

SWDE : A Sub-Word And Document Embedding Based Engine for Clickbait Detection *

Vaibhav Kumar, Mrinal Dhar, Dhruv Khattar
International Institute of Information Technology
Hyderabad
{vaibhav.kumar,mrinal.dhar,dhruv.khattar}@research.iiit.ac.in

Abhimanshu Mishra
Birla Institute of Technology, Mesra
abhimanshu@gmail.com

Yash Kumar Lal[†]
Manipal Institute of Technology, Manipal
yash.kumar4@learner.manipal.edu

Manish Shrivastava, Vasudeva Varma
International Institute of Information Technology
Hyderabad
{m.shrivastava,vv}@iiit.ac.in

ABSTRACT

In order to expand their reach and increase website ad revenue, media outlets have started using clickbait techniques to lure readers to click on articles on their digital platform. Having successfully enticed the user to open the article, the article fails to satiate his curiosity serving only to boost click-through rates. Initial methods for this task were dependent on feature engineering, which varies with each dataset. Industry systems have relied on an exhaustive set of rules to get the job done. Neural networks have barely been explored to perform this task. We propose a novel approach considering different textual embeddings of a news headline and the related article. We generate sub-word level embeddings of the title using Convolutional Neural Networks and use them to train a bidirectional LSTM architecture. An attention layer allows for calculation of significance of each term towards the nature of the post. We also generate Doc2Vec embeddings of the title and article text and model how they interact, following which it is concatenated with the output of the previous component. Finally, this representation is passed through a neural network to obtain a score for the headline. We test our model over 2538 posts (having trained it on 17000 records) and achieve an accuracy of 83.49% outscoring previous state-of-the-art approaches.

KEYWORDS

Clickbait, Neural Networks, Attention Mechanism, Sub-Word Embeddings, Document Embeddings

*The first five authors have equal contribution.

[†]The author was an intern at International Institute of Information Technology Hyderabad when this work was done.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Comps'18, SIGIR, July 8-12, Ann Arbor, MI USA
© 2018 Copyright held by the owner/author(s).
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

ACM Reference Format:

Vaibhav Kumar, Mrinal Dhar, Dhruv Khattar, Yash Kumar Lal, Abhimanshu Mishra, and Manish Shrivastava, Vasudeva Varma. 2018. SWDE : A Sub-Word And Document Embedding Based Engine for Clickbait Detection . In *Proceedings of SIGIR 2018 Workshop on Computational Surprise in Information Retrieval (Comps'18, SIGIR)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

In recent years, content delivery has changed drastically, shifting from offline methods to the Internet. It is now the primary source of information for a majority of the populace, especially for ever-changing news updates. This has also caused a shift in users' preferred sources. Previously, these preferences were static, sticking to a particular news source. Now, with the plethora of information available easily, there is no differentiation in the source it has been gathered from, with users opting to go for whatever is convenient.

Keeping up with the times, news agencies have expanded their digital presence, increasing their reach exponentially. They generate revenue by (1) advertisements on their websites, or (2) a subscription based model for articles that might interest users. Since multiple agencies offer similar content, the user has his pick. To lure in more readers and increase the number of clicks on their content, subsequently enhancing their agency's revenue, writers have begun adopting a new technique - clickbait.

Merriam-Webster defines clickbait as something (such as a headline) to encourage readers to click on hyperlinks based on snippets of information accompanying it, especially when those links lead to content of dubious value or interest. It is built to create, and consequently capitalise, on the Loewenstein information gap [12] by purposefully misrepresenting or promising what can be expected while reading a story on the web, be it through a headline, image or related text.

We propose a two-pronged approach to detect such headlines. The first component leverages distributional semantics of the title text and models its temporal and sequential properties. The article title is represented as a concatenation of its sub-word level embeddings. The sub-word representation serves as input to a bidirectional LSTM network. The contribution of a sub-word towards the clickbait nature of the headline is calculated in a differential manner since the output of the LSTM is passed into an attention

layer [2], following which it goes through a dense layer. The second component focuses on Doc2Vec embeddings of the title and article content, performing an element wise multiplication of the two. This is concatenated with the dense layer output from the previous component. The obtained output is then passed through multiple hidden layers which performs the final classification.

