





A Unique Training Strategy to Enhance Language Models Capabilities for Health Mention Detection from Social Media Content

Pervaiz Iqbal Khan^{1,2}, Muhammad Nabeel Asim¹, Andreas Dengel^{1,2} and Sheraz Ahmed¹

¹German Research Center for Artificial Intelligence (DFKI), Kaiserslautern, Germany

²RPTU Kaiserslautern-Landau, Germany

Keywords: Language Models, Contrastive Learning, Social Media Content Analysis, Health Mention Detection, Meta Predictor.

Abstract: An ever-increasing amount of social media content requires advanced AI-based computer programs capable of extracting useful information. Specifically, the extraction of health-related content from social media is useful for the development of diverse types of applications including disease spread, mortality rate prediction, and finding the impact of diverse types of drugs on diverse types of diseases. Language models are competent in extracting the syntactic and semantics of text. However, they face a hard time extracting similar patterns from social media texts. The primary reason for this shortfall lies in the non-standardized writing style commonly employed by social media users. Following the need for an optimal language model competent in extracting useful patterns from social media text, the key goal of this paper is to train language models in such a way that they learn to derive generalized patterns. The key goal is achieved through the incorporation of random weighted perturbation and contrastive learning strategies. On top of a unique training strategy, a meta predictor is proposed that reaps the benefits of 5 different language models for discriminating posts of social media text into non-health and health-related classes. Comprehensive experimentation across 3 public benchmark datasets reveals that the proposed training strategy improves the performance of the language models up to 3.87%, in terms of F1-score, as compared to their performance with traditional training. Furthermore, the proposed meta predictor outperforms existing health mention classification predictors across all 3 benchmark datasets.


1 INTRODUCTION


Social media platforms like Facebook¹, Twitter², and Reddit³ have played a significant role in connecting people worldwide, effectively shrinking the world into a global community (Abbas et al., 2019). Through these platforms, people can communicate and share their opinions regarding product quality, service standards, and even their emotions and concerns about health issues. Various brands are utilizing the power of artificial intelligence (AI) methods and social media platforms for enhancing the qual-


ity of their products and services. To empower the process of brand monitoring by analyzing social media content, there is a marathon for developing more robust and precise deep learning predictors (Wassan et al., 2021; Jagdale et al., 2019; Ray and Chakrabarti, 2017). Similar to brand monitoring, healthcare systems can be improved by closely monitoring social media content; however, little attention has been paid in this regard and a few predictors have been proposed for discriminating health-related content from other conversations. Accurate categorization of social media content into health and general discussion related classes known as Health Mention Classification (HMC), can provide significant insights about the spread of diseases, mortality rate analysis, the impact of drugs on individuals (Huang et al., 2022; Jahanbin et al., 2020), etc.

Among diverse types of predictors BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019) and XL-

^a <https://orcid.org/0000-0002-1805-335X>

^b <https://orcid.org/0000-0001-5507-198X>

^c <https://orcid.org/0000-0002-6100-8255>

^d <https://orcid.org/0000-0002-4239-6520>

¹<https://www.facebook.com/>

²<https://twitter.com>

³<https://www.reddit.com/>

Net (Yang et al., 2019), etc. are producing state-of-the-art (SotA) performance values for diverse types of text classification tasks. The prime reason behind their better performance is the utilization of transfer learning that helps them understand the contextual information of content. Trained language models offer better performance on small datasets where they are trained in a supervised fashion. These language models are also being utilized for the development of diverse types of real-world useful applications such as DNA (Zeng et al., 2023), RNA sequence classification (Yang et al., 2022), Fake URL detection (Maneriker et al., 2021), hate speech detection (Mozafari et al., 2020), fake news detection (Tariq et al., 2022), and brand monitoring (Jagdale et al., 2019). Researchers are trying to utilize these language models to strengthen healthcare systems by developing diverse types of applications including mortality rate prediction, disease spread prediction (Li et al., 2021), and disease symptoms prediction (Luo et al., 2021). Development of these applications requires discrimination of non-health mention related social media content from health mention content. Precise discrimination of this content is difficult as compared to the aforementioned tasks as people use non-standardized and abbreviated forms of words while describing diseases, symptoms, and drugs. Furthermore, the posts indicating the presence of disease and general conversation about the diseases may contain similar words, hence learning the appropriate context becomes more non-trivial. This paper proposes a random weighted perturbation (RWP) method for finetuning language models on the HMC task, where it generates perturbations based on the statistical distribution of the model's parameters. Furthermore, it jointly optimizes language model training by taking both perturbed and unperturbed parameters. Moreover, it introduces contrastive learning as an additional objective function to further improve text representation and the classifier's performance.

