

Benchmarking for Public Health Surveillance tasks on Social Media with a Domain-Specific Pretrained Language Model

Usman Naseem¹, Byoung Chan Lee², Matloob Khushi¹, Jinman Kim¹, Adam G. Dunn³

¹School of Computer Science, University of Sydney, Australia

²School of Medicine, University of Sydney, Australia

³School of Medical Sciences, University of Sydney, Australia

blee9781@uni.sydney.edu.au

{usman.naseem,matloob.khushi,jinman.kim,adam.dunn}@sydney.edu.au

Abstract

A user-generated text on social media enables health workers to keep track of information, identify possible outbreaks, forecast disease trends, monitor emergency cases, and ascertain disease awareness and response to official health correspondence. This exchange of health information on social media has been regarded as an attempt to enhance public health surveillance (PHS). Despite its potential, the technology is still in its early stages and is not ready for widespread application. Advancements in pretrained language models (PLMs) have facilitated the development of several domain-specific PLMs and a variety of downstream applications. However, there are no PLMs for social media tasks involving PHS. We present and release PHS-BERT, a transformer-based PLM, to identify tasks related to public health surveillance on social media. We compared and benchmarked the performance of PHS-BERT on 25 datasets from different social medial platforms related to 7 different PHS tasks. Compared with existing PLMs that are mainly evaluated on limited tasks, PHS-BERT achieved state-of-the-art performance on all 25 tested datasets, showing that our PLM is robust and generalizable in the common PHS tasks. By making PHS-BERT available¹, we aim to facilitate the community to reduce the computational cost and introduce new baselines for future works across various PHS-related tasks.

1 Introduction

Public health surveillance (PHS) is defined by the World Health Organization² as the ongoing, systematic collection, assessment, and understanding of health-related required information for the planning, implementation, and assessment of healthcare (Aiello et al., 2020). PHS aims to design and

assist interventions; it acts as a primary warning system in health emergencies (epidemics, i.e., acute events), it reports and records public health interventions (i.e., monitoring health), and it observes and explains the epidemiology of health issues, allowing for the prioritization of necessary details for health policy formulation (i.e., targeting chronic events). Traditional PHS systems are often limited by the time required to collect data, restricting the quick or even instantaneous identification of outbreaks (Hope et al., 2006).

Social media is growingly being used for public health purposes and can disseminate disease risks and interventions and promote wellness and healthcare policy. Social media data provides an abundant source of timely data that can be used for various public health applications, including surveillance, sentiment analysis, health communication, and analyzing the history of a disease, injury, or promote health. Systematic reviews of studies that examine personal health experiences shared online reveal the breadth of application domains, which include infectious diseases and outbreaks (Charles-Smith et al., 2015), illicit drug use (Kazemi et al., 2017), and pharmacovigilance support (Golder et al., 2015). These applied health studies are motivated by their potential in supporting PHS, augmenting adverse event reporting, and as the basis of public health interventions (Dunn et al., 2018).

The use of deep learning in natural language processing (NLP) has advanced the development of pretrained language models (PLMs) that can be used for a wide range of tasks in PHS. However, directly applying the state-of-the-art (SOTA) PLMs such as Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2019), and its variants (Liu et al., 2019; Lan et al., 2019; Sanh et al., 2019; Naseem et al., 2021c) that are trained on general domain corpus (e.g., Bookcorpus, Wikipedia, etc.) may yield poor per-

¹<https://huggingface.co/publichealthsurveillance/PHS-BERT>

²<https://www.euro.who.int/en/health-topics/Health-systems/public-health-services>

performances on domain-specific tasks. To address this limitation, several domain-specific PLMs have been presented. Some of the well-known in the biomedical field include the following: biomedical BERT (BioBERT) (Lee et al., 2019) and biomedical A Lite BERT (BioALBERT) (Naseem et al., 2020, 2021a). Recently, other domain-specific LMs such as BERTweet (Nguyen et al., 2020) for 3 downstream tasks, i.e., part-of-speech tagging, named-entity-recognition, and text classification and COVID Twitter BERT (CT-BERT) (Müller et al., 2020) for 5 text classification tasks have been trained on datasets from Twitter.

Despite the number of PLMs that have been released, none have been produced specifically for PHS from online text. Furthermore, all these LMs were evaluated with the selected dataset, and therefore their generalizability is unproven. To benchmark and fill the gap, we present PHS-BERT, a new domain-specific contextual PLM trained and fine-tuned to achieve benchmark performance on various PHS tasks on social media. PHS-BERT is trained on a health-related corpus collected from user-generated content. Our work is the first large-scale study to train, release and test a domain-specific PLM for PHS tasks on social media. We demonstrated that PHS-BERT outperforms other SOTA PLMs on 25 datasets from different social media platforms related to 7 different PHS tasks, showing that PHS is robust and generalizable.

2 Related Work

2.1 Pretrained Language Models

Transformer-based PLMs such as BERT (Devlin et al., 2019) and its variants (Liu et al., 2019; Lan et al., 2019) have altered the landscape of research in NLP domain. These PLMs are trained on a huge corpus but may not provide a good representation of specific domains (Müller et al., 2020). To improve the performance in domain-specific tasks, various domain-specific PLMs have been presented. Some of the famous in the biomedical domain are BioBERT (Lee et al., 2019) and BioALBERT (Naseem et al., 2020). Recently, for tasks on social media-specific, other PLMs such as BERTweet (Nguyen et al., 2020), COVID Twitter BERT (CT-BERT) (Müller et al., 2020) have been trained on datasets from Twitter. For various downstream tasks, these domain-specific PLMs were demonstrated to be effective alternatives for PLMs trained on a general corpus for a variety of down-

stream tasks (Müller et al., 2020). The assumption is that the LMs trained on the user-generated text on Twitter can handle the short and unstructured text in tweets. Despite this progress, their generalizability is unproven, and there is no PLM for public health surveillance using social media.