Previous work in this field that has exploited the power of embeddings has considered either word vectors, for their ability to create context-sensitive word representations, or character-level word embeddings to model the orthographic features of a word. We propose the use of sub-word level representations since it incorporates the word's morphological features. Attaching an attention mechanism to it helps us identify the surprise associated with each representation within the clickbait. One of the identifying characteristics of clickbait is that the article title differs from the text attached to it. For this reason, we define a component to capture the interaction between these attributes and augment our model.

2 RELATED WORK

The importance of detecting clickbait headlines has increased exponentially in recent years. Initial work in this domain can be traced back to [3], relying on heavy feature engineering on a specific news dataset. These works define the various types of clickbait and focus on the presence of linguistic peculiarities in the headline text, including various informality metrics and the use of forward references. Applying such techniques over a social media stream was first attempted by [16] as the authors crowdsourced a dataset of tweets [15] and performed feature engineering to accomplish the task. [5] have tried to expand the work done for news headlines they collected from trusted sources.

[1] used the same collection of headlines as [5] and proposed the first neural network based approach in the field. They employed various recurrent neural network architectures to model sequential data and its dependencies, taking as its inputs a concatenation of the word and character-level embeddings of the headline. Their experiments yielded that bidirectional LSTMs [19] were best suited for the same. [20] built BiLSTMs to model each textual attribute of the post (post-text, target-title, target-paragraphs, target-description, target-keywords, post-time) available in the corpus [15], concatenating their outputs and feeding it to a fully connected layer to classify the post. Attention mechanisms [2] have grown popular for various text classification tasks, like aspect based sentiment analysis. Utilising this technique, [22] deployed a self-attentive bidirectional GRU to infer the importance of each tweet token and model the annotation distribution of headlines in the corpus.

Word vectors and character vectors have been used across various approaches proposed to solve this problem. However, we suggest the use of subword representations to better analyse the morphology of possible clickbait-y words. We also attempt to model the interaction between the title of an article and its text.

3 MODEL ARCHITECTURE

We now describe our approach to clickbait detection and the reasons behind devising such a model. Our approach is a fusion of multiple components, each exploiting a particular type of embedding: (1)




Figure 1: Architecture for learning Sub-word Level Representations using CNN

BiLSTM with attention, and (2) Doc2Vec enrichment. Figure 2 lays out our proposed architecture.

We start with an explanation of the various types of embeddings we have used and proceed to describe the various components of our model, both individually and together. Finally, we cover how the parameters are learned.

3.1 Sub-word Level Representation

Word2Vec [13] has fast become the most popular text embedding method for text since it models a word based on its context. [9] proposed a convolutional neural network architecture to generate subword-level representations of words in order to capture word orthography. Sub-word level embeddings learn representations for character n-grams and represent words as the sum of the n-gram vectors [4]. Such representations also take into account word roots and inflections, rather than just word context. They work well even with highly noisy text with containing misspellings due to the model learning morpheme-level feature maps. They have proven to be extremely useful in tasks such as sentiment analysis [17], PoS tagging [14] and language modeling [9]. These intermediate sub-word feature representations are learned by the filters during the convolution operation. We generate such an embedding by passing the characters of a sentence individually into 3 layer 1D convolutional neural network. Each filter then acts as a learned sub-word level feature. A representation for this architecture can be found in Figure 1.

3.2 Document Embeddings

Doc2Vec [11] is an unsupervised approach to generate vector representations for slightly larger bodies of text, such as sentences, paragraphs and documents. It has been adapted from Word2Vec [13] which is used to generate vectors for words in large unlabeled corpora. The vectors generated by this approach come handy in tasks like calculating similarity metrics for sentences, paragraphs and documents. In sequential models like RNNs, the word sequence is captured in the generated sentence vectors. However, in Doc2Vec, the representations are order independent. We use GenSim [18] to learn 300 dimensional Doc2Vec embeddings for each target description and post title available.