Manifold contributions of this paper can be summarized as:

- It presents a unique strategy to train language models that avoids over-fitting and improves their generalization capabilities.
- Over 3 public benchmark datasets, it performs large-scale experimentation, where the aim is to highlight the impact of the proposed training strategy over diverse types of language models.
- With an aim to develop a robust and precise predictor for health mention classification, it presents a meta predictor that benefits from 5 different language models.

- The proposed meta predictor along with the unique training strategy produces SotA performance values over 3 public benchmark datasets.

2 RELATED WORK

This section presents 11 different types of predictors that are proposed for the HMC task. The presented predictors employ different approaches, such as convolutional neural networks (CNNs), language models, and hybrid approaches.

(Iyer et al., 2019) proposed a two-step approach in which binary features were extracted to indicate whether a disease word was used figuratively. This information was passed to the CNN model as an additional feature for tweet classification. (Luo et al., 2022) presented a health mention classifier that used dual CNN to find COVID-19 mentions. The dual CNN had two components: an additional network called an auxiliary network (A-Net), which helped the primary network (P-Net) overcome the class imbalance. (Karisani and Agichtein, 2018) proposed WESPAD, which divided and distorted the embedding space to make the model generalize better on examples that had not been seen. (Jiang et al., 2018) used non-contextual embeddings to extract representations of the tweets, and then passed those representations to Long Short-Term Memory Networks (Hochreiter and Schmidhuber, 1997). By adding 14k new tweets related to 10 diseases, (Biddle et al., 2020) extended the PHM-2017 (Karisani and Agichtein, 2018) dataset. Moreover, they incorporated sentiment information with contextual and non-contextual embeddings into their study. The incorporation of sentiment information significantly enhanced the classification performance of the classification algorithms. (Khan et al., 2020) improved the performance of the health mention classifier by incorporating emojis into tweet text and using XLNet-based (Yang et al., 2019) word representations. (Naseem et al., 2022a) presented a new health mention dataset related to Reddit posts, and classified the posts based on a combination of disease or symptom terms and user behavior. A method for learning the context and semantics of HMC was proposed by (Naseem et al., 2022b). The method used domain-specific word representations and Bi-LSTM with an attention module. (Naseem et al., 2022c) proposed the PHS-BERT, a domain-specific pretrained language model for social media posts, and evaluated the model on various social media datasets. However, pre-training language models require a significant amount of time and computing resources.

A contrastive adversarial training method was pre-

sented by (Khan et al., 2022b) using the Fast Gradient Sign Method (Goodfellow et al., 2014) to perturb the embedding matrix parameters of the model. Then, they jointly trained the clean and adversarial examples and employed contrastive learning. (Khan et al., 2022a) presented a training method for HMC, in which they added Gaussian noise to the hidden representations of language models. The cleaned and perturbed examples were trained simultaneously. The empirical evidence demonstrated that perturbing the earlier layers of the model enhanced its performance. This paper proposes a new training strategy for language models that uses random weighted perturbations in the parameters of language models and additionally employs contrastive learning to improve text representations produced by the models.

