ReINTEL: A Multimodal Data Challenge for Responsible Information Identification on Social Network Sites

Duc-Trong Le[1], **Xuan-Son Vu**[2], **Nhu-Dung To**[3], **Huu-Quang Nguyen**[4], **Thuy-Trinh Nguyen**[4], **Linh Le**[4], **Anh-Tuan Nguyen**[4], **Minh-Duc Hoang**[4], **Nghia Le**[4] **Huyen Nguyen**[5], **Hoang D. Nguyen**[6]

[1]University of Engineering and Technology, Vietnam National University, Vietnam.
trongld@vnu.edu.vn
[2]Dept. of Computing Science, Umeå University, Sweden.
sonvx@cs.umu.se
[3]School of Computer Science, University of Sydney, Australia.
duto3894@uni.sydney.edu.au
[4]ReML.AI - Reliable Machine Learning Lab, International.
{quang, trinh, linh, tuan, duc, nghia}@reml.ai
[5]Hanoi University of Science, Vietnam National University, Vietnam.
huyenntm@hus.edu.vn
[6]School of Computing Science, University of Glasgow, Singapore.
harry.nguyen@glasgow.ac.uk

Abstract

This paper reports on the ReINTEL Shared Task for Responsible Information Identification on social network sites, which is hosted at the seventh annual workshop on Vietnamese Language and Speech Processing (VLSP 2020). Given a piece of news with respective textual, visual content and metadata, participants are required to classify whether the news is *'reliable'* or *'unreliable'*. In order to generate a fair benchmark, we introduce a novel human-annotated dataset of over 10,000 news collected from a social network in Vietnam. All models will be evaluated in terms of AUC-ROC score, a typical evaluation metric for classification. The competition was run on the Codalab platform. Within two months, the challenge has attracted over 60 participants and recorded nearly 1,000 submission entries.

1 Introduction

This challenge aims at identifying the reliability of information shared on social network sites (SNSs). With the blazing-fast spurt of SNSs (e.g. Facebook, Zalo and Lotus), there are approximately 65 million Vietnamese users on board with the annual growth of 2.7 million in the recent year, as reported by the Digital 2020 [1]. SNSs have become widely accessible for users to not only connect friends but also freely create and share diverse content (Shu et al., 2017; Zhou et al., 2019). A number of users, however, has exploited these social platforms to distribute fake news and unreliable information to fulfill their personal or political purposes (e.g. US election 2016 (Allcott and Gentzkow, 2017)). It is not easy for other ordinary users to realize the unreliability, hence, they keep spreading the fake content to their friends. The problem becomes more seriously once the unreliable post becomes popular and gains belief among the community. Therefore, it raises an urgent need for detecting whether a piece of news on SNSs is reliable or not. This task has gained significant attention recently (Ruchansky et al., 2017; Shu et al., 2019a,b; Yang et al., 2019).

The shared task focuses on the responsible (i.e. reliable) information identification on Vietnamese SNSs, referred to as ReINTEL. It is a part of the 7th annual workshop on Vietnamese Language and Speech Processing, VLSP 2020[2] for short. As a binary classification task, participants are required to propose models to determine the reliability of SNS posts based on their content, image and metadata information (e.g. number of likes, shares, and comments). The shared task consists of three phases namely *Warm up, Public Test, Private Test*, which is hosted on Codalab from October 21st, 2020 to November 30th, 2020. In summary, there are around 1000 submissions created by 8 teams and over 60 participants during the challenge period.