3.3 Bidirectional LSTM with Attention

Recurrent Neural Network (RNN) is a class of artificial neural networks which utilizes sequential information and maintains history through its intermediate layers. A standard RNN has an internal state whose output at every time-step which can be expressed in terms of that of previous time-steps. However, it has been seen that standard RNNs suffer from a problem of vanishing gradients [8]. This means it will not be able to efficiently model dependencies and interactions between sub-word representations that are a few steps apart. LSTMs are able to tackle this issue by their use of gating mechanisms. We convert each article headline into its corresponding sub-word level representation to act as input to our bidirectional LSTMs.

$(\vec{h}_1, \vec{h}_2, \dots, \vec{h}_R)$ represent forward states of the LSTM and its state updates satisfy the following equations:

$$[\vec{f}_t, \vec{i}_t, \vec{o}_t] = \sigma[\vec{W}[\vec{h}_{t-1}, \vec{r}_t] + \vec{b}] \quad (1)$$

$$\vec{l}_t = \tanh[\vec{V}[\vec{h}_{t-1}, \vec{r}_t] + \vec{d}] \quad (2)$$

$$\vec{c}_t = \vec{f}_t \cdot \vec{c}_{t-1} + \vec{i}_t \cdot \vec{l}_t \quad (3)$$

$$\vec{h}_t = \vec{o}_t \cdot \tanh(\vec{c}_t) \quad (4)$$

here σ is the logistic sigmoid function, $\vec{f}_t, \vec{i}_t, \vec{o}_t$ represent the forget, input and output gates respectively. \vec{r}_t denotes the input at time t and \vec{h}_t denotes the latent state, \vec{b}_t and \vec{d}_t represent the bias terms. The forget, input and output gates control the flow of information throughout the sequence. \vec{W} and \vec{V} are matrices which represent the weights associated with the connections.

$(\overleftarrow{h}_1, \overleftarrow{h}_2, \dots, \overleftarrow{h}_R)$ denote the backward states and its updates can be computed similarly.

The number of bidirectional LSTM units is set to a constant K , which is the maximum length of all title lengths of records used in training. The forward and backward states are then concatenated to obtain (h_1, h_2, \dots, h_K) , where

$$h_i = \begin{bmatrix} \vec{h}_i \\ \overleftarrow{h}_i \end{bmatrix} \quad (5)$$

Finally, we are left with the task of figuring out the significance of each word in the sequence i.e. how much a particular sub-word representation influences the clickbait-y nature of the post. The effectiveness of attention mechanisms have been proven for the task of neural machine translation [2] and it has the same effect in this case. The goal of attention mechanisms in such tasks is to derive context vectors which capture relevant source side information and help predict the current target representation. The sequence of annotations generated by the encoder to come up with a context vector capturing how each sub-word contributes to the record's clickbait quotient is of paramount importance to this model. In a typical RNN encoder-decoder framework [2], a context vector is generated at each time-step to predict the target sub-word. However, we only need it for calculation of context vector for a single time-step.

$$c_{attention} = \sum_{j=1}^K \alpha_j h_j \quad (6)$$

where, h_1, \dots, h_K represents the sequence of annotations to which the encoder maps the post title vector and each α_j represents the respective weight corresponding to each annotation h_j . This is represented as the left most component in Figure 2.

3.4 Doc2Vec Enrichment

Each record in the dataset has a target description attached with it. This is the entire text of the article whose title has been given. By definition, clickbait articles differ from the content described in their headline. We generate document embeddings for both the title and the article text and perform element wise multiplication over the two. This allows us to capture the interaction between the two, something which has not been used before. Since the title is supposed to mislead the reader with respect to the content, modeling this interaction in terms of their similarity gives an added dimension to our approach. It augments the output obtained from the first component.

3.5 Fusion of Components

The outputs from the aforementioned components are now concatenated and passed through two dense layers and finally goes into a fully connected layer. This layer finally gives out the probability that a post can be marked clickbait.

3.6 Learning the Parameters

We use binary cross-entropy as the loss optimization function for our model. The cross-entropy method [6] is an iterative procedure where each iteration can be divided into two stages:

- (1) Generate a random data sample (vectors, trajectories etc.) according to a specified mechanism.
- (2) Update the parameters of the random mechanism based on the data to produce a "better" sample in the next iteration.

