

“Attention” for Detecting Unreliable News in the Information Age

Venkatesh Duppada

Seernet Technologies, LLC
venkatesh.duppada@seernet.io

Abstract

An Unreliable news is any piece of information which is false or misleading, deliberately spread to promote political, ideological and financial agendas. Recently the problem of unreliable news has got a lot of attention as the number instances of using news and social media outlets for propaganda have increased rapidly. This poses a serious threat to society, which calls for technology to automatically and reliably identify unreliable news sources. This paper is an effort made in this direction to build systems for detecting unreliable news articles. In this paper, various NLP algorithms were built and evaluated on Unreliable News Data 2017 dataset. Variants of hierarchical attention networks (HAN) are presented for encoding and classifying news articles which achieve the best results of 0.944 ROC-AUC. Finally, Attention layer weights are visualized to understand and give insight into the decisions made by HANs. The results obtained are very promising and encouraging to deploy and use these systems in the real world to mitigate the problem of unreliable news.

1 Introduction

We are living in the digital age where information is disseminated with ease at great speeds using social media and online news outlets. The ability to make information freely available has gifted people from all walks of life equal access to information. As good as this sounds, it has adverse effects when used for spreading false and misleading information. Recently we have been witnessing the instances where social media and news outlets are being used for propaganda on day to day basis. 2016 US presidential election is one of the many notable cases which has witnessed the real menace of unreliable news. As more and more people spend increasing amount of time online for consumption of news, it is becoming crucial to prevent the spread of false and unreliable news. The unreliable news sources are also a serious threat to journalism as it casts a shadow over the credibility of media. People stop trusting any factual information and think of it as one of many possible truths offered by media. Because of these reasons, the task of identifying unreliable news is more important now than ever before. This paper is an effort in the direction of building automated systems that can reliably identify unreliable news.

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The rest of the paper is laid out as follows: Section 2 discusses the annotation statistics and the data format of Unreliable News Data 2017 (unr 2018) dataset on which the system evaluations are made. Section 3 lists down the traditional NLP baseline systems. In section 4 variants of hierarchical attention networks for encoding and classifying news articles are presented. Section 5 presents the results of all the systems developed. In section 6 predictions made by the systems are analyzed. Finally, the paper is concluded with Section 7.

2 Dataset

For the purpose of the task to identify unreliable media and build automatic detectors an existing corpus of data containing news articles is chosen. These articles are then marked with labels signifying whether or not each article is reliable as determined by Open Source Unreliable Media List¹. From now on authors refer to **Unreliable News Data 2017** (unr 2018) dataset as **UND17**. The dataset provided consists of predefined train and test splits and authors use the same splits for training and evaluation. Table 1 and 2 provide annotation counts and format of the dataset.

Unreliable News Data 2017		
	Reliable Articles	Unreliable Articles
Train	95,295	34,524
Test	1,00,247	38,494

Table 1: UND17 dataset annotation counts

Unreliable News Data 2017	
Field	Description
uid	Unique identifier for news article
title	Title of news article
text	Body text of news article
normalizedText	Cleansed body text of news article
label	Ground truth of news article

Table 2: UND17 dataset format

¹www.opensources.co

3 Baselines

In this section, all the baseline systems evaluated on the UND17 dataset are presented. For each system, 4 different scenarios are evaluated. In the first scenario, only the feature vectors created from title are used for classification. In the second scenario, only the feature vectors created from the body are used for classification. In the third scenario, the title is added to the body to treat it as just another sentence of the body and then the resulting body is used to create a feature vector for classification. In the final scenario feature vectors separately created from title and body are concatenated and used for classification. We will call these scenarios as M_{title} , M_{body} , $M_{title+body}$ and, $M_{title||body}$ respectively.

3.1 TF-IDF

TF-IDF is a well known method for information retrieval and text mining which evaluates how important a word is to a document in a collection or corpus. In this method, first a vocabulary is created from all the text available in the corpus. For our case, the text from title and body are used for creating the vocabulary. Now using this vocabulary sparse representations (TF-IDF vectors) of each document are created from the words contained in the document. Gensim (Bird, Klein, and Loper 2009) and XGBoost’s (Chen and Guestrin 2016) gradient boosting classifier have been used for creating TF-IDF model from UND17 corpus. Cross-validation and evaluation metrics are computed using Scikit-learn (Pedregosa et al. 2011) toolkit.

3.2 Doc2Vec

Doc2Vec (Le and Mikolov 2014) is an unsupervised learning algorithm for learning fixed length representations of variable length textual inputs like sentences, paragraphs, and documents similar to word2vec (Mikolov et al. 2013) for fixed learning representation of words. Doc2Vec model tries to overcome the weaknesses of Bag-Of-Words models where the word ordering and semantics of words are lost. Gensim library is used for creating the Doc2Vec model and gradient boosting classifier for classification.