[1] https://wearesocial.com/digital-2020

[2] https://vlsp.org.vn/vlsp2020

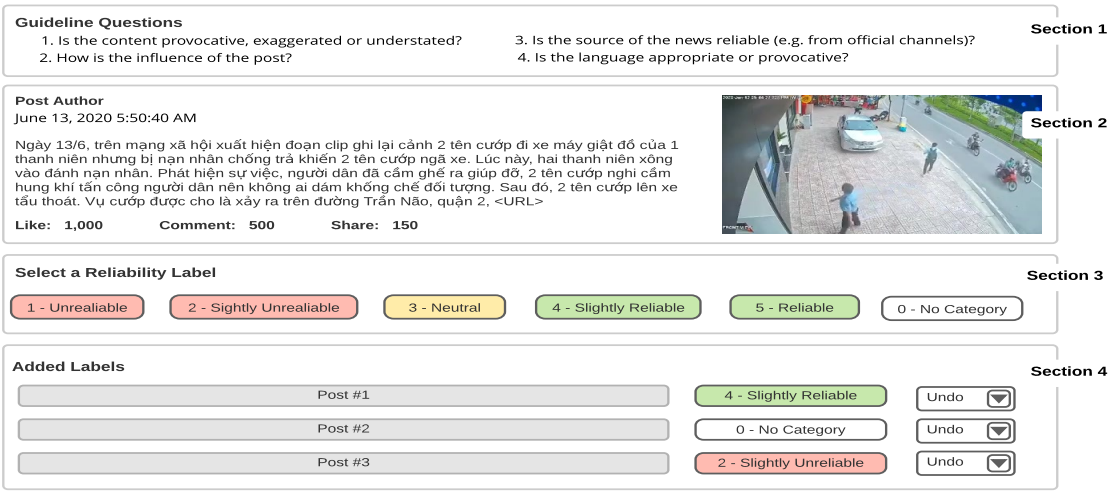

Figure 1: Data Annotation Tool

As our first contribution, this shared task provides an evaluation framework for the reliable information detection task, where participants could leverage and compare their innovative models on the same dataset. Their knowledge contribution may help improve safety on online social platforms. Another valuable contribution is the introduction of a novel dataset for the reliable information detection task. The dataset is built based on a fair human annotation of over 10,000 news from SNSs in Vietnam. We hope this dataset will be a useful benchmark for further research. In this shared task, AUC-ROC is utilized as the primary evaluation metric.

The remainder of the paper is organized as follows. The next section describes the data collection and annotation methodologies. Subsequently, the shared task description and evaluation are summarized in Section 3. In Section 4, we discusses the potentials of language and vision transfer learning for the detection task. Section 5 describes the competition, approaches and respective results. Finally, Section 6 concludes the paper by suggesting potential applications for future studies and challenges.

2 The ReINTEL 2020 Dataset

2.1 Data Collection

We collect the data for two months from August to October 2020. There are two main sources of the data: SNSs and Vietnamese newspapers. As for the former source, public social media posts are retrieved from news groups and key opinion leaders (KOLs). Many fake news, however, has been flagged and removed from the social networking sites since the enforcement of Vietnamese cybersecurity law in 2019 (Son, 2018). Therefore, to include the deleted fake news, we gather newspaper articles reporting these posts and recreate their content.

All the collected data were originally posted in the period of March - June 2020. During this time, Vietnam was facing a second wave of Covid-19 with a drastic increase from 20 to 355 cases (WHO, 2020). The spread of Covid-19 results in an 'infodemic' in which misleading information is disseminated rapidly especially on social media (Hou et al., 2020; Huynh et al., 2020). Hence, this period is a potential source of fake news. Besides Covid-19, the items in our dataset cover a wide range of domains including entertainment, sport, finance and healthcare. The result of the data collection stage is 10,007 items that are prepared for the annotation process.

2.2 Data Annotation

2.2.1 Annotator and Training

We recruit 23 human annotators to participate in the annotation process. The annotators receive one week training to identify fact-related posts and how to evaluate the reliability of the post based on primary features including the news source, its image and content.

2.2.2 Annotation Tool

Figure 1 demonstrates the annotation tool interface, which is designed to support quick and easy annotation. The first section contains guideline questions

to remind the annotators of the labeling criterion including the news source credibility, the language appropriateness and fact accuracy. The second section is the post content, image and influence (i.e. number of likes, comments and shares). In Section 3, the annotators select a Reliability score for the post. There is a 5-point reliability Likert scale for fact-based posts with the following labels: 1 - Unreliable, 2 - Slightly unreliable, 3 - Neutral, 4 - Slightly reliable, 5 - Reliable. On the other hand, if the post is opinion-based and does not contain facts, the annotators should select label '0 - No category' instead.