3 METHOD

3.1 Proposed Training Method and Meta Predictor

Fig. 1 illustrates the workflow of the proposed training strategy, where it can be seen that textual data is processed at two different streams of transformer models, i.e., one stream with normal transformer model setting, and the other with random weighted perturbations. Here, transformer models represent 5 language models, namely, BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019), DeBERTa (He et al., 2020), XLNet (Yang et al., 2019), and GPT-2 (Radford et al., 2018). Furthermore, for both streams, loss is computed using cross-entropy. Moreover, an additional third stream is introduced, where CLS token representations are extracted from both the original and perturbed models for each training example and then projected to lower dimensions using a projection network. Then contrastive loss is computed using the Barlow Twins (Zbontar et al., 2021) that takes these two lower dimensional representations as inputs. Finally, the total loss \mathcal{L}_{total} is computed as follows:

$$\mathcal{L}_{total} = \frac{(1 - \lambda)}{2} (\mathcal{L}_{CE_1} + \mathcal{L}_{CE_2}) + \lambda \mathcal{L}_{BT} \quad (1)$$

where \mathcal{L}_{BT} is the Barlow Twins (Zbontar et al., 2021) loss, \mathcal{L}_{CE_1} , and \mathcal{L}_{CE_2} are two cross-entropy losses, and λ is the hyperparameter that controls the weight between the three losses.

To utilize the strength of various transformer models and improve the prediction score, a meta predictor is proposed at the test time. Specifically, the class prediction probabilities of each of the 5 models are taken for every test-set post, and then averaged. Then the

maximum probability for each post is treated as the predicted class. Results suggest that an ensemble of model predictions improves the overall classification scores for all the datasets.

A more comprehensive detail about different modules of the proposed method is presented in the following subsections:

3.2 Transformers

Let $\mathcal{D} = \{x_i, y_i\}_{i=1, \dots, N}$ be a dataset having N examples with X data points and their corresponding Y labels. Let \mathbf{M} be a pre-trained language model. Each training example is represented as T tokens, i.e., $x_i = \{CLS, t_1, t_2, \dots, t_T, SEP\}$, where CLS, and SEP are special tokens. This token representation is passed as an input to \mathbf{M} that produces the embeddings, i.e., $\{h_{CLS}^L, h_1^L, \dots, h_T^L, h_{SEP}^L\}$ for the given input example. Here, L represents the total layers in model M. To finetune M, a classification layer is added on top of M, and then the cross-entropy loss is minimized as follows:

$$\mathcal{L}_{CE} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C y_{i,c} \log(p(y_i, c | h_{[CLS]}^i)) \quad (2)$$

Here, C represents the total no. of classes, and $h_{[CLS]}$ contains the representations of the complete example.

3.3 Random Weighted Perturbations and Contrastive Learning

Let $w = \{w_1, w_2, \dots, w_L\}$ be the model parameters in all the layers. For each model parameter w_L , perturbations are generated following the Gaussian distribution $N(0, \epsilon \|w_L\|_2)$, elementally, where ϵ controls the strength of perturbations, and $\|w_L\|_2$ is L_2 norm. The parameters with higher norm values will have a higher perturbation amount. The generated perturbations are added into the model's parameters.

Contrastive Learning (CL) methods take two inputs, one of which is clean and the other is its perturbed version. The objective of CL is to learn embeddings for the inputs such that both the clean and its perturbed versions are pushed close, and other examples are pushed apart in the representation space. Barlow Twins (Zbontar et al., 2021) (BT) is a CL method that works on the redundancy reduction principle. Instead of perturbing the data point itself, in our experiments, the output of the perturbed model is treated as a perturbed input to the BT. Let E^c be the embedding of the clean example and E^p be the embedding of its perturbed version. Then, the objective function of the