3.3 EmoInt

EmoInt (Duppada and Hiray 2017)² is a library which can extract lexical, stylistic and semantic features from a text body. These feature vectors extracted can then be used on any downstream NLP task. This is a good baseline to consider if affective features like sentiment and emotion are needed and useful for classification. This baseline is included and evaluated so as to cover the official baseline provided for this task. Gradient boosting classifier is again used for classification.

4 Hierarchical Attention Networks

Recently proposed Hierarchical Attention Networks (Yang et al. 2016) (HAN) for document classification has outperformed previous methods by a substantial margin in many classification tasks. The authors of HANs have attributed

the effectiveness of these networks to hierarchical network structure which reflects the structure of documents and multilevel attention applied on words and sentences separately. Visualizing attention layer weights on the word and sentence level show that network attends differentially based on their importance to document classification task at hand. This is particularly useful because it gives insights about decision making process for classification in contrast to other methods which act as black box optimization techniques.

In this paper, the default hierarchical attention network architecture is modified to encode the news articles effectively and results of experimented variants are presented. Before proceeding further, the terminology used to describe the model is explained here. Given a corpus which has N documents, each document i , $i \in [1, N]$ is assumed to have L_i sentences and each sentence j , $j \in [1, L_i]$ of document i has T_j words. Now a word of document i in sentence j can be represented as w_{kj} where $k \in [1, T_j]$, $j \in [1, L_i]$. Now each component of the architecture and its associated equations are introduced in bottom-up manner.

4.1 Word Encoder

Each word w_{kj} of sentence j in document i is first converted to a word vector using word embedding matrix W_e . At this step, either a bidirectional or unidirectional RNNs can be used to encode context into word annotations. In this paper Bidirectional gated recurrent units (GRUs) (Bahdanau, Cho, and Bengio 2014) are being used even though other types of RNNs like long short-term memory units (LSTMs) (Hochreiter and Schmidhuber 1997) can also be used. Bidirectional GRU contain a forward GRU and backward GRU which encodes the word context from both the directions. The equations for encoding words of sentence j of document i are presented below.

$$\begin{aligned}x_{kj} &= W_e w_{kj} \\ \vec{h}_{kj} &= GRU(x_{kj}), \overleftarrow{h}_{kj} = \overleftarrow{GRU}(x_{kj}) \\ h_{kj} &= [h_{kj}, \overleftarrow{h}_{kj}] \\ k &\in [1, T_j], j \in [1, L_i]\end{aligned}$$

4.2 Word Attention

Each word doesn’t contribute equally to the meaning of a sentence. This varied contribution of words can be modeled using attention mechanism (Bahdanau, Cho, and Bengio 2014). This way words which are important to the meaning of a sentence can be aggregated to form a sentence vector. The attention mechanism used here is as follows: the word representations obtained in the previous step h_{kj} are passed through one layer MLP to get a hidden representation u_{kj} . Now the importance of each word is modeled as the similarity of the hidden representation with word-level context vector u_w and are normalized using softmax. Finally, a sentence vector is generated as the sum of word annotations with attention generated normalized importance weights. The vector u_w can be thought of as a high-level representation of informative words and is jointly learned

²<https://github.com/SEERNET/EmoInt>

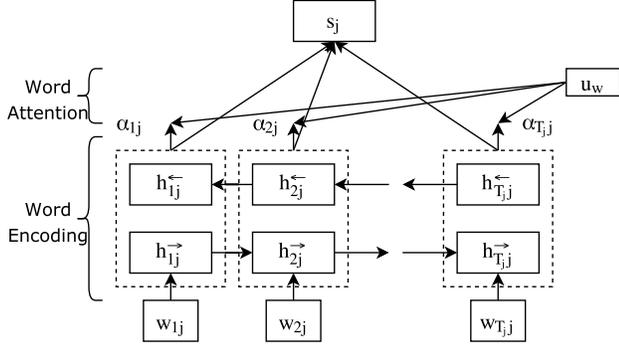


Figure 1: Sentence encoding schema from word vectors

with the model. Refer to figure 1 for the sentence encoding schema.

$$u_{kj} = \tanh(W_w h_{kj})$$

$$\alpha_{kj} = \frac{\exp(u_{kj}^T u_w)}{\sum_t \exp(u_{tj}^T u_w)}$$

$$s_j = \sum_t \alpha_{tj} h_{tj}$$

$$t, k \in [1, T_j], j \in [1, L_i]$$

4.3 Sentence Encoder

Similar to the word encoder, in section 4.1, sentence encoder has a forward GRU and backward GRU which creates sentence annotations by encoding sentence context from both sides of a document i i.e.

