International Journal of Electrical and Computer Engineering (IJECE)
Vol. 13, No. 3, June 2023, pp. 2921~2930
ISSN: 2088-8708, DOI: 10.11591/ijece.v13i3.pp2921-2930

Flagging clickbait in Indonesian online news websites using fine-tuned transformers

Muhammad Noor Fakhruzzaman[1], Sa'idah Zahrotul Jannah[2], Ratih Ardiati Ningrum[1], Indah Fahmiyah[1]
[1]Data Science Technology Study Program, Faculty of Advanced Technology and Multidiscipline, Universitas Airlangga, Surabaya, Indonesia
[2]Statistics Study Program, Faculty of Science and Technology, Universitas Airlangga, Surabaya, Indonesia

Article Info

Article history:

Received Aug 15, 2022
Revised Sep 6, 2022
Accepted Oct 1, 2022

Keywords:

Adult literacy
Clickbait
Natural language processing
Online news

ABSTRACT

Click counts are related to the amount of money that online advertisers paid to news sites. Such business models forced some news sites to employ a dirty trick of click-baiting, i.e., using hyperbolic and interesting words, sometimes unfinished sentences in a headline to purposefully tease the readers. Some Indonesian online news sites also joined the party of clickbait, which indirectly degrade other established news sites' credibility. A neural network with a pre-trained language model multilingual bidirectional encoder representations from transformers (BERT) that acted as an embedding layer is then combined with a 100 node-hidden layer and topped with a sigmoid classifier was trained to detect clickbait headlines. With a total of 6,632 headlines as a training dataset, the classifier performed remarkably well. Evaluated with 5-fold cross-validation, it has an accuracy score of 0.914, an F1-score of 0.914, a precision score of 0.916, and a receiver operating characteristic-area under curve (ROC-AUC) of 0.92. The usage of multilingual BERT in the Indonesian text classification task was tested and is possible to be enhanced further. Future possibilities, societal impact, and limitations of clickbait detection are discussed.

Corresponding Author:

Muhammad Noor Fakhruzzaman
Data Science Technology Study Program, Faculty of Advanced Technology and Multidiscipline
Universitas Airlangga
St. Airlangga No. 4-6, Airlangga, Surabaya, East Java 60115, Indonesia
Email: ruzza@ftmm.unair.ac.id

1. INTRODUCTION

Journalism has changed. Before the emergence of internet news, we bought newspapers because we were enticed by the headline on the front page which usually leads to the truth, but not anymore. The emergence of online news outlets created a whole new scheme for making money in the journalism world, the online ad. With internet advertising, a single click means money, even though it is not as much as a newspaper sale or advertising money from sponsors, like in the olden days. Now, a post headline has to rake in engagement, the metric that measures ratings in online news.

The scheme of online advertising that bases on engagement have a negative influence on the original journalism idea. Sadly, online news organization now hunts for click money instead of the truth. This phenomenon promotes a unique style of headline writing, infamously known as clickbait. The more people click the post, the more engagement that post has, and the more advertising value the site will gain. A study found that most online news organization relies on clickbait's ad money to support their daily activities. With an increasing number of online news sites in recent years, they have to contest for reader's clicks [1]–[4].

Although, in a post-truth world, people tend to click on what they believed [5]. What makes it worse is that some news sources that are once credible are also retreating to the means of click-baiting. Further obscuring the integrity of the Indonesian online news organization, previous studies found that the usage of clickbait worsens the news site's reputation [6]–[9].

Clickbait refers to a headline sentence that contains hyperbolic words to persuade its reader to click the following link but mostly did not reveal any major information. It may also contain a message that is controversial but did not disclose complete information about it in the sentence [7], [10], [11]. Some clickbait headlines often use trending buzzwords, but most of its following link leads to complete misunderstanding. An example of clickbait and non-clickbait headlines is depicted in Figure 1. The first headline, which is a non-clickbait headline, translates to "21.084 Vehicles Ticketed Due to Inability to Show SIKM Jakarta". While the second headline, which is a clickbait headline, translates to "Trump Calls Putin While George Floyd Protest Enrages in the USA, What Did They Talk About?". The example shows a clear distinction between a non-clickbait headline versus a clickbait headline. A non-clickbait headline delivers straightforward key information while a clickbait headline entices us to seek more.

Figure 1. Clickbait and regular headline example, translates to "21.084 Vehicles Ticketed Due to Inability to Show SIKM Jakarta" and "Trump Calls Putin While George Floyd Protest Enrages in the USA, What Did They Talk About?"

The usage of clickbait capitalized on the human nature of curiosity. Such curiosity arises when human wants to know about something new, and they feel the gap between something they already know and something they want to know [7], [11], [12]. That curiosity gap is exploited by providing teaser messages in clickbait headline which then signals the reader about new information, provoking the reader's curiosity, and leading them to click the headline [13].