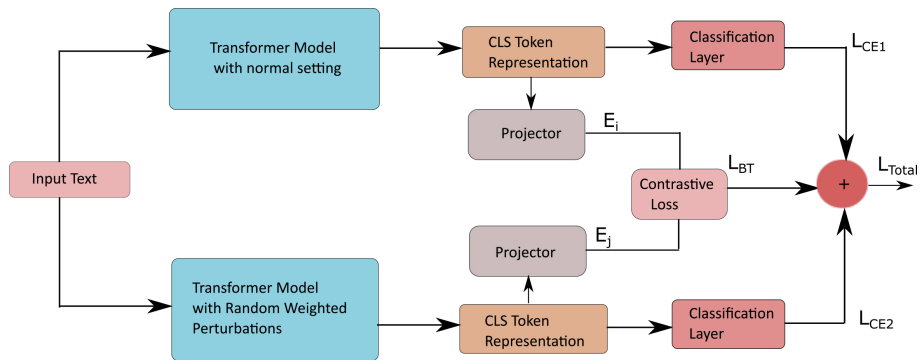


Figure 1: Graphical illustration of the proposed strategy for language models training.

BT is minimized as follows:

$$\mathcal{L}_{BT} = \sum_{i=1} (1 - A_{ii})^2 + \beta \sum_{i=1} \sum_{j \neq i} A_{ij}^2 \quad (3)$$

where $\sum_{i=1} (1 - A_{ii})^2$, and $\sum_{i=1} \sum_{j \neq i} A_{ij}^2$ represent invariance, and redundancy reduction terms respectively, and β is the trade-off parameter controlling weights between the two terms. The matrix A computes the cross-correlation between E^c , and E^p . A is computed as follows:

$$A_{ij} = \frac{\sum_{b=1}^N E_{b,i}^c E_{b,i}^p}{\sqrt{\sum_{b=1}^N (E_{b,i}^c)^2} \sqrt{\sum_{b=1}^N (E_{b,i}^p)^2}} \quad (4)$$

where b is the batch size, and A_{ij} represents the entry of the i -th row and j -th column of the matrix A .

4 EXPERIMENTS

4.1 Benchmark Datasets

The selection of appropriate data is an essential part of evaluating the predictive performance of a proposed predictor. Following the experimental settings and evaluation criteria of existing studies (Khan et al., 2022b), we evaluated our proposed predictor using 3 different public benchmark datasets. A comprehensive description of the development process of these datasets is provided in existing studies (Khan et al., 2022b), so here we only summarize the statistics of these datasets as shown in Table 1.

Furthermore, we performed preprocessing over these datasets and removed user mentions, hashtags, and URLs from the COVID-19 PHM (Luo et al., 2022) and HMC-2019 (Biddle et al., 2020) datasets. Following the existing studies (Khan et al., 2020; Khan et al., 2022b), we converted the emojis in the tweet into text and removed special characters such as “:” and “-” and used those textual representations

as a part of the tweet. Moreover, we used maximum sequence length of 64, 68, and 215 for HMC-2019, COVID-19 PHM, and RHMD datasets, respectively.

4.2 Training Details

We finetuned 5 different language models on the 3 datasets. We used AdamW (Loshchilov and Hutter, 2018) as an optimizer for all experiments. To obtain optimal values of hyperparameters, we tweaked several hyperparameters such as batch size $b \in \{16, 32\}$, $\epsilon \in \{5e^{-4}, 1e^{-4}, 5e^{-3}, 1e^{-3}\}$, $\lambda \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$, and a learning rate of $1e^{-5}$ using grid-search. Furthermore, we trained all models for 40 epochs with an early stopping strategy based on the validation set F1-score. The projection network used for contrastive loss consists of two linear layers with a hidden unit size of 1024, and an output layer with a hidden unit size of 300, respectively. As an activation function, we used ReLU and applied 1-D batch normalization between the two linear layers. For BT (Zbontar et al., 2021) loss, we used the default hyperparameters.

5 RESULTS AND ANALYSIS

5.1 Contrastive Learning and Random Weighted Perturbations Impact on Models Predictive Performance

Table 2 shows the results of 5 different language models with traditional and proposed training strategy. Overall, 4 models out of 5 performed better with the proposed method on all 3 datasets in terms of F1-score as an evaluation measure. On the contrary, DeBERTa (He et al., 2020) showed different behavior.