2.2 NLP for Public Health Surveillance

The use of social media in conjunction with advances in NLP for PHS tasks is a growing area of study (Paul and Dredze, 2017). NLP can assist researchers in the surveillance of mental disorders, such as identifying depression diagnosis, assessing suicide risk and stress identification, vaccine hesitancy and refusal, identifying common health-related misconceptions, sentiment analysis, and the health-related behaviors they support (Naseem et al., 2022a,b).

Rao et al. (2020) presented a hierarchical method that used BERT with attention-based BiGRU and achieved competitive performance for depression detection. For vaccine-related sentiment classification, Zhang et al. (2020) classified tweet-level HPV vaccine sentiment using three transfer learning techniques (ELMo, GPT, and BERT) and found that a finely tuned BERT produced the best results. Biddle et al. (2020) presented a method (BiLSTM-Senti) that leveraged contextual word embeddings (BERT) with word-level sentiment to improve performance. Naseem et al. (2021b) presented a model that uses domain-specific LM and captures commonsense knowledge into a context-aware bidirectional gated recurrent network. Sawhney et al. (2021) presented an ordinal hierarchical attention model for Suicide Risk Assessment where text embeddings obtained by Longformer were fed to BiLSTM with attention and ordinal loss as an objective function. However, there is no PLM trained on health-related text collected from social media that directly benefit the applications related to PHS.

3 Method

PHS-BERT has the same architecture as BERT. Fig. 1 illustrates an overview of pretraining, fine-tuning, and datasets used in this study. We describe BERT and then the pretraining and fine-tuning process employed in PHS-BERT.

3.1 BERT

PHS-BERT has the same architecture as BERT. BERT was trained on 2 tasks: mask language mod-

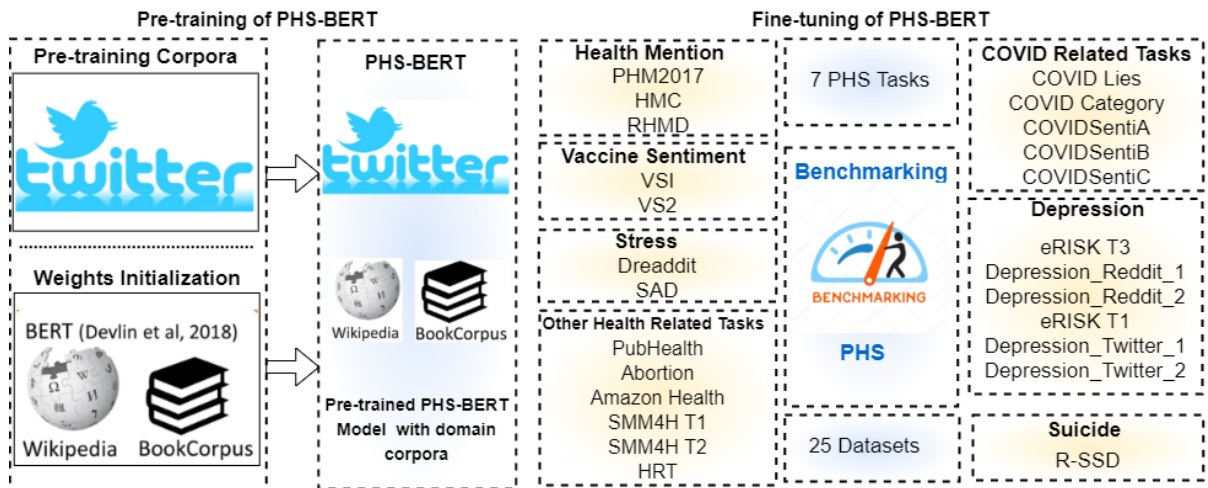


Figure 1: An overview of pretraining, fine-tuning, and the various tasks and datasets used in PHS benchmarking

eling (MLM) (15% of tokens were masked and next sentence prediction (NSP) (Given the first sentence, BERT was trained to predict whether a selected next sentence was likely or not). BERT is pretrained on Wikipedia and BooksCorpus and needs task-specific fine-tuning. Pretrained BERT models include $BERT_{Base}$ (12 layers, 12 attention heads, and 110 million parameters), as well as $BERT_{Large}$ (24 layers, 16 attention heads, and 340 million parameters).

3.2 Pretraining of PHS-BERT

We followed the standard pretraining protocols of BERT and initialized PHS-BERT with weights from BERT during the training phase instead of training from scratch and used the uncased version of the BERT model.

PHS-BERT is the first domain-specific LM for tasks related to PHS and is trained on a corpus of health-related tweets that were crawled via the Twitter API. Focusing on the tasks related to PHS, keywords used to collect pretraining corpus are set to disease, symptom, vaccine, and mental health-related words in English. Pre-processing methods similar to those used in previous works (Müller et al., 2020; Nguyen et al., 2020) were employed prior to training. Retweet tags were deleted from the raw corpus, and URLs and usernames were replaced with HTTP-URL and @USER, respectively. Additionally, the Python emoji³ library was used to replace all emoticons with their associated meanings. The HuggingFace⁴, an open-source python library, was used to segment tweets. Each sequence of BERT LM inputs is converted to 50,265 vocab-

ulary tokens. Twitter posts are restricted to 200 characters, and during the training and evaluation phase, we used a batch size of 8. Distributed training was performed on a TPU v3-8.