The last section is a list of labeled items for the annotators to review and update their decision, if necessary, using the 'Undo' button.

2.2.3 Annotation Process

The annotation process is conducted from 9th to 19th October 2020. The annotators are divided into three groups to annotate 10,007 items independently. Therefore, each item will be annotated three times by different annotators.

Once the annotators finish 30,021 annotations (i.e. 10,007 items annotated three times), we filter and summarise the result based on majority vote basis. Firstly, we combine labels of the same essence: Category 1 and 2 (Unreliable and Sightly unreliable) and Category 4 and 5 (Slightly reliable and Reliable). After merging the categories, we select the majority votes to be the final labels. If the majority vote is 1 or 2, the final label should be 1 - Unreliable. If the majority vote is 4 or 5, the final label should be 0 - Reliable. When the majority vote is 3 - Neutral, we finalise using ground truth labels. Lastly, if the majority agrees that the post is not fact-based (i.e. 0 - No Category), we remove it from the set.

For items with no majority votes (i.e. three annotators have different opinions), we follow an alternate procedure. If the ground truth label is 1 - unreliable, the final label should be 1. On the other hand, if the ground truth label is 0 - reliable, we double check to separate reliable news from opinion-based items. The process is illustrated in Figure 2.

2.2.4 Content Filtering

Once the annotation process is finished, data needs to go through the last step before being published for the competition – the content filtering. In this step, we manually check to ensure that data, including both text and image, published for the competition:

1. Does not violate any law, statue, ordinance, or regulation
2. Will not give rise to any claims of invasion of privacy or publicity
3. Does not contain, depict, include or involve any of the following:
    - Political or religious views or other such ideologies
    - Explicit or graphic sexual activity
    - Vulgar or offensive language and/or symbols or content
    - Personal information of individuals such as names, telephone numbers, and addresses
    - Other forms of ethical violations

3 The ReINTEL 2020 Challenge

3.1 Dataset Splitting

Data splitting for data challenge is a difficult process in order to avoid evidence ambiguity and concept drifting which are the main cause of unstable ranking issue in data challenges.

In this competition, we apply RDS (Nguyen et al., 2020) to split ReINTEL data into three sets including public train, validation, and private test sets. It is worth to mention that, RDS is a method to approximate optimum sampling for model diversification with ensemble rewarding to attain maximal machine learning potentials. It has a novel stochastic choice rewarding is developed as a viable mechanism for injecting model diversity in reinforcement learning.

3.1.1 Baselines

To apply RDS (Nguyen et al., 2020) for the data splitting process, it requires to have baseline learners to obtain rewards for the reinforced process. It is recommended to choose representative baseline learners, to let the reinforced learner better capture different learning behaviors. The use of these baseline learners is important since each learner will behave differently depending on the patterns contained in the target data. As a result, RDS helps to increase the diversity of the data samples in different sets. Here we employ three models to classify reliable news using textual features as follows:

