



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/175036/>

Version: Published Version

Proceedings Paper:

Maronikolakis, A., Schütze, H. and Stevenson, R. (2021) Identifying automatically generated headlines using transformers. In: Feldman, A., Da San Martino, G., Leberknight, C. and Nakov, P., (eds.) Proceedings of the Fourth Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda. 4th Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda, 06 Jun 2021, Virtual conference. Association for Computational Linguistics (ACL). ISBN: 9781954085268.

<https://doi.org/10.18653/v1/2021.nlp4if-1.1>

© 2021 The Authors. Article available under the terms of the CC-BY licence (<https://creativecommons.org/licenses/by/4.0/>).

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Identifying Automatically Generated Headlines using Transformers

Antonis Maronikolakis
CIS, LMU Munich
antmarakis@cis.lmu.de

Hinrich Schutze
CIS, LMU Munich

Mark Stevenson
University of Sheffield
mark.stevenson@sheffield.ac.uk

Abstract

False information spread via the internet and social media influences public opinion and user activity, while generative models enable fake content to be generated faster and more cheaply than had previously been possible. In the not so distant future, identifying fake content generated by deep learning models will play a key role in protecting users from misinformation. To this end, a dataset containing human and computer-generated headlines was created and a user study indicated that humans were only able to identify the fake headlines in 47.8% of the cases. However, the most accurate automatic approach, transformers, achieved an overall accuracy of 85.7%, indicating that content generated from language models can be filtered out accurately.

1 Introduction

Fake content has been rapidly spreading across the internet and social media, misinforming and affecting users' opinion (Kumar and Shah, 2018; Guo et al., 2020). Such content includes fake news articles¹ and truth obfuscation campaigns². While much of this content is being written by paid writers (Luca and Zervas, 2013), content generated by automated systems is rising. Models can produce text on a far greater scale than it is possible to manually, with a corresponding increase in the potential to influence public opinion. There is therefore a need for methods that can distinguish between human and computer-generated text, to filter out deceiving content before it reaches a wider audience.

While text generation models have received consistent attention from the public as well as from the academic community (Dathathri et al., 2020; Subramanian et al., 2018), interest in the detection of automatically generated text has only arisen more

recently (Jawahar et al., 2020). Generative models have several shortcomings and their output text has characteristics that distinguish it from human-written text, including lower variance and smaller vocabulary (Holtzman et al. (2020); Gehrmann et al. (2019)). These differences between real and generated text can be used by pattern recognition models to differentiate between the two. In this paper we test this hypothesis by training classifiers to detect headlines generated by a pretrained GPT-2 model (Radford et al., 2019). Headlines were chosen as it has been shown that shorter generated text is harder to identify than longer content (Ippolito et al., 2020).

The work described in this paper is split into two parts: the creation of a dataset containing headlines written by both humans and machines and training of classifiers to distinguish between them. The dataset is created using real headlines from the *Million Headlines* corpus³ and headlines generated by a pretrained GPT-2. The training and development sets consist of headlines from 2015 while the testing set consists of 2016 and 2017 headlines. A series of baselines and deep learning models were tested. Neural methods were found to outperform humans, with transformers being almost 35% more accurate.

Our research highlights how difficult it is for humans to identify computer-generated content, but that the problem can ultimately be tackled using automated approaches. This suggests that automatic methods for content analysis could play an important role in supporting readers to understand the veracity of content. The main contributions of this work are the development of a novel fake content identification task based on news headlines⁴ and analysis of human evaluation and machine learning approaches to the problem.

¹For example, [How a misleading post went from the fringes to Trump's Twitter](#).

²For example, [Can fact-checkers save Taiwan from a flood of Chinese fake news?](#)

³Accessed 25/01/2021.

⁴Code available at http://bit.ly/ant_headlines.

2 Relevant Work

Kumar and Shah (2018) compiled a survey on fake content on the internet, providing an overview of how false information targets users and how automatic detection models operate. The sharing of false information is boosted by the natural susceptibility of humans to believe such information. Pérez-Rosas et al. (2018) and Ott et al. (2011) reported that humans are able to identify fake content with an accuracy between 50% and 75%. Information that is well presented, using long text with limited errors, was shown to deceive the majority of readers. The ability of humans to detect machine-generated text was evaluated by Dugan et al. (2020), showing that humans struggle at the task.

Holtzman et al. (2020) investigated the pitfalls of automatic text generation, showing that sampling methods such as Beam search can lead to low quality and repetitive text. Gehrmann et al. (2019) showed that automatic text generation models use a more limited vocabulary than humans, tending to avoid low-probability words more often. Consequently, text written by humans tends to exhibit more variation than that generated by models.