3.3 Fine-tuning for downstream tasks

We applied the pretrained PHS-BERT in the binary and multi-class classification of different PHS tasks such as stress, suicide, depression, anorexia, health mention classification, vaccine, and covid related misinformation and sentiment analysis. We fine-tuned the PLMs in downstream tasks. Specifically, we used the `kttrain` library (Maiya, 2020) to fine-tune each model independently for each dataset. We used the embedding of the special token [CLS] of the last hidden layer as the final feature of the input text. We adopted the multilayer perceptron (MLP) with the hyperbolic tangent activation function and used Adam optimizer (Kingma and Ba, 2014). The models are trained with a one cycle policy (Smith, 2017) at a maximum learning rate of $2e-05$ with momentum cycled between 0.85 and 0.95.

4 Experimental Analysis

4.1 Tasks and Datasets

We evaluated and benchmarked the performance of PHS-BERT on 7 different PHS classification tasks (e.g., stress, suicidal ideation, depression, health mention, vaccine, covid related sentiment analysis, and other health-related tasks) collected from popular social platforms (e.g., Reddit and Twitter). We used 25 datasets (see Table 1) crawled from social media platforms (e.g., Reddit and Twitter). We relied on the datasets that are widely used in the community and described each of these tasks and

³<https://pypi.org/project/emoji/>

⁴<https://huggingface.co/>

Table 1: Statistics of the datasets used. We used the Stratified 5-Folds cross-validation (CV) strategy for train/test split if original datasets do not have an official train/test split.

Task (Classification)	Dataset	Platform	# of Samples	# of Classes	Training Strategy Used
Suicide	R-SSD (Cao et al., 2019)	Reddit	500 Users	5	Stratified 5-Folds CV
Stress	Dreaddit (Turcan and McKeown, 2019)	Reddit	3553 Posts	2	Official Split
	SAD (Mauriello et al., 2021)	SMS-like	6850 SMS	2	Official Split
Health Mention	PHM (Karisani and Agichtein, 2018)	Twitter	4635 Posts	4	Stratified 5-Folds CV
	PHM (Karisani and Agichtein, 2018)	Twitter	4635 Posts	2	Stratified 5-Folds CV
	HMC2019 (Biddle et al., 2020)	Twitter	15393 Posts	3	Stratified 5-Folds CV
	RHMD (Naseem et al., 2022b)	Reddit	3553 Posts	4	Stratified 5-Folds CV
Vaccine Sentiment	VS1 (Dunn et al., 2020)	Twitter	9261 Posts	3	Stratified 5-Folds CV
	VS2 (Müller and Salathé, 2019)	Twitter	18522 Posts	3	Stratified 5-Folds CV
COVID Related	Covid Lies (Hossain et al., 2020)	Twitter	3204 Posts	3	Stratified 5-Folds CV
	Covid Category (Müller et al., 2020)	Twitter	4328 Posts	2	Stratified 5-Folds CV
	COVIDSentiA (Naseem et al., 2021d)	Twitter	30000 Posts	3	Stratified 5-Folds CV
	COVIDSentiB (Naseem et al., 2021d)	Twitter	30000 Posts	3	Stratified 5-Folds CV
	COVIDSentiC (Naseem et al., 2021d)	Twitter	30000 Posts	3	Stratified 5-Folds CV
Depression	eRISK T3 (Losada and Crestani, 2016)	Reddit	190 Users	4	Stratified 5-Folds CV
	Depression_Reddit_1 (Naseem et al., 2022a)	Reddit	3553 Posts	4	Stratified 5-Folds CV
	eRISK19 T1 (Losada and Crestani, 2016)	Reddit	2810 Users	2	Official Split
	Depression_Reddit_2 (Pirina and Çöltekin, 2018)	Reddit	1841 Posts	2	Stratified 5-Folds CV
	Depression_Twitter_1	Twitter	1793 Posts	3	Stratified 5-Folds CV
	Depression_Twitter_2	Twitter	10314 Posts	2	Stratified 5-Folds CV
Other Health related	PubHealth (Kotonya and Toni, 2020)	News Websites	12251 Posts	4	Official Split
	Abortion (Mohammad et al., 2016)	Twitter	933 Posts	3	Official Split
	Amazon Health (He and McAuley, 2016)	Amazon	2003 Posts	4	Official Split
	SMM4H T1 (Weissenbacher et al., 2018)	Twitter	14954 Posts	2	Official Split
	SMM4H T2 (Weissenbacher et al., 2018)	Twitter	13498 Posts	3	Official Split
	HRT (Paul and Dredze, 2012)	Twitter	2754 Posts	4	Stratified 5-Folds CV

datasets. Below we briefly discussed each task and dataset used in our study (appendix A for details).

- Suicide:** The widespread use of social media for expressing personal thoughts and emotions makes it a valuable resource for assessing suicide risk on social media. We used the following dataset to evaluate the performance of our model. We used R-SSD (Cao et al., 2019) dataset to evaluate the performance of our model on suicide risk detection.
- Stress:** It is desirable to detect stress early in order to address the growing problem of stress. To evaluate stress detection using social media, we evaluated PHS-BERT on the Dreaddit (Turcan and McKeown, 2019) and SAD (Mauriello et al., 2021) datasets.
- Health mention:** In social media platforms, people often use disease or symptom terms in ways other than to describe their health. In data-driven PHS, the health mention classification task aims to identify posts where users discuss health conditions rather than using disease and symptom terms for other reasons. We used PHM (Karisani and Agichtein, 2018), HMC2019 (Biddle et al., 2020) and RHMD⁵ health mention-related datasets.