Previous studies about automatic clickbait detection used a neural network that was trained on a specifically labeled corpus of clickbait. Past researchers also tried to find a specific pattern in their corpus of clickbait, but the pattern is constantly changing over time [13], [14]. The need for automatic clickbait detection was addressed in previous studies but mostly trained on English clickbait corpus. There is a gap to be filled in Indonesian clickbait detection [15]. Indonesia needs such a tool to increase the quality of the journalism itself, while also indirectly enhancing public digital literacy as the use of clickbait in Indonesian online news enrages.

A past study that focused on detecting Indonesian fake news using neural networks, utilized frequency–inverse document frequency (TF-IDF) as their feature extraction algorithm. The term TF-IDF algorithm represented the feature of a text by counting the term or word appearance frequency in a document to express its relevance in a corpus [16]–[18]. However, term frequency is simply not enough to capture the characteristics of clickbait headlines. In order to capture semantic and syntactic properties in the headline, the text in this study is represented with word embeddings. Word embeddings is a text feature extraction technique that maps the words into a vector space model, thus representing each word as a vector and enabling computers to measure distances between words, thus returning word similarity [19]–[21].

Recently, a state-of-the-art language representation model was released, named bidirectional encoder representations from transformers (BERT). It performed the best among the available language models in completing natural language processing (NLP) tasks [22]. With the availability of the trained language model in Indonesian and other languages, this study used a pre-trained multilingual BERT (M-BERT) model as the language model. The use of transfer learning from the multilingual model enables the model to extract the features from some headlines that used both English and Indonesian in its sentence.

The approach of using a neural network to classify clickbait was deemed to be feasible due to the dynamic nature of clickbait writing [14], [15]. Moreover, the neural network performed with higher accuracy than other baseline models i.e., support vector machine, decision tree, and random forest for clickbait detection [13]. Therefore, this study attempted to use M-BERT as an embedding layer in a neural network to detect clickbait in Indonesian online news sites.

2. METHOD

The training process of the neural network used in this study is depicted in the flowchart in Figure 2. BERT weights from the flowchart refer to the pre-trained model that is used for an embedding layer. While BERT tokenizer is a built-in tokenizing factory that split the strings of headlines into tokens. Using BERT as the initial transformer preprocessor made sure that the data going into the model is in the right and proper format.

The news headline corpus was retrieved from dataset [23], consisting of 8,613 annotated news headlines from 12 online news sites. The news sites (partially redacted) are listed in Table 1. Most of the news sites are popular news sites, also accredited by the Indonesian Press Board. These news sites are chosen based on their credibility and accountability in the mass media field.

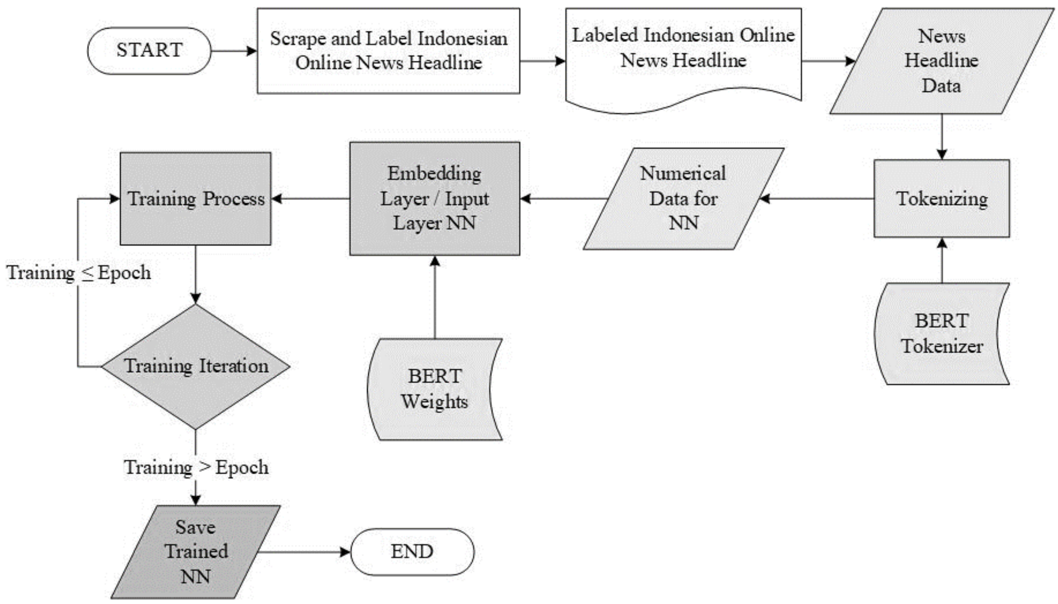

Figure 2. Training pipeline

Table 1. Lists of scraped news sites

No.	Site's URL
1	http://www.de**k.com/
2	http://www.trib**news.com/
3	http://www.pos**tro.com/
4	http://www.repu**ika.co.id/
5	http://www.ka***lagi.com/
6	http://www.k*mp*s.com/
7	http://www.te*po.co/
8	http://www.ok**one.com/
9	http://www.fim**a.com/
10	http://sin***ews.com/
11	http://www.w**ke*en.com/
12	http://www.lip**an6.com/