In Zellers et al. (2019), neural fake news detection and generation are jointly examined in an adversarial setting. Their model, called Grover, achieves an accuracy of 92% when identifying real from generated news articles. Human evaluation though is lacking, so the potential of Grover to fool human readers has not been thoroughly explored. In Brown et al. (2020), news articles generated by their largest model (175B parameters) managed to fool humans 48% of the time. The model, though, is prohibitively large to be applied at scale. Further, Ippolito et al. (2020) showed that shorter text is harder to detect, both for humans and machines. So even though news headlines are a very potent weapon in the hands of fake news spreaders, it has not been yet examined how difficult it is for humans and models to detect machine-generated headlines.

3 Dataset

3.1 Dataset Development

The dataset was created using Australian Broadcasting Corporation headlines and headlines generated from a model. A pretrained⁵ GPT-2 model (Radford et al., 2019) was finetuned on the headlines data. Text was generated using sampling with tem-

⁵As found in the HuggingFace library.

perature and continuously re-feeding words into the model until the end token is generated.

Data was split in two sets, 2015 and 2016/2017, denoting the sets a “defender” and an “attacker” would use. The goal of the attacker is to fool readers, whereas the defender wants to filter out the generated headlines of the attacker. Headlines were generated separately for each set and then merged with the corresponding real headlines.

The “defender” set contains 72,401 real and 414,373 generated headlines, while the “attacker” set contains 179,880 real and 517,932 generated.

3.2 Dataset Analysis

Comparison of the real and automatically generated headlines revealed broad similarities between the distribution of lexical terms, sentence length and POS tag distribution, as shown below. This indicates that the language models are indeed able to capture patterns in the original data.

Even though the number of words in the generated headlines is bound by the maximum number of words learned in the corresponding language model, the distribution of words is similar across real and generated headlines. In Figures 1 and 2 we indicatively show the 15 most frequent words in the real and generated headlines respectively. POS tag frequencies are shown in Table 1 for the top tags in each set. In real headlines, nouns are used more often, whereas in generated headlines the distribution is smoother, consistent with findings in Gehrmann et al. (2019). Furthermore, in generated headlines verbs appear more often in their base (VB) and third-person singular (VBZ) form while in real headlines verb tags are more uniformly distributed. Overall, GPT-2 has accurately learned the real distribution, with similarities across the board.




Figure 1: Top 15 Words for real headlines

Lastly, the real headlines are shorter than the generated ones, with 6.9 and 7.2 words respectively.




Figure 2: Top 15 Words for generated headlines

Real		Generated	
POS	freq	POS	freq
NN	0.372	NN	0.352
NNS	0.129	NNS	0.115
JJ	0.109	JJ	0.113
IN	0.108	IN	0.113
VB	0.045	VB	0.061
TO	0.040	TO	0.056
VBZ	0.033	VBZ	0.047
VBP	0.031	VBP	0.022
VBN	0.020	RB	0.017
VBG	0.020	VBG	0.015

Table 1: Frequencies for the top 10 part-of-speech tags in real and generated headlines

3.3 Survey

A crowd-sourced survey⁶ was conducted to determine how realistic the generated text is. Participants (n=124) were presented with 93 headlines (three sets of 31) in a random order and asked to judge whether they were real or generated. The headlines were chosen at random from the “attacker” (2016/2017) headlines.

In total, there were 3435 answers to the ‘real or generated’ questions and 1731 (50.4%) were correct. When presented with a computer-generated headline, participants answered correctly in 1113 out of 2329 (47.8%) times. In total 45 generated headlines were presented and out of those, 23 were identified as computer-generated (based on average response). This is an indication that GPT-2 can indeed generate realistic-looking headlines that fool readers. When presented with actual headlines, participants answered correctly in 618 out of 1106 times (55.9%). In total 30 real headlines were presented and out of those, 20 were correctly identified as real (based on average response).