⁵<https://github.com/usmaann/RHMD-Health-Mention-Dataset>

- Vaccine sentiment:** Vaccines are a critical component of public health. On the other hand, vaccine hesitancy and refusal can result in clusters of low vaccination coverage, diminishing the effectiveness of vaccination programs. Identifying vaccine-related concerns on social media makes it possible to determine emerging risks to vaccine acceptance. We used VS1 (Dunn et al., 2020) and VS2 (Müller and Salathé, 2019) vaccine-related Twitter datasets to show the effectiveness of our model.
- COVID related:** Due to the ongoing pandemic, there is a higher need for tools to identify COVID-19-related misinformation and sentiment on social media. Misinformation can have a negative impact on public opinion and endanger the lives of millions of people if precautions are not taken. We used COVID Lies (Hossain et al., 2020), Covid category (Müller et al., 2020), and COVIDSenti (Naseem et al., 2021d)⁶ datasets to test our model.
- Depression:** User-generated text on social media has been actively explored for its feasibility in the early identification of depression. We used following eRisk T3 (Losada and Crestani, 2016), eRisk T1 (Losada and Crestani, 2016), Depression_Reddit_1 (Naseem et al.,

⁶we used 3 subsets (COVIDSentiA, COVIDSentiB and COVIDSentiC)

2022a)⁷, Depression_Reddit_2 (Pirina and Çöltekin, 2018), Depression_Twitter_1⁸, and Depression_Twitter_2⁹ depression-related datasets in our experiments.

7. **Other health related tasks:** We also evaluated the performance of our PHS-BERT on other health-related 6 datasets. We used PUBHEALTH (Kotonya and Toni, 2020), Abortion (Mohammad et al., 2016)¹⁰, Amazon Health dataset (He and McAuley, 2016), SMM4H T1 (Weissenbacher et al., 2018), SMM4H T2 (Weissenbacher et al., 2018) and HRT (Paul and Dredze, 2012).

4.2 Evaluation Metric

To evaluate the performance, we used F1-score and the relative improvement in marginal performance (ΔMP) used in a previous similar study (Müller et al., 2020).

4.3 Baselines

We evaluated the performance of PHS-BERT with various SOTA existing PLMs in different domains. We compared the performance with BERT (Devlin et al., 2019), ALBERT (Lan et al., 2019), and DistilBERT (Sanh et al., 2019) pretrained with general corpus, BioBERT (Lee et al., 2019) pretrained in the biomedical domain, CT-BERT (Müller et al., 2020) and BERTweet (Nguyen et al., 2020) pretrained on covid related tweets and MentalBERT (Ji et al., 2021) pretrained on corpus from Reddit from mental health-related subreddits.

4.4 Results

Table 2 summarizes the results of the presented PHS-BERT in comparison to the baselines. We observe that the performance of PHS-BERT is higher than SOTA PLMs on all tested tasks and datasets. Below we discuss the performance comparison of PHS-BERT with BERT and the results of the second-best PLM.

Suicide Ideation Task: We observed that the marginal increases in performance of PHS-BERT is 18.45% when compared to BERT and 12.79% when compared to second best results.

⁷https://github.com/usmaann/Depression_Severity_Dataset

⁸<https://github.com/AshwanthRamji/Depression-Sentiment-Analysis-with-Twitter-Data>

⁹<https://github.com/viritaromero/Detecting-Depression-in-Tweets>

¹⁰The SemEval 2016 stance detection task has 5 target domains. We used the legalization of abortion.

Stress Detection Task: We showed that PHS-BERT achieved higher performance than the best baseline on both datasets. The average marginal increase in performance of PHS-BERT is 3.80% compared to BERT and 2% when compared to second-best results.

Health Mention Task: PHS-BERT outperformed all the baselines on all health mention classification datasets. The average marginal increase in performance of PHS-BERT is 3.34% compared to BERT and 1.76% when compared to second-best results.

Depression Detection Task: We demonstrated that PHS-BERT outperformed all the baselines on all 6 depression datasets to identify depression on social media. We observed that the average marginal increase in performance of PHS-BERT is 6.03% compared to BERT and 2.76% when compared to second-best results.

Vaccine Sentiment Task: For the vaccine sentiment task, PHS-BERT achieved higher performance compared to all baselines on both datasets. Results showed that the average marginal increase in performance of PHS-BERT is 7.70% than BERT and 0.34% compared to second-best results.

COVID Related Task: PHS-BERT outperformed all baselines on all 5 datasets for COVID-related tasks. On average, the marginal increase in performance is 11.82% compared to BERT and 4.471% compared to the second-best results.

Other Health Related Task: We showed that PHS-BERT outperformed all the baselines on all 6 datasets to identify other health-related tasks on social media. We observed that the average marginal increase in performance of PHS-BERT is 11.82% compared to BERT and 4.71% when compared to second-best results.

4.5 Discussion

We demonstrated the effectiveness of our domain-specific PLM on a downstream classification task related to PHS. Compared to previous SOTA PLMs, PHS-BERT improved the performance on all datasets (7 tasks). Our experimental results showed that BERT, a PLM trained in the general domain, gets competitive results on downstream classification tasks. However, for domain-specific tasks, general domain PLMs (BERT, ALBERT, distilBERT) might need more training on relevant corpora to achieve better performance on the domain-specific downstream classification task. Further, we observed that using a domain-specific PLM trained

Table 2: Comparison of PHS-BERT (Ours) v/s SOTA PLMs. Best results (F1-score) are represented in bold, whereas second-best results are underlined. ΔMP_{BERT} and ΔMP_{SB} represent the marginal increase in performance compared to the BERT and the second-best PLM (under-lined).