The dataset contains 15,000 headlines. They were labeled by 3 undergraduate students per headline and then were deemed moderately reliable with Fleiss' Kappa Interrater agreement of 0.42 [23]. However, in this study, only the headlines which every rater agreed that it is a clickbait, are selected. With that, the dataset is now consisted of 8,613 headlines, with Fleiss' Kappa of 1 which means full agreements between raters.

Hence, the dataset was deemed strongly reliable. The used dataset was available online for reproducibility. To explain further and make a clear distinction between clickbait and non-clickbait. Various sources listed the criteria of clickbait headlines [10], [24], [25]. The criteria are listed in Table 2.

Table 2. Clickbait criteria

No.	Criterion
1	Contains teaser message to entice curiosity
2	Contains pointing word(s) following images or video
3	Contains controversial word and/or buzzword
4	Contains hyperbolic sentence and/or word
5	Contains question for the reader
6	Contains emotion-provoking word

The dataset loaded directly to the python code which is also available on the GitHub link specified at the end of this article. Because the class is imbalanced, the data was normalized. 3,316 non-clickbait headlines were randomly picked to balance the dataset. The clickbait headlines still consist of 3,316 data. The total counts of data for training were 6,632 headlines. Then, using the BERT Tokenizer from Hugging Face, the text is stemmed, tokenized into words, padded, and indexed while also formatted accordingly for BERT layer input specification. The texts then had their stop words removed by using an open-source Indonesian stop word remover PySastrawi [26]. Finally, each headline in the dataset is transformed into a list of sequences of token ids. The sequence of integers, which refers to respective words in the dictionary, is ready to be fed to the neural network [27].

2.1. Neural network configuration

The neural network configuration is as follows. The input layers are two Keras input layers, each responsible for handling a list of sequences of token ids, and attention masks (the marker for pad and non-pad tokens), it is then passed forward to the BERT layer. The embedding layer is a BERT multilingual model that is trained on a multilingual Wikipedia dump dataset, which included the Indonesian language [22]. The usage of a pre-trained language model allows the researcher to capture semantics features in the headline, it also enables the researcher to extract features from the headline corpus in a short time, without wasting a lot of computing hours to train a language model.

The hidden layer consisted of 100 densely connected neurons, activated with the rectified linear unit (ReLU) function. Before the sequence is passed into the dense layer, it passes through a global average pooling layer and is flattened to fit the dense layer's dimension. Finally, the sequence is passed into a final dense layer, activated with a sigmoid function to classify it as either a clickbait class or a non-clickbait class. The model is compiled with Adam optimizer and a learning rate of 1e-05, then trained in three iterations. The network architecture is depicted in Figure 3.

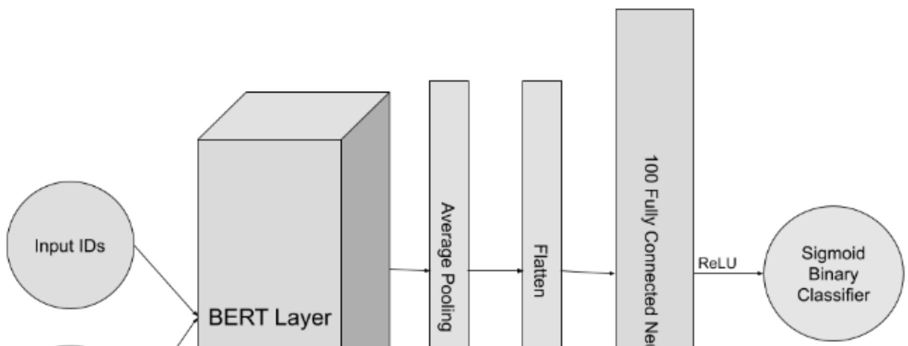

Figure 3. Network configuration

2.2. Performance evaluation

The model was evaluated using a 5-fold cross-validation method to identify its accuracy, confusion matrices, and its receiver operating characteristic-area under curve (ROC-AUC) plot. An additional evaluation was also employed. A total of 3,237 labeled headlines collected in May 2020, which is different than the training dataset collection time, was used as test data. It is then evaluated using the same metric. The additional evaluation was employed to test whether the model can detect clickbait in another dataset with different topics and possibly different clickbait sentence structures.