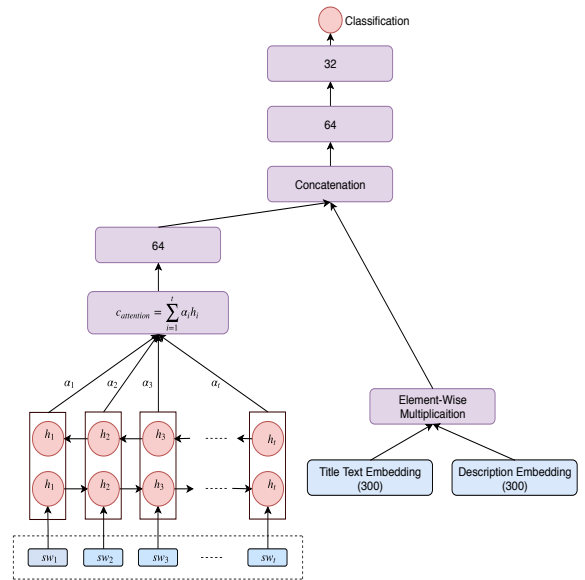


Figure 2: Full Model Architecture

Table 1: Model Performance Comparison

Model	F1 Score	Accuracy
Proposed Approach	0.63	83.49%
BiLSTM [1]	0.61	83.28%
Feature Engineering SotA [16]	0.55	83.24%
Concatenated NN Architecture [20]	0.39	74%

4 EVALUATION AND RESULTS

[15] crowdsourced the annotation of 19538 tweets they had curated, into various levels of their clickbait-y nature. These tweets contained the title and text of the article and also included supplementary information such as target description, target keywords and linked images. We trained our model over 17000 records in the described dataset and test it over 2538 disjoint instances from the same. We performed our experiments with the aim of increasing the accuracy and F1 score of the model. Other metrics like mean squared error (MSE) were also considered.

4.1 Training

We randomly partition the training set of over 17000 posts into training and validation set in a 4:1 ratio. This ensures that the two sets do not overlap. The model hyperparameters are tuned over the validation set. We initialise the fully connected network weights with the uniform distribution in the range $-\sqrt{6/(fanin + fanout)}$ and $\sqrt{6/(fanin + fanout)}$ [7]. We used a batch size of 256 and adadelta [21] as a gradient based optimizer for learning the model parameters.

4.2 Model Comparison

In Table 1, we evaluate our model against the existing state-of-the-art for the dataset used and other models which have employed similar techniques to accomplish the task. It is clear that our proposed model outperforms the previous feature engineering benchmark and other work done in the field both in terms of F1 score and accuracy of detection. Feature engineering models rely on a selection of handcrafted attributes which may not be able to consider all the factors involved in making a post clickbait. The approach proposed in [20] takes into account each of the textual features available in an individual fashion, considering them to be independent of each other, which is not the case since, by definition of clickbait, the content of the article title and text are not mutually exclusive. [10] proposed the integration of multimodal embeddings. [1] utilise word and character embeddings which do not capture morpheme-level information that may incorporate a surprise element.

5 CONCLUSION AND FUTURE WORK

We have devised an approach to detecting clickbait that puts emphasis on utilising the linguistic value of words by learning its morphological features through its sub-word representations. These embeddings and their dependencies are, in turn, modeled by the LSTM. Attention mechanism allows us to understand the importance of individual representations towards the nature of the post. Using the document embeddings for title and article text allows us to augment the generated embeddings and use as input to a neural

network to finally classify the post. In the future, we would like to explore the possibility of integrating the sub-word representations with deep neural networks to better model the temporal and sequential properties of text.