Suicide Ideation Task										
Dataset	BERT	ALBERT	distilBERT	CT-BERT	BioBERT	BERTweet	MentalBERT	Ours	ΔMP_{BERT}	ΔMP_{SB}
R-SSD	25.72	23.07	<u>26.96</u>	18.67	23.51	24.82	17.35	30.28	18.45 \uparrow	12.79 \uparrow
Stress Detection Task										
Dreaddit	78.55	79.43	78.22	<u>81.46</u>	78.34	80.03	80.89	82.89	5.60 \uparrow	1.78 \uparrow
SAD	92.66	91.11	91.47	91.11	93.92	<u>94.17</u>	93.23	94.75	2.28 \uparrow	0.62 \uparrow
Average	85.61	85.27	84.85	86.29	86.13	87.10	87.06	88.82	3.80 \uparrow	2.00 \uparrow
Health Mention Task										
PHM (Multi-class)	86.21	80.05	85.06	82.02	82.22	85.59	87.76	89.38	3.72 \uparrow	1.87 \uparrow
PHM (Binary)	91.89	90.53	90.64	92.17	89.62	92.12	<u>92.29</u>	93.27	1.52 \uparrow	1.07 \uparrow
HMC2019	88.99	87.22	88.01	<u>90.82</u>	86.27	90.65	90.17	91.71	3.09 \uparrow	0.99 \uparrow
RHMD	74.20	69.02	73.22	<u>72.87</u>	72.25	74.66	75.28	77.16	5.48 \uparrow	2.53 \uparrow
Average	85.07	81.71	84.23	84.47	82.59	85.76	86.38	87.38	3.34 \uparrow	1.76 \uparrow
Depression Detection Task										
eRisk T3	64.56	64.78	67.33	63.17	64.86	63.56	67.75	68.98	6.95 \uparrow	1.84 \uparrow
Depression_Reddit_1	22.39	21.09	21.95	<u>24.21</u>	24.00	20.84	21.95	28.75	29.73 \uparrow	19.56 \uparrow
eRisk T1	93.72	93.79	93.34	86.74	91.73	91.92	94.30	94.52	0.86 \uparrow	0.24 \uparrow
Depression_Reddit_2	91.33	90.72	91.01	68.16	90.53	91.75	92.70	93.36	2.25 \uparrow	0.72 \uparrow
Depression_Twitter_1	64.17	51.70	66.71	57.11	64.12	64.24	<u>72.95</u>	76.18	19.01 \uparrow	4.49 \uparrow
Depression_Twitter_2	96.99	96.79	96.70	96.96	96.59	96.87	97.09	97.12	0.14 \uparrow	0.03 \uparrow
Average	72.19	69.81	72.84	66.06	71.97	71.53	74.46	76.49	6.03 \uparrow	2.76 \uparrow
Vaccine Sentiment Task										
VS1	74.14	70.00	73.95	<u>79.92</u>	73.30	76.81	71.56	79.96	7.96 \uparrow	0.05 \uparrow
VS2	76.60	74.82	75.91	<u>81.73</u>	76.77	79.10	77.65	82.24	7.46 \uparrow	0.63 \uparrow
Average	75.37	72.41	74.93	80.84	75.04	77.96	74.61	81.10	7.70 \uparrow	0.34 \uparrow
COVID Related Task										
Covid Lies	92.96	91.53	92.14	92.24	93.79	91.07	94.60	95.35	2.60 \uparrow	0.80 \uparrow
COVID Category	93.98	93.94	94.35	<u>95.29</u>	93.72	93.45	94.97	95.83	1.99 \uparrow	0.57 \uparrow
COVIDSentiA	90.90	90.81	90.90	<u>78.96</u>	90.41	66.30	<u>91.55</u>	93.97	3.41 \uparrow	2.67 \uparrow
COVIDSentiB	91.31	89.88	91.06	86.85	91.02	89.46	<u>92.06</u>	93.44	2.36 \uparrow	1.52 \uparrow
COVIDSentiC	91.24	83.72	90.77	84.83	90.55	61.78	<u>91.66</u>	93.11	2.03 \uparrow	1.60 \uparrow
Average	92.08	89.98	91.84	87.63	91.90	80.41	92.97	94.34	2.48 \uparrow	1.49 \uparrow
Other Health Related Task										
PubHealth	60.30	61.43	60.77	<u>63.97</u>	58.85	60.57	57.30	64.77	7.54 \uparrow	1.27 \uparrow
Abortion	58.79	58.59	68.09	<u>70.39</u>	62.53	62.82	63.03	72.31	23.40 \uparrow	2.77 \uparrow
Amazon Health	63.45	63.18	62.30	54.84	60.27	65.50	65.57	68.09	7.43 \uparrow	3.90 \uparrow
SMM4H T1	33.33	33.86	35.80	45.50	39.45	45.87	39.81	46.49	40.71 \uparrow	1.38 \uparrow
SMM4H T2	75.54	72.76	75.12	79.19	73.43	80.20	77.54	80.34	6.44 \uparrow	0.18 \uparrow
HRT	78.67	76.97	78.35	<u>80.90</u>	76.13	80.48	80.46	81.12	3.15 \uparrow	0.28 \uparrow
Average	61.68	61.13	63.41	65.80	61.78	65.91	63.95	68.85	11.82 \uparrow	4.71 \uparrow

on biomedical corpora (BioBERT) is less effective than pretraining on the target domain. We also observed that using CT-BERT, BERTweet, and MentalBERT, which are trained on social media-based text, performs better compared to PLMs trained in the general and biomedical domain. These results also demonstrated the effectiveness of training in a target domain. In particular, CT-BERT has the second-best performance on 9 datasets, and MentalBERT has the second-best performance on 13 datasets. The results of domain-specific PLMs demonstrated that continued pretraining in the relevant domain improves performance on downstream tasks.

5 Conclusion

We present PHS-BERT, a domain-specific PLM trained on health-related social media data. Our results demonstrate that using domain-specific corpora to train general domain LMs improves per-

formance on PHS tasks. On all 25 datasets related to 7 different PHS tasks, PHS-BERT outperforms previous state-of-the-art PLMs. We expect that the PHS-BERT PLM will benefit the development of new applications based on PHS NLP tasks.