3. RESULTS AND DISCUSSION

A total of 6,632 headlines were used as training data. Specifically, 3,316 headlines were labeled as clickbait, and 3,316 headlines were randomly picked from a total of 5,297 non-clickbait headlines to balance the class. The undersampling is done to avoid complications in training machine learning models for imbalanced datasets.

3.1. Exploratory data analysis

In the clickbait category, the word "*ini*" or "this" in English, appears highly frequently, as seen in Figure 4. Usually, the word "*ini*" is used as a pointing word that leads the reader to the curiosity gap e.g. "This 5 Kinds Fruit is Really Good for Your Skin!" which is translated from the Indonesian sentence "*5 Macam Buah Ini Sangat Bagus Untuk Kulitmu!*". Notice the use of "*ini*" in the headline.

(a) (b)

Figure 4. Generated word cloud based on term frequency from (a) clickbait headlines and (b) regular headlines

Although it may be just one signal word of clickbait, it seems that click baiting uses this word very often. Hence, it appears as a top word in the clickbait category. Additionally, the word "*bikin*" (to make) is also appearing frequently in the clickbait category. The word "*bikin*" is considered conversational slang in Indonesia, mostly used among urban citizens. It is well-suited for clickbait because slang words are often used in clickbait headlines to bring the headline to an "easier level" so that the readers can relate to the headline with relative ease.

The word "*jadi*" can mean two different things. One means "to become" if it is used in a less formal setting. It can also be used as a conjunction word, often coupled with other sentences, which then the word "*jadi*" translates to "so" in English. The word "*jadi*" appears often both in the clickbait category and non-clickbait category, this may confuse the model because the word has different meanings in Indonesian. However, M-BERT can solve this confusion because of its word embedding-based structure that can capture semantic meaning by the context of the sentence it appeared in.

Other words in the word cloud depicted in Figure 4(a) also show some hints on which topic often appeared in the clickbait category. Celebrity names appeared in the clickbait word cloud which indicated that clickbait is often used in gossip and tabloid headlines, or "soft news". Still, some "hard news" words are also appearing often, such as "Jokowi", "BJ Habibie", and "Indonesia", which indicated that clickbait is used among the "hard news" topics as well. Meanwhile, the non-clickbait word cloud depicted in Figure 4(b) shows a lot of "hard news" related words e.g., "Indonesia", "Karhutla", "KPK", "Polisi" and no signs of neither celebrity names nor informal words.

Looking at Figure 5(a), the word "*ini*" appeared 881 times in the corpus, which is a lot more often compared to other words. Also, most of the top 10 words in clickbait did not refer directly to the topic of the article, although a good headline should refer to the topic directly. The clickbait bar chart shows a good description of the corpus and matches the clickbait criteria well, whereas in Figure 5(b), "KPK" appeared most often with 206 occurrences. Looking at all the top 10 words of the non-clickbait corpus, most of them were nouns, e.g., "Indonesia", "Habibie", "DPR". It shows that non-clickbait headline often refers to a specific topic directly, without using pointing words or conjunctions, unlike clickbait. From the word clouds and bar charts, the distinction between the two categories lies in the use of informal words, named entities, and different parts-of-speech usage.

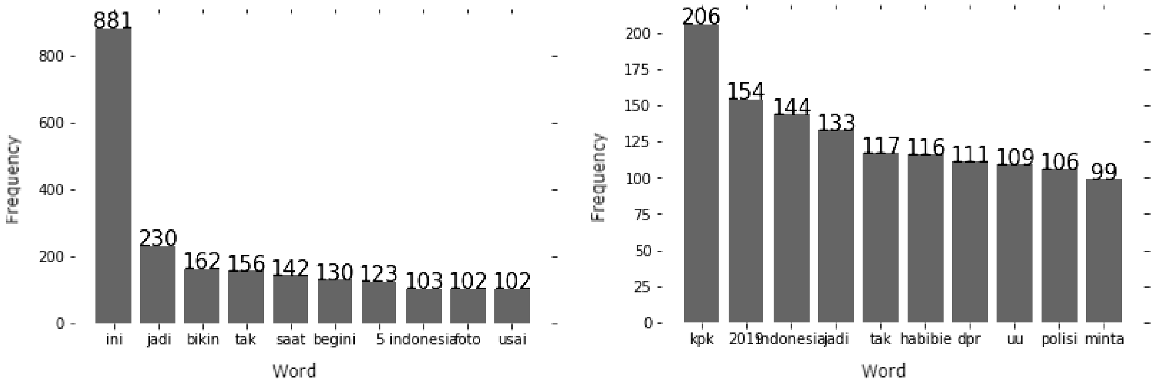

(a) (b)

Figure 5. Generated ten most frequent words for (a) clickbait headlines and (b) non-clickbait headline

3.2. Initial model evaluation

After the dataset is trained on the described neural network model for 3 iterations, it is evaluated using 5-fold cross-validation, with each fold yielded accuracy, precision, recall, and F1-score, which are then used for evaluating the model. Accuracy, precision, recall, and F1-score values indicate the performance of the model, the closer the value is to 1, the better the model in classifying headlines.