$$\vec{h}_j = G\vec{R}U(s_j), \bar{h}_j = G\bar{R}U(s_j)$$

$$h_j = [\vec{h}_j, \bar{h}_j], j \in [1, L_i]$$

4.4 Sentence Attention

Not all sentences are equally important or relevant for encoding and classifying a document. Similar to word level attention, in section 4.2, sentence level attention is employed using sentence-level context vector u_s to weigh the importance of individual sentences in a document. Refer to figure 2 for the document encoding schema.

$$u_j = \tanh(W_s h_j)$$

$$\alpha_j = \frac{\exp(u_j^T u_s)}{\sum_t \exp(u_t^T u_s)}$$

$$d_i = \sum_t \alpha_t h_t$$

$$i \in [1, N] \& t, j \in [1, L_i],$$

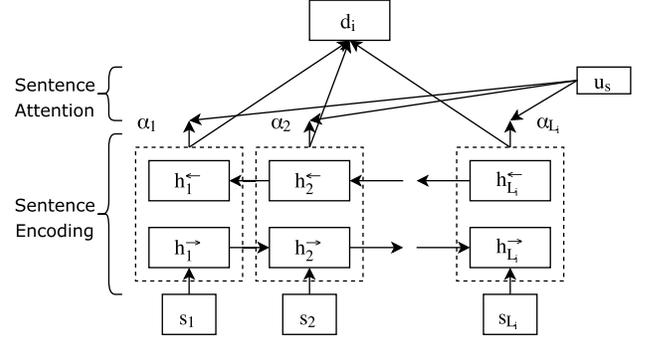


Figure 2: Document encoding schema from sentence vectors

4.5 News Article Encoder

Now that the terminology for describing HANs is established, here various schemes for encoding news articles is presented. As discussed earlier, title and body of news articles can be used in different ways for encoding an article.

In the first method HAN_{v1} , title and body are concatenated to form a document and HAN document encoder schema is used to get a fixed length representation of the news article. This method directly follows from the HAN system description above.

In the second method HAN_{v2} , title is encoded using sentence encoding schema (Figure 1) and body is encoded with HAN document encoding schema (Figure 2). Once title encoding s_t and body encoding d_b are obtained these are passed through bidirectional GRUs. This gives:

$$\vec{h}_t = G\vec{R}U(s_t), \bar{h}_t = G\bar{R}U(s_t)$$

$$h_t = [\vec{h}_t, \bar{h}_t]$$

$$\vec{h}_b = G\vec{R}U(d_b), \bar{h}_b = G\bar{R}U(d_b)$$

$$h_b = [\vec{h}_b, \bar{h}_b]$$

Article level attention for this method is described in section 4.6.

For the sake of completeness in the third method HAN_{v3} recently proposed HAN (Singhania, Fernandez, and Rao) variant for news article encoding is used.

4.6 News Article Attention

Let u_a be the article level context vector, similar to word level context vector u_w and sentence level context vector u_s which captures the importance of title and body for finding the reliability of the article. This vector is jointly learned with the model and is initialized randomly similar to previous context vectors, then article level attention directly follows from word and sentence level attentions. Refer Figure 3 for article encoding schema.

$$u_t = \tanh(W_a h_t), u_b = \tanh(W_a h_b)$$

$$\alpha_t = \frac{\exp(u_t^T u_a)}{\exp(u_t^T u_a) + \exp(u_b^T u_a)}$$

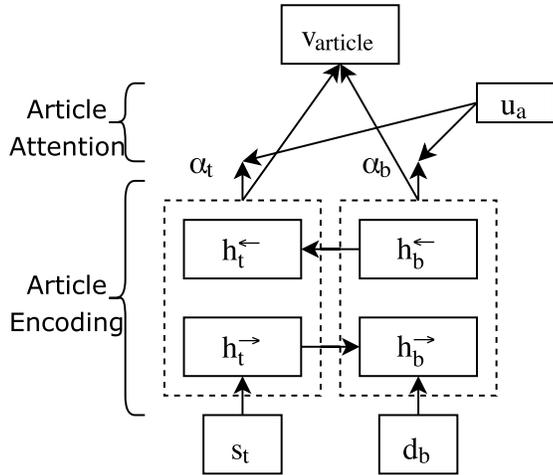


Figure 3: Article encoding schema from title and body vectors

$$\alpha_b = \frac{\exp(u_b^T u_a)}{\exp(u_t^T u_a) + \exp(u_b^T u_a)}$$

$$v_{article} = \alpha_t h_t + \alpha_b h_b$$

4.7 Unreliable News Classification

Once a news article is encoded to a vector $v_{article}$, it is passed to a single layer MLP with softmax activation for obtaining probability of the article being an unreliable one. Minimizing the categorical cross entropy loss is used as objective function for model training.