Of the 45 generated headlines, five were marked as real by over 80% of the participants, while for

⁶Participants were students and staff members in a mailing list from the University of Sheffield.

the real headlines, 2 out of 30 reached that threshold. The five generated headlines were:

Rumsfeld Talks Up Anti Terrorism Campaign
 Cooper Rebounds From Olympic Disappointment
 Jennifer Aniston Tops Celebrity Power Poll
 Extra Surveillance Announced For WA Coast
 Police Crack Down On Driving Offences

At the other end of the spectrum, there were seven generated headlines that over 80% of the participants correctly identified as being computer-generated:

Violence Restricting Rescue Of Australian
 Scientists Discover Gene That May Halt Ovarian
 All Ordinaries Finishes Day On Closing High
 Waratahs Starting Spot Not A Mere Formality Sailor
 Proposed Subdivision Wont Affect Recreational
 Bangladesh To Play Three Tests Five Odis In
 Minister Promises More Resources To Combat Child

Most of these examples contain grammatical errors, such as ending with an adjective, while some headlines contain absurd or nonsensical content. These deficiencies set these headlines apart from the rest. It is worth noting that participants appeared more likely to identify headlines containing grammatical errors as computer-generated than other types of errors.

4 Classification

For our classifier experiments, we used the three sets of data (2015, 2016 and 2017) we had previously compiled. Specifically, for training we only used the 2015 set, while the 2016 and 2017 sets were used for testing. Splitting the train and test data by the year of publication ensures that there is no overlap between the sets and there is some variability between the content of the headlines (for example, different topics/authors). Therefore, we can be confident that the classifiers generalize to unknown examples.

Furthermore, for hyperparameter tuning, the 2015 data was randomly split into training and development sets on a 80/20 ratio. In total, for training there are 129,610 headlines, for development there are 32,402 and for testing there are 303,965.

4.1 Experiments

Four types of classifiers were explored: baselines (Elastic Net and Naive Bayes), deep learning (CNN, Bi-LSTM and Bi-LSTM with Attention), transfer

Method	Ovr. Acc.	Precision	Recall
Human	50.4	66.3	52.2
Naive Bayes	50.6	58.5	56.9
Elastic Net	73.3	58.1	62.3
CNN	81.7	75.3	76.2
BiLSTM	82.8	77.9	77.3
BiLSTM/Att.	82.5	76.9	77.2
ULMFit	83.3	79.1	78.5
BERT	85.7	86.9	81.2
DistilBERT	85.5	86.8	81.0

Table 2: Each run was executed three times with (macro) results averaged. Standard deviations are omitted for brevity and clarity (they were in all cases less than 0.5).

learning via ULMFit (Howard and Ruder, 2018) and Transformers (BERT (Devlin et al., 2019) and DistilBERT (Sanh et al., 2019)). The architecture and training details can be found in Appendix A.

Results are shown in Table 2. Overall accuracy is the accuracy in percentage over all headlines (real and generated), while (macro) precision and recall are calculated over the generated headlines. Precision is the percentage of correct classifications out of all the generated classifications, while recall is the percentage of generated headlines the model classified correctly out of all the actual generated headlines. High recall scores indicate that the models are able to identify a generated headline with high accuracy, while low precision scores show that models classify headlines mostly as generated.

We can observe from the results table that humans are overall less effective than all the examined models, including the baselines, scoring the lowest accuracy. They are also the least accurate on generated headlines, achieving the lowest recall. In general, human predictions are almost as bad as random guesses.

Deep learning models scored consistently higher than the baselines, while transfer learning outperformed all previous models, reaching an overall accuracy of around 83%. Transformer architectures though perform the best overall, with accuracy in the 85% region. BERT, the highest-scoring model, scores around 30% higher than humans in all metrics. The difference between the two BERT-based models is minimal.

Since training and testing data are separate (sampled from different years), this indicates that there are some traits in generated text that are not present in human text. Transformers are able to pick up on these traits to make highly-accurate classifications.

For example, generated text shows lower variance than human text (Gehrmann et al., 2019), which means text without rarer words is more likely to be generated than being written by a human.

4.2 Error Analysis

We present the following two computer-generated headlines as indicative examples of those misclassified as real by BERT:

Extra Surveillance Announced For WA Coast
Violence Restricting Rescue Of Australian

The first headline is not only grammatically sound, but also semantically plausible. A specific region is also mentioned (“WA Coast”), which has low probability of occurring and possibly the model does not have representative embeddings for. This seems to be the case in general, with the mention of named entities increasing the chance of fooling the classifier. The task of predicting this headline is then quite challenging. Human evaluation was also low here, with only 19% of participants correctly identifying it.

In the second headline, the word “restricting” and the phrase “rescue of” are connected by their appearance in similar contexts. Furthermore, both “violence” and “restricting rescue” have negative connotations, so they also match in sentiment. These two facts seem to lead the model in believing the headline is real instead of computer-generated, even though it is quite flimsy both semantically (the mention of violence is too general and is not grounded) and pragmatically (some sort of violence restricting rescue is rare). In contrast, humans had little trouble recognising this as a computer-generated headline; 81% of participants labelled it as fake. This indicates that automated classifiers are still susceptible to reasoning fallacies.