Specifically, accuracy represents the model's ability to classify each headline into its supposed class accurately. Then, precision represents the proportion of true positives i.e., predicted as clickbait turned out true, among the total predicted clickbait class. Furthermore, recall represents the proportion of true positives related to the actual clickbait class, while F1 score is a combination of precision and recall that can represent both values in equal proportion. Table 3 shows evaluation metrics from all folds in the cross-validation.

Table 3. The 5-fold cross validation evaluation

Fold	Accuracy	Precision	Recall	F1-Score
1	0.89	0.90	0.89	0.89
2	0.93	0.93	0.93	0.93
3	0.92	0.92	0.92	0.92
4	0.91	0.91	0.91	0.91
5	0.92	0.92	0.92	0.92

Then, the ROC curve and AUC are calculated to further evaluate the model. The ROC curve shows the relation between the true positive rate and the false positive rate. Ideally, the model is expected to maximize the true positive rate and minimize the false positive rate. A wider AUC with a score closer to 1 is considered a better model. Figure 6 depicts that the ROC curve of each fold is close to the y-axis, followed by a decent AUC score with an average of 0.92.

Furthermore, other classifiers, i.e. bidirectional long-short term memory (Bi-LSTM) and convolutional neural network (CNN) using configurations in [13], [14], and also an XGBoost classifier with TF-IDF vectors, were used as a comparison. All classifiers used are tested using 5-fold cross-validation, and the average accuracy of those folds is compared in Table 4. In conclusion, M-BERT combined with a simple dense layer performed better compared to other sophisticated neural network architectures (Bi-LSTM and CNN), and XGBoost with TF-IDF feature extraction in this task.

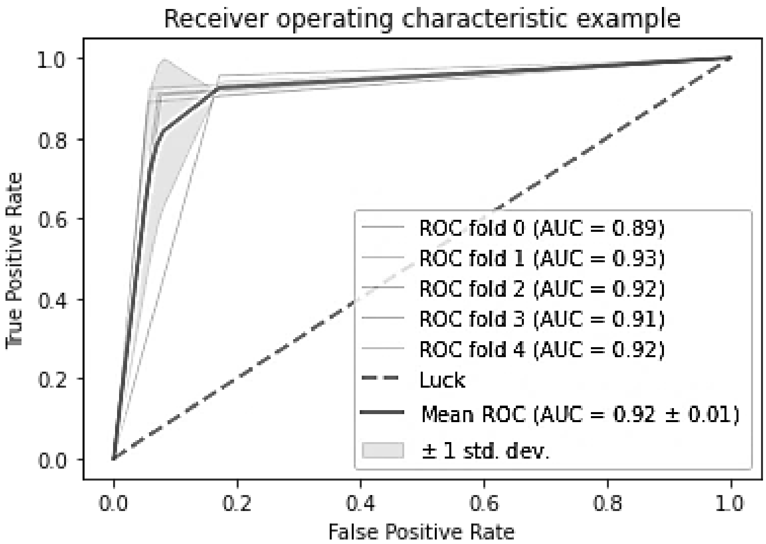

Figure 6. ROC-AUC plot

Table 4. Performance comparison with other models

Model Name	Parameters	Average Accuracy
Fine-tuned BERT	As described	0.90
Bi-LSTM	TF-IDF, Dropout 0.3, Epoch 200, Batch size 64	0.93
CNN	3 Convolution Filters (3, 4, 5), Dropout 0.5, Feature map 100, Epoch 200, Batch size 128	0.92
XGBoost	TF-IDF, n Estimators 250, seed 2, sub sample 0.7	0.91

3.3. Additional evaluation

Additionally, a further evaluation is employed to see whether or not the model can adapt to a dataset from the different collection time periods with different media agendas. Specifically, the training dataset is collected pre-coronavirus disease (COVID-19) pandemic while the additional test dataset is collected during the COVID-19 pandemic. A total of 3,237 new annotated headlines from May 2020 were pre-processed through the same method and then fed to the model to classify, then evaluated with the ground truth.

The additional evaluation shows the average accuracy of 0.83, the precision of 0.82, recall of 0.83, and F1-score of 0.83. The score indicated that the model can still detect newer headlines, given different topics, and term frequencies, although, with decreased performance. Finally, considering all of the evaluation metrics, ROC-AUC scores, and the additional evaluation, the model was deemed to perform well.

