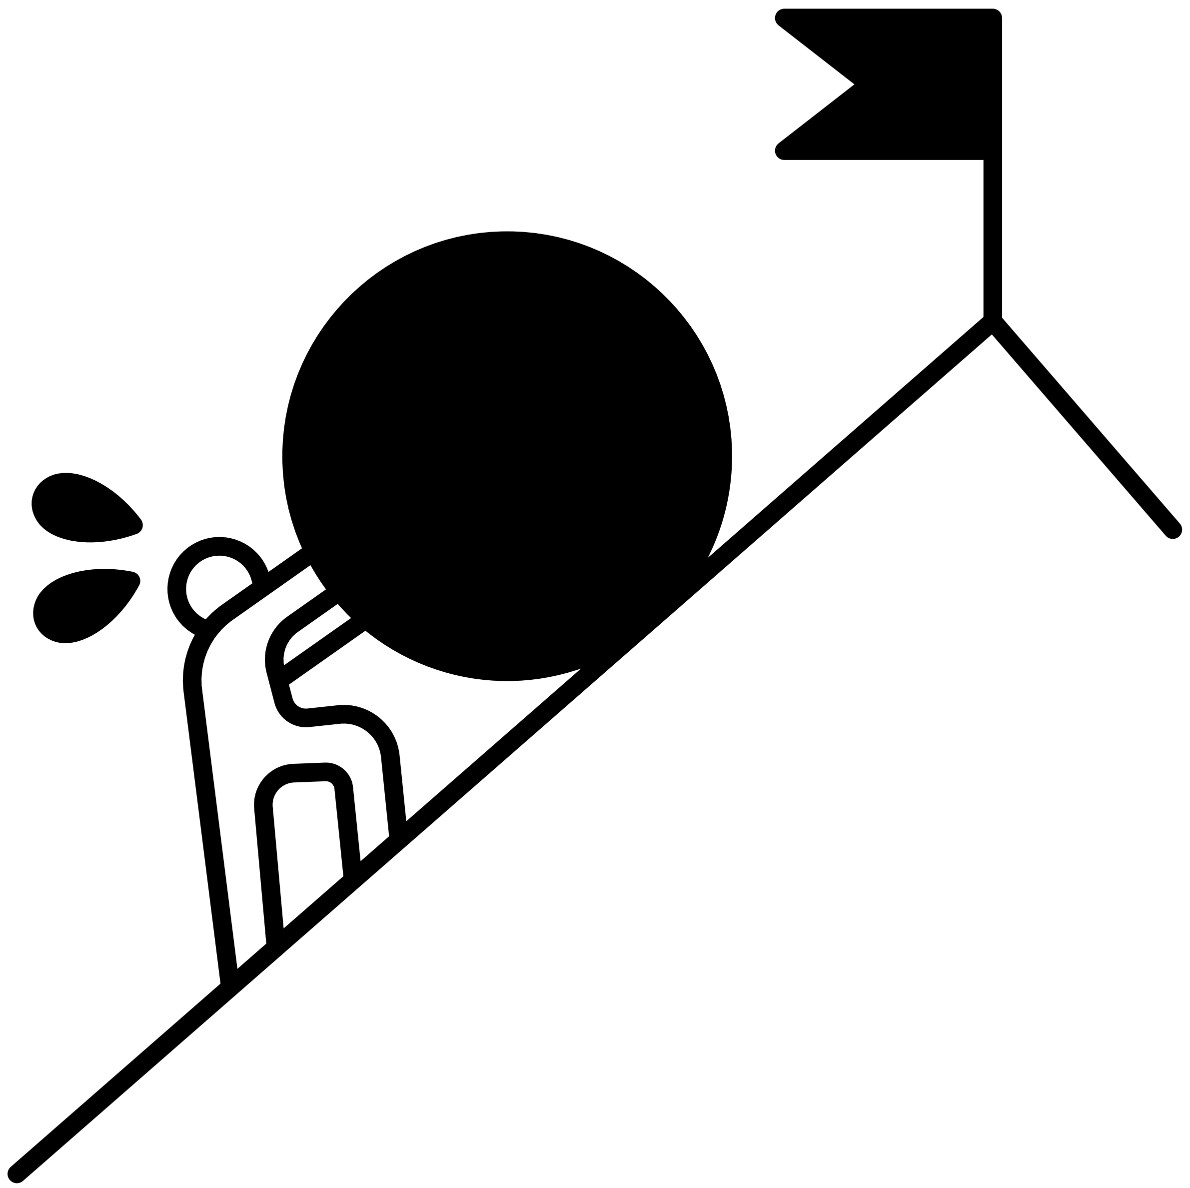

Future:

1. **Multi-Source Transfer Learning-** An avenue we had attempted to look at that would allow us to use multiple pre-trained models to get more accurate outputs. We could follow through on that to get better results.
2. **Attention Mechanisms-** Another thing we tried implementing to no avail, could be a big avenue for innovation.
3. **Integrating More Features**- For now we used only the statement to make predictions, we wish to understand how to represent and use the other features effectively in the future.
4. **More Diverse Dataset**- The dataset covers only American statements, we want to explore how to get around this.