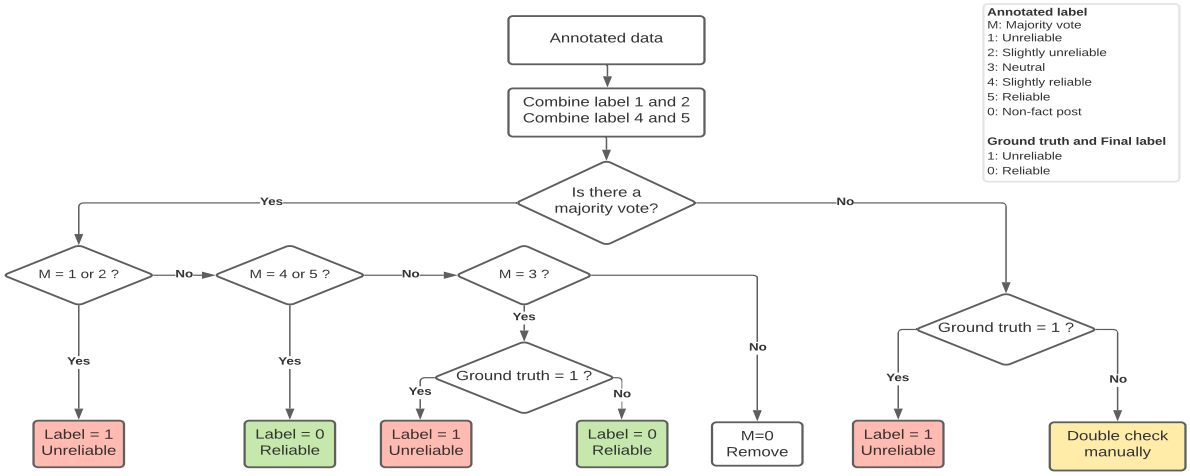

Figure 2: Data Annotation Process

- **Bi-LSTM** (Schuster and Paliwal, 1997) is a bi-directional LSTM model. It has two LSTMs in which, one LSTM takes input sequence in a forward direction, and another LSTM takes input sequence in a backward direction. The use of Bi-LSTM architecture helps to increase the amount of information available to the network, to gain better performance in most of sequence related tasks. Bi-LSTM network is a standard baseline for most of text classification tasks.

- **CNN-Text** (Kim, 2014) is the use of CNN (LeCun et al., 1989) network on word embeddings to perform the classification tasks. The simple architecture outperformed all other models at the publication time.

- **EasyEnsemble** (Liu et al., 2009) is used to represent a tradition approach in dealing with im-balanced dataset. For the vectorization, we trained a Sent2Vec (Pagliardini et al., 2018) using the combined 1GB texts of Vietnamese Wikipedia data (Vu et al., 2019) and 19 GB texts of Vuong (2018).

3.1.2 Learning Dynamics

To disentangle dataset shift and evidence ambiguity of the data splitting strategy, we apply RDS stochastic choice reward mechanism (Nguyen et al., 2020) to create public training, public- and private testing sets. Figure 3 illustrates the learning dynamic towards the goal.

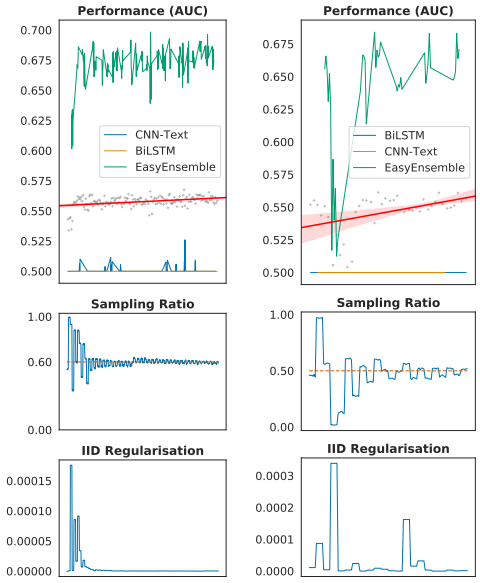

(a) Training Set (b) Val & Test Sets

Figure 3: Learning Dynamics for splitting data into 3 sets (public training, public testing, and private testing) using RDS Stochastic Choice Reward Mechanism (Nguyen et al., 2020).

4 Transfer Learning

Knowledge transfer has been found to be essential when it comes to downstream tasks with new datasets. If this transfer process is done correctly, it would greatly improve the performance of learning. Since ReINTEL challenge is a multimodal challenge, both visual based knowledge transfer and language based knowledge transfer are used by different teams.

To be fair between participants, we required all teams to register for the use of pre-trained models.