5 Experiments & Results

5.1 HAN System Parameters

The data was pre-processed before use to clean and standardize the article’s raw text. NLTK’s (Bird, Klein, and Loper 2009) word and sentence tokenizers are used to split words and sentences. Keras (Chollet and others 2015) deep learning library is used for building and training attention neural networks. GloVe embeddings (Pennington, Socher, and Manning 2014) of 100 dimensions trained on Wikipedia and Gigaword data is used for constructing word embedding matrix. The count of words per sentence and the number of sentences per document counts are calculated on training data and are presented in Table 3. The sentences and words were trimmed or padded to a sequence length of 64, which covers 95% of the data. A vocabulary of 65510 most occurring words from the training data is used for word embedding matrix of embedding layer in Keras. A dropout (Srivastava et al. 2014) of 0.5 is used after each attention layer for regularizing the network during training. Class weights are used to scale the training loss to mitigate the data imbalance problem. Finally, the models are trained using a batch size of 64 using Adam (Kingma and Ba 2014) optimizer to minimize categorical cross entropy loss.

Unreliable News Data 2017		
	Words/sentence	Sentences/document
Average	23	25
95 Percentile	64	55
99 Percentile	124	84

Table 3: Number of words per sentence and number of sentences per document distribution on training data.

	Precision	Recall	ROC-AUC
HAN_{v1}	0.742	0.78	0.944
HAN_{v2}	0.826	0.753	0.933
HAN_{v3}	0.876	0.692	0.931

Table 4: Evaluation metrics of hierarchical attention network variants.

5.2 Baseline Parameters

We used Gradient boosting classifier with default parameters for all the baselines. Baseline results are presented in Tables 5, 6, 7.

	Precision	Recall	ROC-AUC
M_{title}	0.726	0.09	0.642
M_{body}	0.597	0.229	0.766
$M_{title+body}$	0.621	0.249	0.772
$M_{title body}$	0.625	0.233	0.775

Table 5: Doc2Vec baseline results.

	Precision	Recall	ROC-AUC
M_{title}	0.732	0.249	0.753
M_{body}	0.706	0.507	0.857
$M_{title+body}$	0.703	0.501	0.853
$M_{title body}$	0.736	0.496	0.860

Table 6: EmoInt baseline results.

	Precision	Recall	ROC-AUC
M_{title}	0.866	0.279	0.759
M_{body}	0.861	0.451	0.877
$M_{title+body}$	0.895	0.503	0.892
$M_{title body}$	0.908	0.535	0.898

Table 7: TF-IDF baseline results.

6 Analysis

The hypothesis that titles play an important role in finding unreliable news is tested by running exhaustive experiments on the data. Similar studies have been done recently (Horne and Adali 2017) supporting the same. It can be observed from the in Tables 7, 6, and 5, training a model using just the title M_{title} gives good precision and poor recall. This means if an article is declared as unreliable by a model trained just using titles it is highly probable that it is unreliable, even

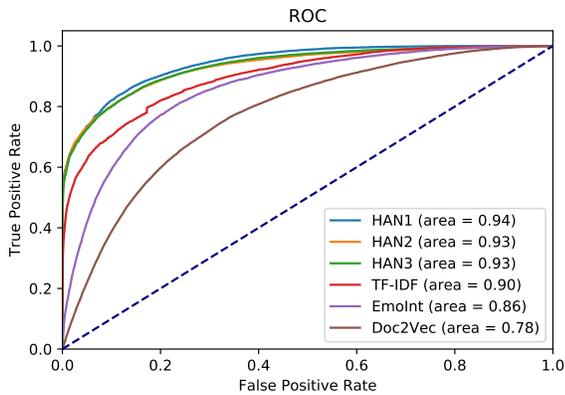


Figure 4: ROC plots of best performing model variants.