Ethics and Societal Impact

Ethics: No additional ethics approval was sought for the analysis of data in this study because data were drawn from already published studies.

Societal Impact: We train and release a PLM to accelerate the automatic identification of tasks related to PHS on social media. Our work aims to develop a new computational method for screening users in need of early intervention and is not intended to use in clinical settings or as a diagnostic tool.

Reproducibility: For reproducibility and future works, PHS-BERT is publicly released and is available at <https://huggingface.co/publichealthsurveillance/PHS-BERT>.

References

- Allison E Aiello, Audrey Renson, and Paul N Zivich. 2020. Social media—and internet-based disease surveillance for public health. *Annual review of public health*, 41:101–118.
- Aaron T Beck, Calvin H Ward, Mock Mendelson, Jeremiah Mock, and John Erbaugh. 1961. An inventory for measuring depression. *Archives of general psychiatry*, 4(6):561–571.
- Rhys Biddle, Aditya Joshi, Shaowu Liu, Cecile Paris, and Guandong Xu. 2020. Leveraging sentiment distributions to distinguish figurative from literal health reports on twitter. In *Proceedings of The Web Conference 2020*, pages 1217–1227.
- Lei Cao, Huijun Zhang, Ling Feng, Zihan Wei, Xin Wang, Ningyun Li, and Xiaohao He. 2019. Latent suicide risk detection on microblog via suicide-oriented word embeddings and layered attention. *arXiv preprint arXiv:1910.12038*.
- Lauren E Charles-Smith, Tera L Reynolds, Mark A Cameron, Mike Conway, Eric HY Lau, Jennifer M Olsen, Julie A Pavlin, Mika Shigematsu, Laura C Streichert, Katie J Suda, et al. 2015. Using social media for actionable disease surveillance and outbreak management: a systematic literature review. *PloS one*, 10(10):e0139701.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Adam G Dunn, Kenneth D Mandl, and Enrico Coiera. 2018. Social media interventions for precision public health: promises and risks. *NPJ digital medicine*, 1(1):1–4.
- Adam G Dunn, Didi Surian, Jason Dalmazzo, Dana Rezazadegan, Maryke Steffens, Amalie Dyda, Julie Leask, Enrico Coiera, Aditi Dey, and Kenneth D Mandl. 2020. Limited role of bots in spreading vaccine-critical information among active twitter users in the united states: 2017–2019. *American Journal of Public Health*, 110(S3):S319–S325.
- Su Golder, Gill Norman, and Yoon K Loke. 2015. Systematic review on the prevalence, frequency and comparative value of adverse events data in social media. *British journal of clinical pharmacology*, 80(4):878.
- Ruining He and Julian McAuley. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *proceedings of the 25th international conference on world wide web*, pages 507–517.
- Kirsty Hope, David N Durrheim, Edouard Tursan d’Espaignet, and Craig Dalton. 2006. Syndromic surveillance: is it a useful tool for local outbreak detection?
- Tamanna Hossain, Robert L Logan IV, Arjuna Ugarte, Yoshitomo Matsubara, Sean Young, and Sameer Singh. 2020. Covidlies: Detecting covid-19 misinformation on social media. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*.
- Shaoxiong Ji, Tianlin Zhang, Luna Ansari, Jie Fu, Prayag Tiwari, and Erik Cambria. 2021. Mentalbert: Publicly available pretrained language models for mental healthcare. *arXiv preprint arXiv:2110.15621*.
- Payam Karisani and Eugene Agichtein. 2018. Did you really just have a heart attack? towards robust detection of personal health mentions in social media. In *Proceedings of the 2018 World Wide Web Conference*, pages 137–146.
- Donna M Kazemi, Brian Borsari, Maureen J Levine, and Beau Dooley. 2017. Systematic review of surveillance by social media platforms for illicit drug use. *Journal of Public Health*, 39(4):763–776.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Neema Kotonya and Francesca Toni. 2020. [Explainable automated fact-checking for public health claims](#). *CoRR*, abs/2010.09926.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. [Albert: A lite bert for self-supervised learning of language representations](#).
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [Biobert: a pre-trained biomedical language representation model for biomedical text mining](#).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- David E Losada and Fabio Crestani. 2016. A test collection for research on depression and language use. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 28–39. Springer.
- Arun S. Maiya. 2020. ktrain: A low-code library for augmented machine learning. *arXiv*, arXiv:2004.10703 [cs.LG].
- Matthew Louis Mauriello, Thierry Lincoln, Grace Hon, Dorien Simon, Dan Jurafsky, and Pablo Paredes. 2021. Sad: A stress annotated dataset for recognizing everyday stressors in sms-like conversational