The finding indicates a possible future of using a pre-trained language model in classifying clickbait. With such a carefully curated dataset, clickbait detection can be further expanded to different NLP tasks as well [23], [28]–[30]. By using BERT, the whole model looks simplified, using only a BERT layer and a hidden standard dense layer, finally topped with a sigmoid-activated neuron, the classifier worked remarkably well with an average accuracy of 92%.

However, with additional evaluation using a newer dataset, the model performance decreased. It may be due to the different main issues in the new dataset. Further study is needed to evaluate the model's versatility. Moreover, training a neural network with M-BERT took a lot of computing resources. If efficiency is the priority, XGBoost can perform moderately well (80% average accuracy).

This study adds to the body of research in mass communication studies by providing an alternative tool for detecting clickbait in online news headlines, specifically in Indonesian. There are other methods, however, both quantitatively and qualitatively that scientifically assess clickbait detection, mostly in linguistic and mass communication studies. Indeed, this study chose to contribute from the perspective of NLP and expand the body of research in applied machine learning by implementing a neural network in the Indonesian mass media domain. Future researchers may delve deeper into the area of applied machine learning for mass media by looking into other methods or optimizing the current method.

The presence of clickbait detection may help the public to increase digital literacy by giving extra attention to the flagged headline so that they know what specific characteristics a clickbait headline has. Likewise, it might be able to discourage some news sites from using clickbait and shift their practices slowly into more high-quality reporting. Although, it seems unlikely due to the current funding scheme in the online advertisement domain. Furthermore, if the model is deployed into usable software, it can help to alleviate the bias given by the priming words that are widely used in writing clickbait headlines. Flagging the clickbait headline gives a chance for the reader to rethink their urge to click and fill their curiosity gap [7], [11], [12], [31]–[33].

3.4. Limitations and future research

This study has some limitations, one of which is the broad topics of news in the dataset. Such broad topics may provide noise because clickbait headlines often dominated celebs and gossip topics. It may affect future prediction due to the noise provided by the mentioned topic. Future research may fill the gap by focusing on specific hard news topics, such as politics, to learn further about clickbait usage in different topics.

Additionally, the labeling of clickbait is relatively hard, even for humans. Future research may expand the dataset and select only the data where every rater agreed so that there is a clear distinction between clickbait and non-clickbait headlines. Also, the initial briefing of the rater should be conducted systematically, so that the dataset is reliable and unambiguous. The training dataset also needed expansions by adding more data from multiple time periods of collection. This may increase the model's versatility by enabling it to capture the dynamic pattern of clickbait structure in various time periods.

Furthermore, this study used pre-trained BERT using Indonesian Wikipedia dumps as well as other 100 languages, which may not contain sensational wording commonly used in clickbait headlines, due to Wikipedia writing rules. Therefore, future researchers may collect a bigger Indonesian corpus that includes offensive and rude words, slang words, and sensational words to fine-tune the BERT model, which might increase the performance of the model. Yet, a qualitative approach regarding clickbait assessment is also needed to define clickbait characteristics thoroughly. With a detailed specification of clickbait headlines, the labeling process of headlines can be less biased. Qualitative study can also confirm the descriptive analysis of this study which stated that clickbait headlines often used non-topic related words and teasing words.

4. CONCLUSION

Using the neural network in classifying clickbait has been fairly common in the English language, but not in Indonesian. This study contributes to showing that multilingual BERT, a state-of-the-art model is able to classify Indonesian clickbait headlines, given a proper model to fine-tune the language model.

Furthermore, we would like to explore more about the methods for detecting clickbait of online news headlines. We also want to deploy the model into a usable component, like a browser extension program, so that the clickbait detector can be used by the public. Finally, future researchers can look into the effect of the clickbait detector on the general public. Whether or not it influences adult literacy and its ability to inhibit the spread of misinformation.

ADDITIONAL RESOURCES

The complete Python notebook and datasets are stored on *https://github.com/ruzcmc/ClickbaitIndo-textclassifier*.

REFERENCES

[1] Y. Chen, N. K. Conroy, and V. L. Rubin, “News in an online world: The need for an ‘automatic crap detector,’” *Proceedings of the Association for Information Science and Technology*, vol. 52, no. 1, pp. 1–4, Jan. 2015, doi: 10.1002/pra2.2015.145052010081.

[2] A. Chakraborty, B. Paranjape, S. Kakarla, and N. Ganguly, “Stop clickbait: Detecting and preventing clickbaits in online news media,” in *ASONAM ’16: Proceedings of the 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2016, pp. 9–16.

[3] T. Jiang, Q. Guo, S. Chen, and J. Yang, “What prompts users to click on news headlines? Evidence from unobtrusive data analysis,” *Aslib Journal of Information Management*, vol. 72, no. 1, pp. 49–66, Oct. 2019, doi: 10.1108/AJIM-04-2019-0097.