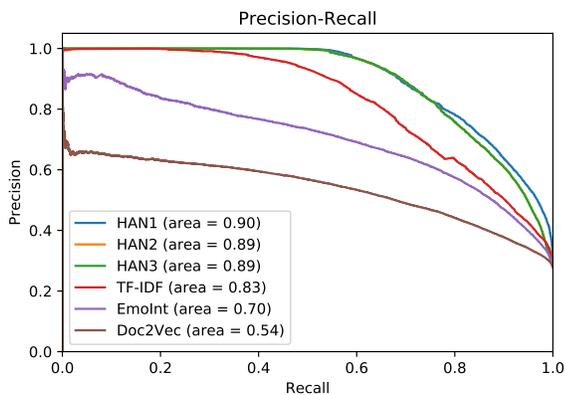


Figure 5: Precision-Recall plots of best performing model variants.

though it cannot find all the unreliable articles exhaustively. The precision values with models trained on the title are as good as the model trained on the body, if not better. Once it is established that the title of an article is important, the effective usage of the title and data for unreliable article classification is studied. For this a model $M_{title+body}$ where the title is concatenated to the body to treat it as another sentence of the article body and a model $M_{title||body}$ where separate features computed from title and body are used is trained. The results show that model $M_{title||body}$ outperforms other models.

In HANs the attention layer weights of variant 1 HAN_{v1} are visualized to understand how the decisions are made. For instance, Figure 6 shows a heat map of sentence level and word level attention weights where an unreliable article with title “Breaking! The Manchester ARENA BOMBER Has THIS in COMMON with ISLAMIC Terrorists” is used to plot word and sentence level attention weights. In Figure 6 for brevity only top 5 important sentences are shown in order. As you can see, the model assigns high weight to the flashy words like “Breaking !” and “ISLAMIC Terrorists” in the first two sentences. This is one of many cases where

0.416	breaking	!				
0.207	the	manchester	arena	bomber	has	
	this	in	common	with	islamic	terrorists
0.047	the	attacker	reportedly	carried	a	bomb
	in	a	backpack	and	may	have
	ignited	it	by	the	venue's	box
	office	as	attendees	began	to	stream
	out	of	the	ariana	grande	concert
	.					
0.059	2	explosions	occurred	which	contained	nails
	as	a	cruel	weapon	,	according
	to	a	cbs	report	.	
0.085	the	bbc	has	reported	that	great
	britain's	north	west	counter	terrorism	unit
	is	treating	the	incident	as	possible
	terrorism	.				

Figure 6: Attention Weights Visualization

the online news or social media jump to conclusions without knowing the full facts for various reasons like click baiting, getting website traffic etc.

7 Conclusion

In this paper, first traditional NLP algorithms are evaluated for unreliable news classification on UND17 dataset. Later variants of hierarchical attention networks (HANs) are presented for encoding news articles effectively by taking advantage of title and body of the article. HANs attention layer is visualized so as to reveal insights about how a classification decision is made. The results obtained are promising and encourages for further exploration of novel methods. The authors would like to thank AICS 2018 Organizing Committee for timely support.

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Bird, S.; Klein, E.; and Loper, E. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. O'Reilly Media, Inc.™.
- Chen, T., and Guestrin, C. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794. ACM.
- Chollet, F., et al. 2015. Keras. <https://github.com/fchollet/keras>.
- Duppada, V., and Hiray, S. 2017. Seernet at emoint-2017: Tweet emotion intensity estimator. In *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, 205–211.

0.05	rice	and	obama	conspired	to	wiretap
	trump	and	hide	the	evidence	
0.058	not	only	has	obama	been	caught
	illegally	wiretapping	,	he	even	had
	the	gall	to	hide	the	evidence
	.					
0.074	however	,	you	should	be	aware
	that	under	the	presidential	records	act
	,	presidential	records	remain	closed	to
	the	public	for	five	years	after
	an	administration	has	left	office	
	.					
0.119	unbelievable	!				
0.08	what	compounds	the	seriousness	is	
	that	the	place	he	was	caught
	hiding	his	dirty	secrets	is	one
	huge	lie	,	which	proves	he
	is	trying	to	avoid	prosecution	.

Figure 7: Attention Weights Visualization

Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.

Horne, B. D., and Adali, S. 2017. This just in: Fake news packs a lot in title, uses simpler, repetitive content in text body, more similar to satire than real news. *arXiv preprint arXiv:1703.09398*.

Kingma, D., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Le, Q., and Mikolov, T. 2014. Distributed representations of sentences and documents. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, 1188–1196.

Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, 3111–3119.

Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12:2825–2830.

Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.

Singhania, S.; Fernandez, N.; and Rao, S. 3han: A deep neural network for fake news detection.

Srivastava, N.; Hinton, G. E.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to pre-

vent neural networks from overfitting. *Journal of machine learning research* 15(1):1929–1958.

2018.

Yang, Z.; Yang, D.; Dyer, C.; He, X.; Smola, A. J.; and Hovy, E. H. 2016. Hierarchical attention networks for document classification. In *HLT-NAACL*, 1480–1489.