On 2 out of 3 datasets, it performed better with

Table 1: Statistics of the benchmark datasets used for validating the proposed training strategy.

Dataset	No. of Samples	No. of Diseases	No. of Classes
HMC-2019 (Biddle et al., 2020)	15,742	10	2
RHMD (Naseem et al., 2022a)	10,015	15	3
COVID-19 PHM (Luo et al., 2022)	9,219	1	2

Table 2: showing that random weighted perturbation with contrastive learning (RWP + CL) improves F1-scores (macro) on 3 HMC datasets. Results are reported on the test sets of all the datasets.

Model	HMC-2019	RHMD	COVID-19 PHM
BERT baseline (Devlin et al., 2018)	92.00	79.81	78.22
BERT (Devlin et al., 2018) + RWP + CL	92.47	80.30	79.65
RoBERTa baseline (Liu et al., 2019)	93.58	80.20	76.80
RoBERTa (Liu et al., 2019) + RWP + CL	93.95	80.72	78.65
DeBERTa baseline (He et al., 2020)	93.66	81.03	79.91
DeBERTa (He et al., 2020) + RWP + CL	93.91	82.36	78.93
XLNet baseline (Yang et al., 2019)	92.53	79.91	77.43
XLNet (Yang et al., 2019) + RWP + CL	93.28	81.21	79.02
GPT-2 baseline (Radford et al., 2018)	88.05	71.63	71.22
GPT-2 (Radford et al., 2018) + RWP + CL	88.46	74.52	75.09

noise. However, its performance on the COVID-19 PHM (Luo et al., 2022) dataset decreased with our proposed training method. Similarly, the difference between DeBERTa (He et al., 2020) performance with our proposed method and its corresponding baseline method on the HMC-2019 (Biddle et al., 2020) dataset is smaller than the other models. One possible explanation for this behavior is the existence of a built-in virtual adversarial training algorithm in the model, which already adds perturbation. Since COVID-19 PHM (Luo et al., 2022) is an imbalanced dataset, the performance degrades on this dataset with additional perturbations. On the other hand, the length of tweets in the HMC-2019 (Biddle et al., 2020) dataset is smaller, hence, additional perturbation with our proposed method doesn't help this model much as compared to the other models. However, as the text length of the RHMD (Naseem et al., 2022a) dataset is longer, additional perturbations in DeBERTa (He et al., 2020) with our proposed training method outperforms its baseline method significantly. One possible explanation for the improved performance of the models using the proposed training strategy might be the prevention of over-fitting. Generally, models learn and memorize the training set quickly, and are unable to learn contextual and semantic information properly during the training phase. However, perturbations prevent the memorization and encourage the generalization of the models by learning the

appropriate representations. CL further improves the representations of the given input text hence improves the models performance.

We further extended the experimentation and performed 10-fold cross-validation for the RHMD and HMC-2019 datasets. For these experiments, we used the best validation set hyperparameters of the train/validation/test splits. The results are presented in Table 3. Results suggested that different models exhibit different performance characteristics across 2 datasets. This motivated us to propose a meta predictor that takes ensemble predictions from these models. The results are presented in the section 5.2.

5.2 Proposed Meta Predictor Performance Analysis and Comparison with SotA

Fig. 2 presents the comparison of the proposed meta predictor with existing predictors. Among the existing predictors, (Khan et al., 2022b) produced better performance on the RHMD, COVID-19 PHM datasets, and (Khan et al., 2022a) on the HMC-2019 dataset. However, like (Luo et al., 2022), these methods are biased towards precision on COVID-19 PHM dataset. The proposed predictor reduced the difference between precision and recall by utilizing RWP and benefiting from the capabilities of 5 models.

Table 3: Results (macro F1-score) on 10-fold cross-validation set for the RHMD (Naseem et al., 2022a) and HMC-2019 (Biddle et al., 2020) datasets.