[4] J. R. Collier, J. Dunaway, and N. J. Stroud, “Pathways to deeper news engagement: Factors influencing click behaviors on news sites,” *Journal of Computer-Mediated Communication*, vol. 26, no. 5, pp. 265–283, Sep. 2021, doi: 10.1093/jcmc/zmab009.

[5] M. Makhortykh, C. de Vreese, N. Helberger, J. Harambam, and D. Bountouridis, “We are what we click: Understanding time and content-based habits of online news readers,” *New Media & Society*, vol. 23, no. 9, pp. 2773–2800, Sep. 2021, doi: 10.1177/1461444820933221.

[6] N. Hurst, “To clickbait or not to clickbait? an examination of clickbait headline effects on source credibility,” University of Missouri-Columbia, 2016.

[7] K. Scott, "You won't believe what's in this paper! Clickbait, relevance and the curiosity gap," *Journal of Pragmatics*, vol. 175, pp. 53–66, Apr. 2021, doi: 10.1016/j.pragma.2020.12.023.
[8] K. Munger, "All the news that's fit to click: The economics of clickbait media," *Political Communication*, vol. 37, no. 3, pp. 376–397, May 2020, doi: 10.1080/10584609.2019.1687626.
[9] L. Molyneux and M. Coddington, "Aggregation, clickbait and their effect on perceptions of journalistic credibility and quality," *Journalism Practice*, vol. 14, no. 4, pp. 429–446, Apr. 2020, doi: 10.1080/17512786.2019.1628658.
[10] M. Potthast, S. Köpsel, B. Stein, and M. Hagen, "Clickbait detection," in *ECIR 2016: Advances in Information Retrieval*, 2016, pp. 810–817, doi: 10.1007/978-3-319-30671-1_72.
[11] R. Golman and G. Loewenstein, "Information gaps: A theory of preferences regarding the presence and absence of information," *Decision*, vol. 5, no. 3, pp. 143–164, Jul. 2018, doi: 10.1037/dec0000068.
[12] G. Loewenstein, "The psychology of curiosity: A review and reinterpretation," *Psychological Bulletin*, vol. 116, no. 1, pp. 75–98, Jul. 1994, doi: 10.1037/0033-2909.116.1.75.
[13] A. Anand, T. Chakraborty, and N. Park, "We used neural networks to detect clickbaits: You won't believe what happened next!," in *ECIR 2017: Advances in Information Retrieval*, 2017, pp. 541–547, doi: 10.1007/978-3-319-56608-5_46.
[14] A. Agrawal, "Clickbait detection using deep learning," in *2016 2nd International Conference on Next Generation Computing Technologies (NGCT)*, Oct. 2016, pp. 268–272, doi: 10.1109/NGCT.2016.7877426.
[15] N. A. Zuhroh and N. A. Rakhmawati, "Clickbait detection: A literature review of the methods used," *Register: Jurnal Ilmiah Teknologi Sistem Informasi*, vol. 6, no. 1, pp. 1–10, Oct. 2019, doi: 10.26594/register.v6i1.1561.
[16] A. Rusli, J. C. Young, and N. M. S. Iswari, "Identifying fake news in Indonesian via supervised binary text classification," in *2020 IEEE International Conference on Industry 4.0, Artificial Intelligence, and Communications Technology (IAICT)*, Jul. 2020, pp. 86–90, doi: 10.1109/IAICT50021.2020.9172020.
[17] S. Qaiser and R. Ali, "Text mining: Use of TF-IDF to examine the relevance of words to documents," *International Journal of Computer Applications*, vol. 181, no. 1, pp. 25–29, Jul. 2018, doi: 10.5120/ijca2018917395.
[18] A. Aizawa, "An information-theoretic perspective of TF–IDF measures," *Information Processing & Management*, vol. 39, no. 1, pp. 45–65, Jan. 2003, doi: 10.1016/S0306-4573(02)00021-3.
[19] Y. Chen, B. P., R. Al-Rfou, and S. Skiena, "The expressive power of word embeddings," in *ICLR 2013 conference*, 2013.
[20] R. Al-Rfou, B. Perozzi, and S. Skiena, "Polyglot: Distributed word representations for multilingual NLP," in *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, 2013, pp. 183–192.
[21] Y. Li and T. Yang, "Word embedding for understanding natural language: A survey," in *Guide to Big Data Applications*, vol. 26, S. Srinivasan, Ed. Cham: Springer International Publishing, 2018, pp. 83–104, doi: 10.1007/978-3-319-53817-4_4.
[22] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of NAACL-HLT 2019*, 2019, pp. 4171–4186.
[23] A. William and Y. Sari, "CLICK-ID: A novel dataset for Indonesian clickbait headlines," *Data in Brief*, vol. 32, Oct. 2020, doi: 10.1016/j.dib.2020.106231.
[24] P. Biyani, K. Tsioutsiouliklis, and J. Blackmer, "'8 amazing secrets for getting more clicks': Detecting clickbaits in news streams using article informality," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30, no. 1, Feb. 2016, doi: 10.1609/aaai.v30i1.9966.
[25] M. Potthast *et al.*, "Crowdsourcing a large corpus of clickbait on Twitter," in *Proceedings of the 27th International Conference on Computational Linguistics*, 2018, pp. 1498–1507.
[26] H. Robbani, "Github: Indonesian stemmer. Python port of PHP Sastrawi project," *Github*, 2018. https://github.com/har07/PySastrawi (accessed May 30, 2020).
[27] T. Wolf *et al.*, "Transformers: State-of-the-art natural language processing," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2020, pp. 38–45, doi: 10.18653/v1/2020.emnlp-demos.6.
[28] K. Clark, M.-T. Luong, U. Khandelwal, C. D. Manning, and Q. V. Le, "BAM! Born-again multi-task networks for natural language understanding," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2019, pp. 5931–5937, doi: 10.18653/v1/P19-1595.
[29] J. Worsham and J. Kalita, "Multi-task learning for natural language processing in the 2020s: Where are we going?," *Pattern Recognition Letters*, vol. 136, pp. 120–126, Aug. 2020, doi: 10.1016/j.patrec.2020.05.031.
[30] A. C. Stickland and I. Murray, "BERT and PALs: Projected attention layers for efficient adaptation in multi-task learning," *Prepr. arXiv1902.02671v2*, pp. 1–10, 2019.
[31] S. Pandey and G. Kaur, "Curious to click it?-Identifying clickbait using deep learning and evolutionary algorithm," in *2018 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, Sep. 2018, pp. 1481–1487, doi: 10.1109/ICACCI.2018.8554873.
[32] V. Apresjan and A. Orlov, "Pragmatic mechanisms of manipulation in Russian online media: How clickbait works (or does not)," *Journal of Pragmatics*, vol. 195, pp. 91–108, Jul. 2022, doi: 10.1016/j.pragma.2022.02.003.
[33] S. (Fone) Pengnate, J. Chen, and A. Young, "Effects of clickbait headlines on user responses: An empirical investigation," *Journal of International Technology and Information Management*, vol. 30, no. 3, 2021.