Model	Language	Vision	Description
Word2VecVN (Vu, 2016)	x		Trained on 7GB texts of Vietnamese news
FastText (Vietnamese version) (Joulin et al., 2016)	x		Trained on Vietnamese texts of the CommonCrawl corpus
ETNLP (Vu et al., 2019)	x		Trained on 1GB texts of Vietnamese Wikipedia
PhoBERT (Nguyen and Nguyen, 2020)	x		Trained on 20GB texts of both Vietnamese news and Vietnamese Wikipedia
Bert4News (Nha, 2020)	x		Trained on more than 20GB texts of Vietnamese news
vElectra and ViBERT (The et al., 2020)	x		vElectra was trained on 10GB texts, whereas ViBERT was trained on 60GB texts of Vietnamese news
VGG16 (Simonyan and Zisserman, 2015)		x	Trained on ImageNet (Deng et al., 2009)
YOLO (Redmon et al., 2015)		x	Trained on ImageNet (Deng et al., 2009)
EfficientNet B7 (Tan and Le, 2019)		x	Trained on ImageNet (Deng et al., 2009)

Table 1: List of pre-trained models registered by all participants of ReINTEL challenge in 2020.

Table 1 lists all pre-trained language and vision models registered by all participants.

4.1 Language Transfer Learning

For natural language processing tasks in Vietnamese, there have been many pre-trained language models are available. In 2016, Vu (2016) introduced the first monolingual pre-trained models for Vietnamese based on Word2Vec (Mikolov et al., 2013). The use of pre-trained Word2VecVN models was proved to be useful in various tasks, such as the name entity recognition task (Vu et al., 2018). In 2019, Vu et al. (2019) introduced the use of multiple pre-trained language models to achieve new state-of-the-art results in the name entity recognition task (Nguyen et al., 2019). Up to date, there have been many other new monolingual language models for Vietnamese are available such as PhoBERT (Nguyen and Nguyen, 2020), vElectra and ViBERT (The et al., 2020).

4.2 Vision Transfer Learning

Different from language models, visual models are normally universal and existing pre-trained models can be directly applied in most of image processing tasks. For the use of visual features, there is only one team using multimodal features among top 6 teams of the leader board. This team, in fact, achieved the 1^{st} rank on the public test (see Table 3); but they did not get the same rank on the private test. This hints that the reliability of news mainly depends on content of news and other meta information, such as number of likes on social networks. Moreover, it is yet to be explored to capture the reliability of news using both vision and language information.

4.3 Language and Vision Transfer Learning

The use of both language and vision transfer learning is important for multimodal tasks. This line of research has attracted much attention with various new language-vision models, such as VilBERT (Lu et al., 2019), 12-in-1 (Lu et al., 2020). No participants employ into this approach in the ReINTEL challenge due to the lack of language and vision pre-trained models in Vietnamese. Moreover, it is required to have extensive computer resources for applying this approach in a data challenge. In the future, we expect to see more research done in this direction because both images and texts are essential to SNS issues.

No	Attribute	Description
1	id	Unique ID of each post
2	user_name	Anomynized post owner's identity
3	post_message	Text content of the post
4	timestamp_post	The time when the post is uploaded
5	num_like_post	Number of likes that the post received
6	num_comment_post	Number of comments that the post received
7	num_share_post	Number of shares that the post received
8	image	The image uploaded with the post
9	label	Manually annotated label indicating the reliability of the post 1: Unreliable 0: Reliable

Table 2: Data attributes

5 Results

5.1 Data Format

Each instance includes 8 main attributes with/without a binary target label. Table 2 summarizes the key features of each attribute.

5.2 Training/Testing Data

The challenge provides approximately 8,000 training examples with the respective target labels. The testing set consists of 2,000 examples without labels.

5.3 Result Submission

Participants must submit the result in the same order as the testing set in the following format:

```
id1, label probability 1
Id2, label probability 2
...
```

5.4 Evaluation Metric

The challenge task is evaluated based on Area Under the Receiver Operating Characteristic Curve (AUC-ROC), which is a typical metric for classification tasks. Let us denote X as a *continuous random variable* that measures the 'classification' score of a given a news. As a binary classification task, this news could be classified as *"unreliable"* if X is greater than a threshold parameter T, and *"reliable"* otherwise. We denote $f_1(x), f_0(x)$ as probability density functions that the news belongs to *"unreliable"* and *"reliable"* respectively, hence the true positive rate $TPR(T)$ and the false positive rate $FPR(T)$ are computed as follows:

$$TPR(T) = \int_T^{\infty} f_1(x)dx \quad (1)$$

$$FPR(T) = \int_T^{\infty} f_0(x)dx \quad (2)$$

and the AUC-ROC score is computed as:

$$AUC\_ROC = \int_{-\infty}^{\infty} TPR(T)FPR'(T)dT \quad (3)$$

Here, submissions are evaluated with ground-truth labels using the *scikit-learn*'s implementation [3].