Method	HMC-2019 (Biddle et al., 2020)	RHMD (Naseem et al., 2022a)
BERT (Devlin et al., 2018) + RWP + CL	93.62	82.16
RoBERTa (Liu et al., 2019) + RWP + CL	94.27	83.05
DeBERTa (He et al., 2020) + RWP + CL	94.07	82.85
XLNet (Yang et al., 2019) + RWP + CL	93.47	81.91
GPT-2 (Radford et al., 2018) + RWP + CL	89.61	74.70

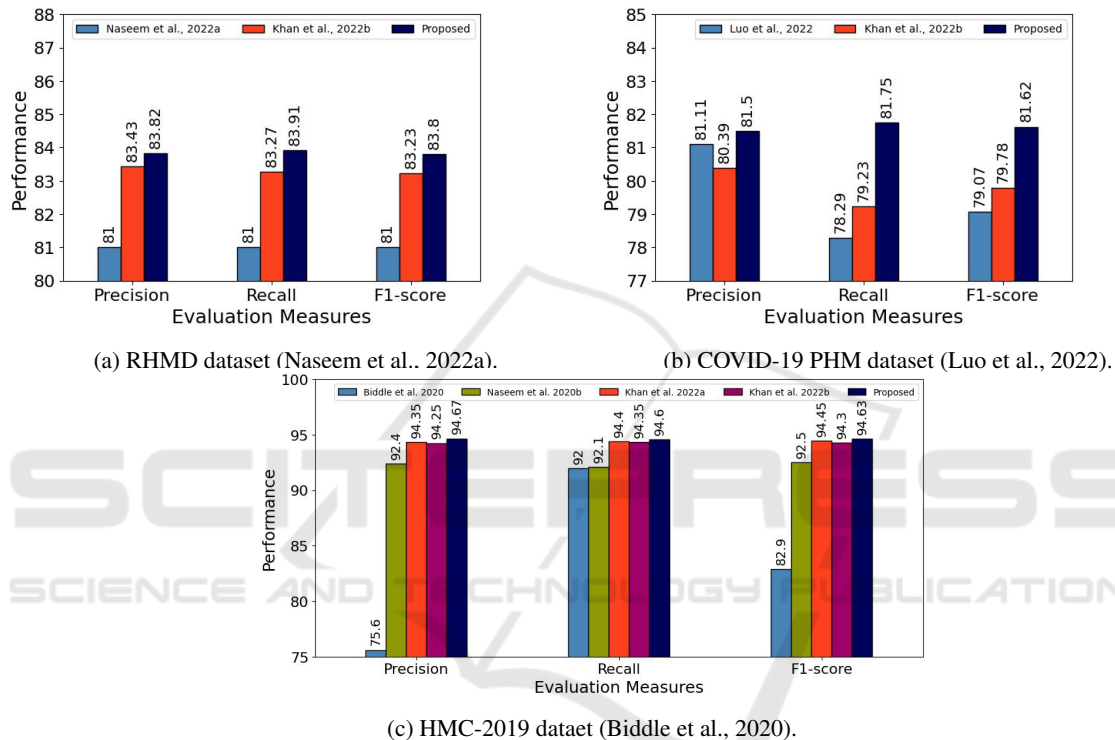


Figure 2: Performance comparison of the proposed meta predictor built on the top of proposed training strategy with SotA on 3 datasets. Results of the SotA are taken from (Khan et al., 2022b). For a fair comparison with existing predictors, RHMD (Naseem et al., 2022a) and HMC-2019 (Biddle et al., 2020) datasets are evaluated using a 10-fold cross-validation.

Table 4: Showing the effectiveness of our proposed method. Results are reported in terms of macro F1-score.

Dataset	Model	Baseline	RWP	RWP + CL
HMC-2019 (Biddle et al., 2020)	XLNet (Yang et al., 2019)	92.53	93.03	93.28
RHMD (Naseem et al., 2022a)	DeBERTa (He et al., 2020)	81.03	80.31	82.36
COVID-19 PHM (Luo et al., 2022)	GPT-2 (Radford et al., 2018)	71.22	71.59	75.09

Overall, our presented meta predictor achieves SotA performance on all 3 datasets.