- systems. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–7.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. [SemEval-2016 task 6: Detecting stance in tweets](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41, San Diego, California. Association for Computational Linguistics.
- Danielle L Mowery, Craig Bryan, and Mike Conway. 2015. Towards developing an annotation scheme for depressive disorder symptoms: A preliminary study using twitter data. In *Proceedings of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, pages 89–98.
- Martin Müller, Marcel Salathé, and Per E Kummervold. 2020. Covid-twitter-bert: A natural language processing model to analyse covid-19 content on twitter. *arXiv preprint arXiv:2005.07503*.
- Martin M Müller and Marcel Salathé. 2019. Crowdbreaks: Tracking health trends using public social media data and crowdsourcing. *Frontiers in public health*, 7:81.
- Usman Naseem, Adam G Dunn, Matloob Khushi, and Jinman Kim. 2021a. Benchmarking for biomedical natural language processing tasks with a domain specific albert. *arXiv preprint arXiv:2107.04374*.
- Usman Naseem, Adam G. Dunn, Jinman Kim, and Matloob Khushi. 2022a. Early identification of depression severity levels on reddit using ordinal classification. In *Proceedings of the Web Conference 2022*, pages 1–10.
- Usman Naseem, Matloob Khushi, Jinman Kim, and Adam G Dunn. 2021b. Classifying vaccine sentiment tweets by modelling domain-specific representation and commonsense knowledge into context-aware attentive gru. *arXiv preprint arXiv:2106.09589*.
- Usman Naseem, Matloob Khushi, Vinay Reddy, Sakthivel Rajendran, Imran Razzak, and Jinman Kim. 2020. Bioalbert: A simple and effective pre-trained language model for biomedical named entity recognition. *arXiv preprint arXiv:2009.09223*.
- Usman Naseem, Jinman Kim, Matloob Khushi, and Adam G. Dunn. 2022b. Identification of disease or symptom terms in reddit to improve health mention classification. In *Proceedings of the Web Conference 2022*, pages 11–19.
- Usman Naseem, Imran Razzak, Shah Khalid Khan, and Mukesh Prasad. 2021c. A comprehensive survey on word representation models: From classical to state-of-the-art word representation language models. *Transactions on Asian and Low-Resource Language Information Processing*, 20(5):1–35.
- Usman Naseem, Imran Razzak, Matloob Khushi, Peter W Eklund, and Jinman Kim. 2021d. Covidsentiment: A large-scale benchmark twitter data set for covid-19 sentiment analysis. *IEEE Transactions on Computational Social Systems*, 8(4):1003–1015.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. Bertweet: A pre-trained language model for english tweets. *arXiv preprint arXiv:2005.10200*.
- Michael J Paul and Mark Dredze. 2012. A model for mining public health topics from twitter. *Health*, 11(16-16):1.
- Michael J Paul and Mark Dredze. 2017. Social monitoring for public health. *Synthesis Lectures on Information Concepts, Retrieval, and Services*, 9(5):1–183.
- Inna Pirina and Çağrı Çöltekin. 2018. Identifying depression on reddit: The effect of training data. In *Proceedings of the 2018 EMNLP Workshop SMM4H: The 3rd Social Media Mining for Health Applications Workshop & Shared Task*, pages 9–12.
- Guozheng Rao, Chengxia Peng, Li Zhang, Xin Wang, and Zhiyong Feng. 2020. A knowledge enhanced ensemble learning model for mental disorder detection on social media. In *International Conference on Knowledge Science, Engineering and Management*, pages 181–192. Springer.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Ramit Sawhney, Harshit Joshi, Saumya Gandhi, and Rajiv Ratn Shah. 2021. Towards ordinal suicide ideation detection on social media. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 22–30.
- Leslie N Smith. 2017. Cyclical learning rates for training neural networks. In *2017 IEEE winter conference on applications of computer vision (WACV)*, pages 464–472. IEEE.
- Elsbeth Turcan and Kathleen McKeown. 2019. Dreddit: A reddit dataset for stress analysis in social media. *arXiv preprint arXiv:1911.00133*.
- Davy Weissenbacher, Abeer Sarker, Michael J. Paul, and Graciela Gonzalez-Hernandez. 2018. [Overview of the third social media mining for health \(SMM4H\) shared tasks at EMNLP 2018](#). In *Proceedings of the 2018 EMNLP Workshop SMM4H: The 3rd Social Media Mining for Health Applications Workshop & Shared Task*, pages 13–16, Brussels, Belgium. Association for Computational Linguistics.
- Li Zhang, Haimeng Fan, Chengxia Peng, Guozheng Rao, and Qing Cong. 2020. Sentiment analysis methods for hpv vaccines related tweets based on transfer learning. In *Healthcare*, volume 8, page 307. Multi-disciplinary Digital Publishing Institute.

A Dataset description

1. **Depression:** We used 6 depression-related datasets in our experiments.

- **eRisk T3:** We used eRISK, a publicly available dataset, released by (Losada and Crestani, 2016) and labeled across 4 depression severity levels using Beck’s Depression Inventory (Beck et al., 1961) criteria to detect the existence of depression and identify its severity level in social media posts. eRISK was later used in the CLEF’s eRISK challenge Task 3¹¹ on early identification of depression in social media. Since in each years’ challenge author released a small number of user’s data (ranging from 70-90 users data), we combined and used the data of the last 3 years, which is equivalent to 190 Reddit users, labeled across 4 depression severity levels.
- **Depression_Reddit_1:** We used new Reddit depression data released by Naseem et al. (2022a). This dataset consists of 3,553 Reddit posts to identify the depression severity on social media. Annotators manually labeled data into 4 depression severity levels i.e., (i) minimal depression; (ii) mild depression, (iii) moderate depression; and (iv) severe depression using Depressive Disorder Annotation scheme (Mowery et al., 2015).
- **eRisk T1:** The third depression data is from eRisk shared task 1 (Losada and Crestani, 2016), which is a public competition for detecting early risk in health-related areas. The eRisk data consists of posts from 2,810 users, with 1,370 expressing depression and 1,440 as a control group without depression.
- **Depression_Reddit_2:** The fourth depression dataset used is released by Pirina and Çöltekin (Pirina and Çöltekin, 2018). The authors used Reddit to collect additional social data, which they combined with previously collected data to identify depression.
- **Depression_Twitter_1:** Our fifth depression dataset is a publicly available¹². This data is collected from Twitter and labeled into 3 labels (e.g., Positive, Negative, and Neutral) for depression sentiment analysis.
- **Depression_Twitter_2:** Our sixth depression

dataset is a public dataset¹³, collected from Twitter and labeled into 2 labels (e.g., Positive and Negative) for depression detection.

2. **Health Mention:** We used 3 health mention-related datasets in our experiments.