5.5 Participation

During the course two months of the competition, 61 participants sign up for the challenge. 30% of the participants compete in groups of 2 (6 teams) and 4 members (2 teams). 19 participants sign our corpus usages agreement.

From top 8 of the Private test leaderboard, 6 teams/participants submit their technical reports that demonstrate their strategies and findings from the challenge. The summary of the competition participation can be seen in Table 4.

5.6 Outcomes

In total, 657 successful entries were recorded. The highest results of the Public test and Private test phase were 0.9427 and 0.9521 respectively. Key descriptive statistics of the results in each phase is illustrated in Table 5.

[3] https://scikit-learn.org/stable/modules/generated/sklearn.metrics.roc_auc_score.html

Table 3: Top 6 teams on public-test and private-test with submitted papers and their final approaches. The rank is based on the ROC-AUC scores on the private-test.

#	Team	ROC-AUC		Final Approach	Ensemble?	Multimodal?
		Public-test	Private-test			
1	Kurtosis	0.9399	**0.9521**	TF-IDF + SVD; Emb + SVD; NB, LightGBM, CatBoost	Yes	No
2	NLP_BK	0.9360	0.9513	Bert4News + phoBERT + XLM + MetaFeatures	Yes	No
3	SunBear	0.9418	0.9462	RoBerta + MLP	Yes	No
4	uit_kt	-	0.9452	phoBERT + Bert4News	Yes	No
5	Toyo-Aime	**0.9427**	0.9449	CNN + Bert + Fully connected	Yes	Yes
6	ZaloTeam	-	0.9378	viBERT + viELECTRA + phoBERT	Yes	No

Metric	Value
Number of participants	61
Number of teams	8
Number of signed agreements	19
Number of submitted papers	6

Table 4: Participation summary

	Public Test	**Private Test**	**Overall**
Total Entries	571	86	657
Highest ROC	0.9427	0.9521	0.9474
Mean ROC	0.8463	0.8942	0.8703
Std. ROC	0.1215	0.1022	0.1119

Table 5: Results summary

6 Conclusion

The rise of misleading information on social media platforms has triggered the need for fact-checking and fake news detection. Therefore, the reliability of news has become a critical question in the modern age. In this paper, we introduce a novel dataset of nearly 10,000 SNSs entries with reliability labels. The dataset covers a great variety of topics ranging from healthcare to entertainment and economics. The annotation and validation process are presented in details with several filtering rounds. With both linguistic and visual features, we believe that the corpus is suitable for future research on fake news detection and news distributor behaviours using NLP and computer vision techniques. In Vietnam, where datasets on SNSs are scarce, our corpus will serve as a reliable material for other research.

Acknowledgment

The authors would like to thank the InfoRE company for the data contribution, the ReML-AI research group[4] for the data contribution and financial support, and the twenty three annotators for their hard work to support the shared task. Without their support, the task would not have been possible.

References

Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives*, 31(2):211–36.

J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.

Zhiyuan Hou, Fanxing Du, Hao Jiang, Xinyu Zhou, and Leesa Lin. 2020. Assessment of public attention, risk perception, emotional and behavioural responses to the covid-19 outbreak: social media surveillance in china. *Risk Perception, Emotional and Behavioural Responses to the COVID-19 Outbreak: Social Media Surveillance in China (3/6/2020)*.

Toan Luu Huynh et al. 2020. The covid-19 risk perception: A survey on socioeconomics and media attention. *Econ. Bull*, 40(1):758–764.

Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, Hérve Jégou, and Tomas Mikolov. 2016. Fasttext.zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.

Yoon Kim. 2014. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar. Association for Computational Linguistics.

[4] https://reml.ai

Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. 1989. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551.

X. Liu, J. Wu, and Z. Zhou. 2009. Exploratory undersampling for class-imbalance learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(2):539–550.

Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *CoRR*, abs/1908.02265.

Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 2020. 12-in-1: Multi-task vision and language representation learning. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space.

Dat Quoc Nguyen and Anh Tuan Nguyen. 2020. PhoBERT: Pre-trained language models for Vietnamese. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1037–1042.

Hoang D. Nguyen, Xuan-Son Vu, Quoc-Tuan Truong, and Duc-Trong Le. 2020. Reinforced data sampling for model diversification.

Huyen Nguyen, Quyen Ngo, Luong Vu, Vu Tran, and Hien Nguyen. 2019. Vlsp shared task: Named entity recognition. *Journal of Computer Science and Cybernetics*, 34(4):283–294.

Nguyen Van Nha. 2020. Pre-trained bert4news. https://github.com/bino282/bert4news.

Matteo Pagliardini, Prakhar Gupta, and Martin Jaggi. 2018. Unsupervised Learning of Sentence Embeddings using Compositional n-Gram Features. In *NAACL 2018 - Conference of the North American Chapter of the Association for Computational Linguistics*.

Joseph Redmon, Santosh Kumar Divvala, Ross B. Girshick, and Ali Farhadi. 2015. You only look once: Unified, real-time object detection. *CoRR*, abs/1506.02640.

Natali Ruchansky, Sungyong Seo, and Yan Liu. 2017. Csi: A hybrid deep model for fake news detection. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, CIKM '17, page 797–806, New York, NY, USA. Association for Computing Machinery.

M. Schuster and K. K. Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11):2673–2681.

Kai Shu, Limeng Cui, Suhang Wang, Dongwon Lee, and Huan Liu. 2019a. defend: Explainable fake news detection. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 395–405.

Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake news detection on social media: A data mining perspective. *ACM SIGKDD explorations newsletter*, 19(1):22–36.

Kai Shu, Suhang Wang, and Huan Liu. 2019b. Beyond news contents: The role of social context for fake news detection. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 312–320.

Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition.

Tuan Son. 2018. Vietnam passes cyber security law.

Mingxing Tan and Quoc V. Le. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. *CoRR*, abs/1905.11946.

Viet Bui The, Oanh Tran Thi, and Phuong Le-Hong. 2020. Improving sequence tagging for vietnamese text using transformer-based neural models.

Thanh Vu, Dat Quoc Nguyen, Dai Quoc Nguyen, Mark Dras, and Mark Johnson. 2018. Vncorenlp: A vietnamese natural language processing toolkit. In *Proceedings of the 2018 NAACL: Demonstrations*, pages 56–60, New Orleans, Louisiana. Association for Computational Linguistics.

Xuan-Son Vu. 2016. Pre-trained word2vec models for vietnamese. https://github.com/sonvx/word2vecVN.

Xuan-Son Vu, Thanh Vu, Son N. Tran, and Lili Jiang. 2019. Etnlp: A visual-aided systematic approach to select pre-trained embeddings for a downstream task. *In: Proceedings of the International Conference Recent Advances in Natural Language Processing (RANLP)*.

Quoc Binh Vuong. 2018. Vietnamese news corpus. https://github.com/binhvq/news-corpus.

WHO. 2020. Who coronavirus disease (covid-19) dashboard.

Shuo Yang, Kai Shu, Suhang Wang, Renjie Gu, Fan Wu, and Huan Liu. 2019. Unsupervised fake news detection on social media: A generative approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5644–5651.

Xinyi Zhou, Reza Zafarani, Kai Shu, and Huan Liu. 2019. Fake news: Fundamental theories, detection strategies and challenges. In *Proceedings of the twelfth ACM international conference on web search and data mining*, pages 836–837.