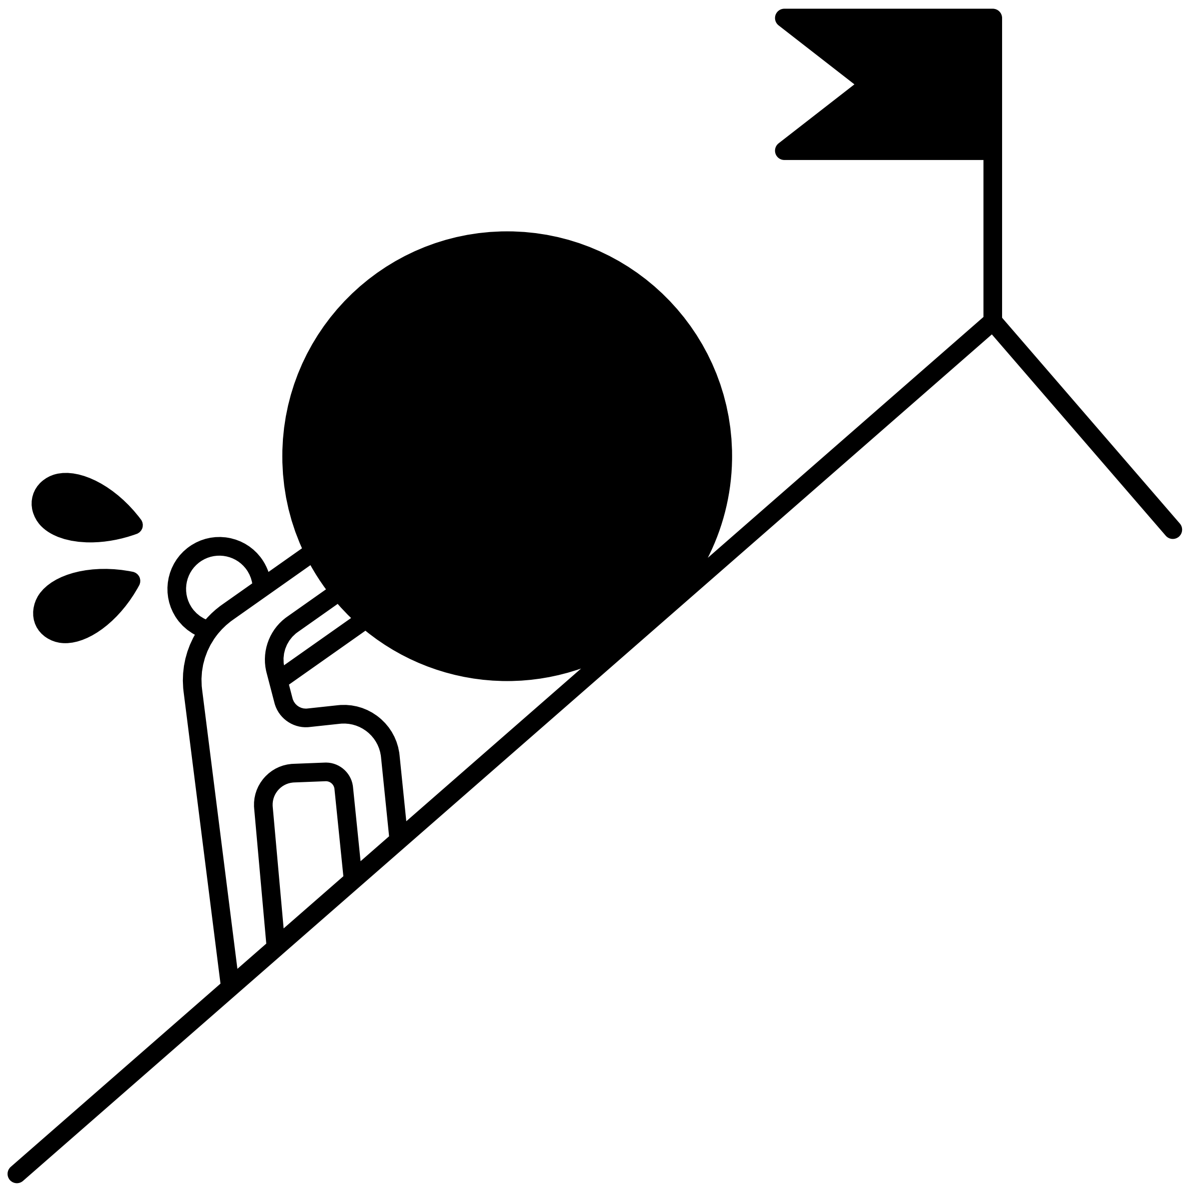