5 Conclusion

This paper examined methods to detect headlines generated by a GPT-2 model. A dataset was created using headlines from ABC and a survey conducted asking participants to distinguish between real and generated headlines.

Real headlines were identified as such by 55.9% of the participants, while generated ones were identified with a 47.8% rate. Various models were trained, all of which were better at identifying generated headlines than humans. BERT scored 85.7%, an improvement of around 35% over human accuracy.

Our work shows that whereas humans cannot differentiate between real and generated headlines, automatic detectors are much better at the task and therefore do have a place in the information consumption pipeline.

Acknowledgments

This work was supported by ERCAdG #740516. We want to thank the anonymous reviewers for their insightful comments and questions, and the members from the University of Sheffield who participated in our survey.

References

- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. [Plug and play language models: A simple approach to controlled text generation](#). In *International Conference on Learning Representations*.
- 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.
- Liam Dugan, Daphne Ippolito, Arun Kirubakaran, and Chris Callison-Burch. 2020. [RoFT: A tool for evaluating human detection of machine-generated text](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 189–196, Online. Association for Computational Linguistics.
- Sebastian Gehrmann, Hendrik Strobelt, and Alexander Rush. 2019. [GLTR: Statistical detection and visualization of generated text](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 111–116, Florence, Italy. Association for Computational Linguistics.
- Bin Guo, Yasan Ding, Lina Yao, Yunji Liang, and Zhiwen Yu. 2020. [The future of false information detection on social media: New perspectives and trends](#). *ACM Comput. Surv.*, 53(4).
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Jeremy Howard and Sebastian Ruder. 2018. [Universal language model fine-tuning for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Daphne Ippolito, Daniel Duckworth, Chris Callison-Burch, and Douglas Eck. 2020. [Automatic detection of generated text is easiest when humans are fooled](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1808–1822, Online. Association for Computational Linguistics.
- Ganesh Jawahar, Muhammad Abdul-Mageed, and Laks Lakshmanan, V.S. 2020. [Automatic detection of machine generated text: A critical survey](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2296–2309, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Srijan Kumar and Neil Shah. 2018. [False information on web and social media: A survey](#). *CoRR*, abs/1804.08559.
- Michael Luca and Georgios Zervas. 2013. [Fake it till you make it: Reputation, competition, and yelp review fraud](#). *SSRN Electronic Journal*.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. [Regularizing and optimizing LSTM language models](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.
- Myle Ott, Yejin Choi, Claire Cardie, and Jeffrey T. Hancock. 2011. [Finding deceptive opinion spam by any stretch of the imagination](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 309–319, Portland, Oregon, USA. Association for Computational Linguistics.
- Verónica Pérez-Rosas, Bennett Kleinberg, Alexandra Lefevre, and Rada Mihalcea. 2018. [Automatic detection of fake news](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3391–3401, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter](#).

Sandeep Subramanian, Sai Rajeswar Mudumba, Alessandro Sordani, Adam Trischler, Aaron C Courville, and Chris Pal. 2018. [Towards text generation with adversarially learned neural outlines](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 7551–7563. Curran Associates, Inc.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. [Defending against neural fake news](#).

dropout of 0.33. After the recurrent layer, we concatenate average pooling and max pooling layers. We also experiment with a Bi-LSTM with self-attention (Vaswani et al., 2017). These models are trained for 5 epochs.

A Classifier Details

ULMFit and the Transformers require their own special tokenizers, but the rest of the models use the same method, a simple indexing over the most frequent tokens. No pretrained word vectors (for example, GloVe) were used for the Deep Learning models.

ULMFit uses pre-trained weights from the AWD-LSTM model (Merity et al., 2018). For fine-tuning, we first updated the LSTM weights with a learning rate of 0.01 for a single epoch. Then, we unfroze all the layers and trained the model with a learning rate of $7.5e-5$ for an additional epoch. Finally, we trained the classifier head on its own for one more epoch with a learning rate of 0.05.

For the Transformers, we loaded pre-trained weights which we fine-tuned for a single epoch with a learning rate of $4e-5$. Specifically, the models we used were base-BERT (12 layers, 110m parameters) and DistilBERT (6 layers, 66m parameters).

The CNN has two convolutional layers on top of each other with filter sizes 8 and 4 respectively, and kernel size of 3 for both. Embeddings have 75 dimensions and the model is trained for 5 epochs.

The LSTM-based models have one recurrent layer with 35 units, while the embeddings have 100. Bidirectionality is used alongside a spatial