- **PHM:** Karisani and Agichtein (2018) constructed and released the PHM dataset consisting of 7,192 English tweets across 6 diseases and symptoms. They used the Twitter API to retrieve the data using the colloquial disease names as search keywords. They manually annotated the tweets and categorized them into 4 labels. In addition to 4 labels, similar to Karisani and Agichtein (2018) we also used binary labels for health mention classification.
 - **HMC2019:** HMC2019 is presented by Biddle et al. (2020) by extending the PHM dataset to include 19,558 tweets and included labels related to figurative mentions, and included 4 more different disease or symptom terms (10 in total) for health mention classification.
 - **RHMD:** We also used Reddit health mention dataset (RHMD)(Naseem et al., 2022b) for HMC task. RHMD consists of 10K+ Reddit posts manually annotated with 4 labels (personal health mention, non-personal health mention, figurative health mention, hyperbolic health mention). In our study, we used 3 label versions of data released by authors where they merged figurative health mention and hyperbolic health mention into 1 class.
3. **Suicide:** We used the following dataset to evaluate the performance of our model on suicide risk detection.
- **R-SSD:** For suicide ideation, we used a dataset released by Cao et al. (2019), which contains 500 individuals’ Reddit postings categorized into 5 increasing suicide risk classes from 9 mental health and suicide-related subreddits.
4. **Stress:** To evaluate stress detection using social media, we evaluated PHS-BERT on the following datasets.
- **Dreaddit:** For stress detection, we used Dreaddit (Turcan and McKeown, 2019) collected from 5 different Reddit forums.

¹¹<https://erisk.irlab.org/2021/index.html>

¹²<https://github.com/AshwanthRamji/Depression-Sentiment-Analysis-with-Twitter-Data>

¹³<https://github.com/viritaromero/Detecting-Depression-in-Tweets>

Dreaddit consists of 3,553 posts and focuses on three major stressful topics: interpersonal conflict, mental illness, and financial need. Posts in Dreaddit are collected from 10 subreddits, including some mental health domains such as anxiety and PTSD.

- **SAD:** The SAD (Mauriello et al., 2021) dataset, which contains 6,850 SMS-like sentences, is used to recognize everyday stressors. The SAD dataset is derived from stress management articles, chatbot-based conversation systems, crowdsourcing, and web crawling. Some of the more specific stressors are work-related issues like fatigue or physical pain, financial difficulties like debt or anxiety, school-related decisions like final projects or group projects, and interpersonal relationships like friendships and family relationships.
5. **Vaccine sentiment:** We used two vaccine-related Twitter datasets to show the effectiveness of our model.
- **VS1:** Our first dataset consists of tweets about vaccine dissemination on Twitter from January 12, 2017, to December 3, 2019. Dunn et al. (2020) crawled and labeled this data. The total tweets count is 9,212, with 6,683 positive, 1,084 negatives, and 1,445 neutral tweets.
 - **VS2:** The second dataset¹⁴ includes tweets about measles and vaccinations obtained via the Twitter Streaming API between July 2018 and January 2019 and provided by Müller and Salathé (2019). The total number of tweets is 18,503, with 8,965 pro-vaccine tweets, 1,976 anti-vaccine tweets, and 7,562 neutral tweets.
6. **COVID:** We used 5 covid related datasets to test our model.
- **COVID Lies:** Hossain et al. (2020) released COVIDLIES, a dataset (6761 tweets) annotated by experts with known COVID-19 misconceptions and tweets that agree, disagree, or express no stance.
 - **Covid category:** Covid category dataset is released by Müller et al. (2020). Amazon Turk annotators were asked to classify a given tweet
- text as personal narrative or news. Crowdbreaks was used to perform the annotation.
- **COVIDSenti:** We used a newly released large-scale sentiment dataset, COVIDSenti, which contains 90,000 COVID-19-related tweets obtained during the pandemic’s early stages, from February to March 2020. The tweets are labeled into positive, negative, and neutral sentiment classes. In our experiments, we used 3 subsets (COVIDSentiA, COVIDSentiB and COVIDSentiC) released by authors (Naseem et al., 2021d).
7. **Other health related tasks:** We used PUBHEALTH (Kotonya and Toni, 2020), a dataset for automated fact-checking of public health claims that are explainable. PUBHEALTH is labeled with its factuality (true, false, unproven, mixture). (ii) Abortion: In SemEval 2016 stance detection task (Mohammad et al., 2016), 5 target domains are given: legalization of abortion, atheism, climate change, feminism, and Hillary Clinton. We used the legalization of abortion in our experiments. (iii) Amazon Health dataset: The Amazon Health dataset (He and McAuley, 2016) contains reviews of Amazon healthcare products and has 4 classes i.e., strongly positive, positive, negative, and strongly negative. (iv) SMM4H T1: We used Social Media Mining for Health (SMM4H) Shared Task 1 recognizing whether a tweet is reporting an adverse drug reaction (Weissenbacher et al., 2018). (v) SMM4H T2: Drug Intake Classification (SMM4H Task 2) (Weissenbacher et al., 2018) where participants were given tweets manually categorized as definite intake, possible intake, or no intake. (vi) HRT: Health related tweets (HRT) (Paul and Dredze, 2012) were collected using Twitter and manually annotated using Mechanical Turk as related or unrelated to health. Health-related tweets were further labeled as sick (the text implied that the user was suffering from an acute illness, such as a cold or the flu) or health (the text made general comments about the user’s or the other’s health, such as chronic health conditions, lifestyle, or diet) and unrelated tweets were further labeled as unrelated (texts that were not about a specific person’s health, such as news and updates about the swine flu or advertisements for diet pills) and non-English.

¹⁴<https://github.com/digitalepidemiologylab/crowdbreaks-paper>