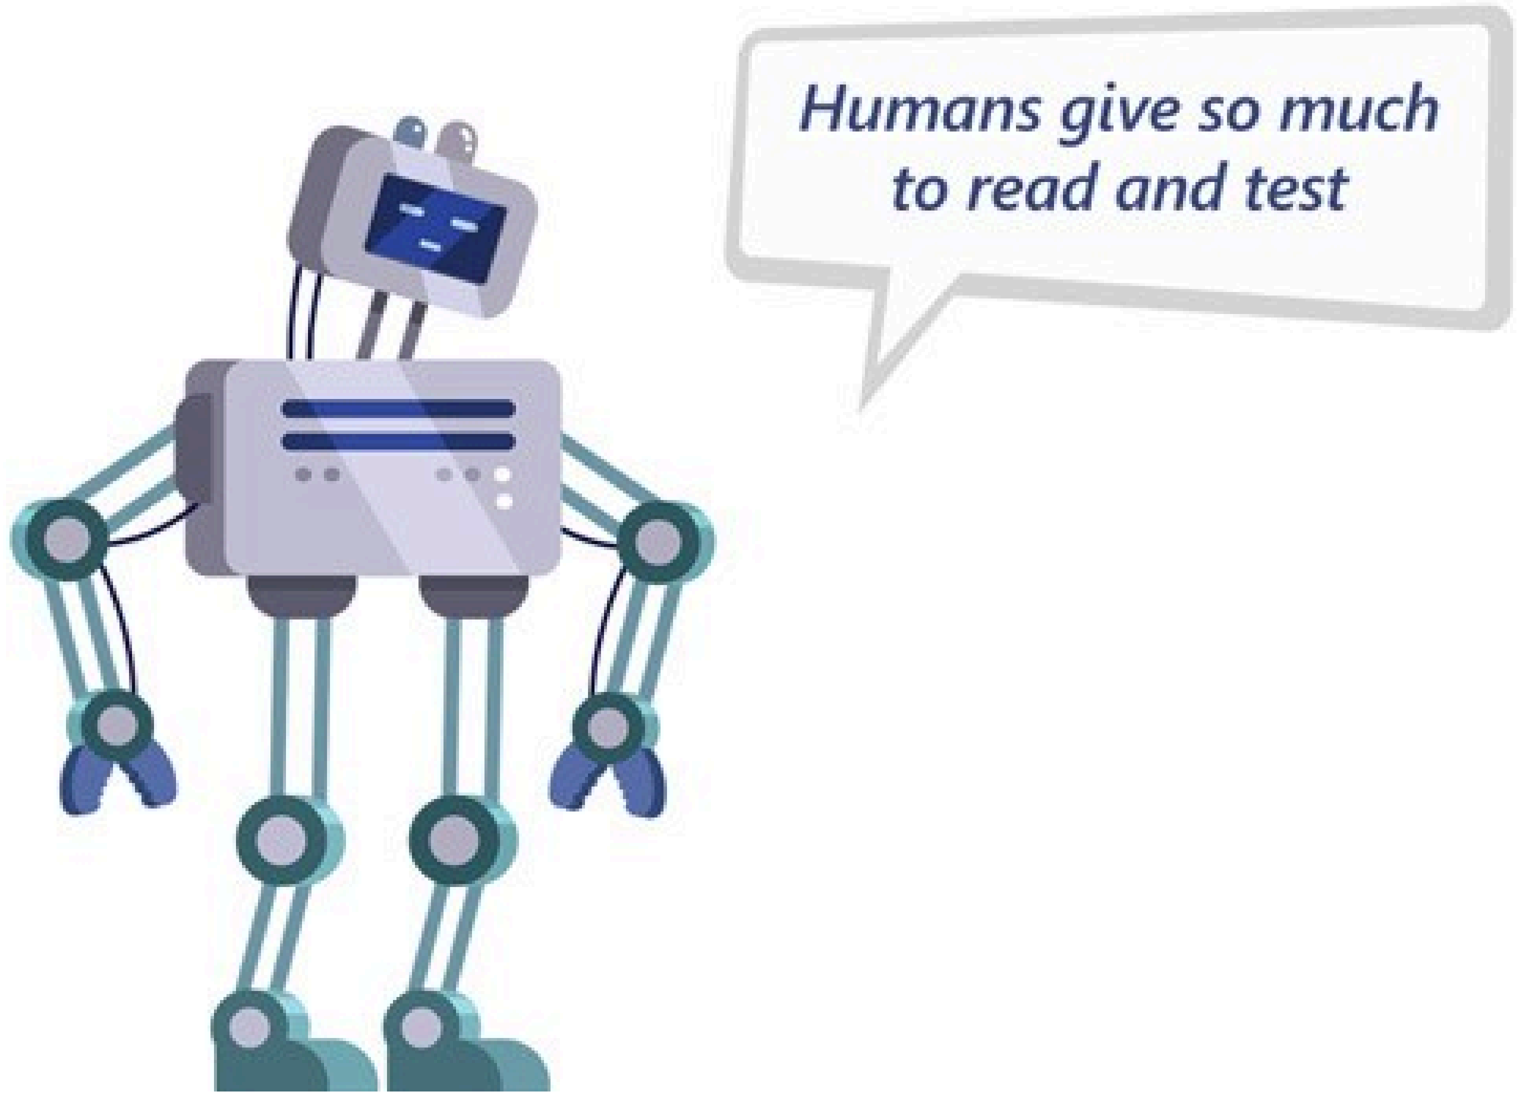

anyways...