BIOGRAPHIES OF AUTHORS

Muhammad Noor Fakhruzzaman was born and raised in Surabaya. He holds an interdisciplinary master's degree in human-computer interaction and journalism and mass communication from Iowa State University. His current research interests fall between data science and mass communication, mainly automated media monitoring using natural language processing. He currently teaches in Data Science Technology Study Program at Universitas Airlangga. His research includes a brain-computer interface, Indonesian natural language processing, and media monitoring. He is a member of Kappa Tau Alpha, an honor society of journalism and mass communication studies. He also loves to train in grappling sports: wrestling, Brazilian Jiu-jitsu (2nd-degree blue belt), and fanatically watches pro-wrestling shows. He can be contacted at ruzza@ftmm.unair.ac.id or ruzza@alumni.iastate.edu.

Sa'idah Zahrotul Jannah was born and raised in Surabaya. She got her bachelor's and master's degree in Statistics from Institut Teknologi Sepuluh Nopember, Indonesia. Currently, she is a lecturer at Universitas Airlangga, Indonesia. Her research interest is data mining, multivariate analysis, and natural language processing. She can be contacted at s.zahrotul.jannah@fst.unair.ac.id.

Ratih Ardiati Ningrum received her bachelor's degree in Statistics from the Institut Teknologi Sepuluh Nopember, Surabaya, Indonesia in 2017. In the same year, she took a master's program in Statistics at the same institute. She had a joint degree with the Institute of Statistics at National Chiao Tung University, Hsinchu, Taiwan in the second year while undergoing a master's program. She received her master's degree in 2019. She is currently a young researcher at the Technology of Data Science, Faculty of Advanced Technology and Multidiscipline, Airlangga University, Indonesia. Her research interests include statistics and machine learning. She can be reached at ratih.an@ftmm.unair.ac.id.

Indah Fahmiyah is a lecturer at the Data Science Technology Study Program of the Faculty of Advanced Technology and Multidiscipline, Universitas Airlangga, Indonesia. Her educational background is bachelor's and master's degrees in Statistics from Institut Teknologi Sepuluh Nopember Surabaya, Indonesia. Her research interests are predictive analytics, machine learning, data visualization, and data mining. She is interested in weather and health data. She can be reached at indah.fahmiyah@ftmm.unair.ac.id.