REFERENCES

- [1] Ankesh Anand, Tanmoy Chakraborty, and Noseong Park. 2017. We used Neural Networks to Detect Clickbaits: You won't believe what happened Next!. In *Advances in Information Retrieval. 39th European Conference on IR Research (ECIR 17) (Lecture Notes in Computer Science)*. Springer.
- [2] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014).
- [3] Prakhar Biyani, Kostas Tsioutsoulis, and John Blackmer. 2016. "8 Amazing Secrets for Getting More Clicks": Detecting Clickbaits in News Streams Using Article Informality. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*.
- [4] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. Enriching Word Vectors with Subword Information. *CoRR abs/1607.04606* (2016). <http://arxiv.org/abs/1607.04606>
- [5] Abhijnan Chakraborty, Bhargavi Paranjape, Sourya Kakarla, and Niloy Ganguly. 2016. Stop Clickbait: Detecting and preventing clickbaits in online news media. *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)* (2016), 9–16.
- [6] Pieter-Tjerk de Boer, Dirk P. Kroese, Shie Mannor, and Reuven Y. Rubinfeld. 2005. A Tutorial on the Cross-Entropy Method. *Annals of Operations Research* 134, 1 (01 Feb 2005), 19–67. <https://doi.org/10.1007/s10479-005-5724-z>
- [7] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Aistats*, Vol. 9. 249–256.
- [8] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [9] Yoon Kim, Yacine Jernite, David Sontag, and Alexander M. Rush. 2016. Character-aware Neural Language Models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI'16)*.
- [10] Vaibhav Kumar, Dhruv Khattar, Yash Kumar Lal, and Vasudeva Varma. 2018. Identifying Clickbait: A Multi-Strategy Approach Using Neural Networks. In *Proceedings of the 41st International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '18)*.
- [11] Quoc Le and Tomas Mikolov. 2014. Distributed representations of sentences and documents. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*. 1188–1196.
- [12] George Loewenstein. 1994. The Psychology of Curiosity: A Review and Reinterpretation. 116 (07 1994), 75–98.
- [13] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *CoRR abs/1301.3781* (2013). <http://arxiv.org/abs/1301.3781>
- [14] Barbara Plank, Anders Søgaard, and Yoav Goldberg. 2016. Multilingual Part-of-Speech Tagging with Bidirectional Long Short-Term Memory Models and Auxiliary Loss. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, 412–418. <https://doi.org/10.18653/v1/P16-2067>
- [15] Martin Potthast, Tim Gollub, Kristof Komlosy, Sebastian Schuster, Matti Wiegmann, Erika Garces, Matthias Hagen, and Benno Stein. 2017. Crowdsourcing a Large Corpus of Clickbait on Twitter. In *(to appear)*.
- [16] Martin Potthast, Sebastian Köppl, Benno Stein, and Matthias Hagen. 2016. Clickbait Detection. In *Advances in Information Retrieval. 38th European Conference on IR Research (ECIR 16) (Lecture Notes in Computer Science)*, Nicola Ferro, Fabio Crestani, Marie-Francine Moens, Josiane Mothe, Fabrizio Silvestri, Giorgio Maria Di Nunzio, Claudia Hauff, and Gianmaria Silvello (Eds.), Vol. 9626. Springer, Berlin Heidelberg New York, 810–817. https://doi.org/10.1007/978-3-319-30671-1_72
- [17] Ameya Prabhu, Aditya Joshi, Manish Shrivastava, and Vasudeva Varma. 2016. Towards Sub-Word Level Compositions for Sentiment Analysis of Hindi-English Code Mixed Text. *CoRR abs/1611.00472* (2016). [arXiv:1611.00472](http://arxiv.org/abs/1611.00472) <http://arxiv.org/abs/1611.00472>
- [18] Radim Řehůřek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. 45–50.
- [19] Mike Schuster and Kuldeep K Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing* 45, 11 (1997), 2673–2681.
- [20] Philippe Thomas. 2017. Clickbait Identification using Neural Networks. *CoRR abs/1710.08721* (2017). [arXiv:1710.08721](http://arxiv.org/abs/1710.08721) <http://arxiv.org/abs/1710.08721>
- [21] Matthew D Zeiler. 2012. ADADELTA: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701* (2012).
- [22] Yiwei Zhou. 2017. Clickbait Detection in Tweets Using Self-attentive Network. *CoRR abs/1710.05364* (2017). [arXiv:1710.05364](http://arxiv.org/abs/1710.05364) <http://arxiv.org/abs/1710.05364>