5.3 Ablation Studies

In Table 4, we present the effectiveness of our method by dropping the CL and RWP components. These

results indicate that performance of the models decreases if proposed components are dropped. In Table 5, we present the noise impact on the performance of the models. Generally, models performances significantly drop for higher noise values, i.e., $\epsilon = 5e^{-3}$. BERT (Devlin et al., 2018) is more sensitive to noise on RHMD and COVID-19 PHM datasets. The per-

Table 5: Showing the amount of noise impact on validation sets performance (macro F1-score) of 3 datasets.

Model	λ	HMC-2019(Biddle et al., 2020)				RHMD (Naseem et al., 2022a)				COVID-19 PHM (Luo et al., 2022)			
		Noise Amount ϵ											
		$1e^{-4}$	$5e^{-4}$	$1e^{-3}$	$5e^{-3}$	$1e^{-4}$	$5e^{-4}$	$1e^{-3}$	$5e^{-3}$	$1e^{-4}$	$5e^{-4}$	$1e^{-3}$	$5e^{-3}$
BERT (Devlin et al., 2018)	0.1	93.75	93.39	93.72	92.60	81.20	81.62	81.26	81.60	77.86	78.24	78.57	78.31
	0.2	92.99	92.91	93.27	93.25	81.42	80.69	81.22	82.13	78.41	78.81	77.60	79.16
	0.3	93.17	93.04	93.36	93.01	80.60	81.58	80.93	81.98	77.60	78.42	78.42	77.76
	0.4	93.20	93.17	93.46	92.78	80.97	80.75	82.05	81.74	77.80	77.54	77.32	78.14
	0.5	93.31	92.98	93.31	93.12	81.21	81.30	80.46	81.85	77.00	76.84	76.53	77.26
RoBERTa (Liu et al., 2019)	0.1	94.48	93.94	93.75	93.27	82.99	83.85	83.08	83.33	80.12	79.31	80.14	78.44
	0.2	93.94	94.26	93.97	92.91	82.01	83.40	83.64	83.12	79.58	79.31	79.86	78.08
	0.3	94.04	94.15	94.40	92.49	82.61	83.77	83.30	82.40	79.99	79.97	79.23	77.27
	0.4	94.32	93.75	94.21	92.67	83.10	83.15	83.17	81.26	79.71	79.31	79.03	61.32
	0.5	94.34	94.10	93.71	91.72	82.91	82.72	82.70	82.17	78.70	78.61	78.69	76.29
DeBERTa (He et al., 2020)	0.1	94.08	93.89	93.84	93.68	83.29	83.81	83.51	83.58	80.44	81.41	80.09	81.33
	0.2	94.13	94.12	93.96	93.91	83.67	83.17	83.03	83.87	80.59	79.94	80.00	80.25
	0.3	93.98	93.86	94.08	93.65	82.45	82.83	82.82	82.86	80.56	80.99	80.53	80.19
	0.4	94.08	94.03	94.08	94.10	83.34	82.92	83.52	82.97	80.34	80.77	80.00	79.35
	0.5	93.82	93.84	93.65	93.69	82.70	82.76	82.84	82.63	79.28	80.34	80.23	79.01
XLNet (Yang et al., 2019)	0.1	93.74	93.60	93.07	92.62	81.95	82.37	82.30	80.72	79.26	78.70	78.52	77.06
	0.2	92.65	93.28	93.23	92.13	81.44	81.99	81.81	81.56	78.35	79.83	78.03	77.32
	0.3	93.22	93.18	93.50	92.05	82.78	82.73	82.64	80.40	79.07	77.58	78.43	77.36
	0.4	92.72	93.55	93.63	92.59	82.03	80.94	81.75	78.47	79.34	77.76	77.97	73.96
	0.5	93.37	93.37	93.32	91.09	81.58	80.67	80.97	79.48	77.98	79.41	78.59	75.81
GPT-2 (Radford et al., 2018)	0.1	89.85	89.92	89.67	87.99	75.15	74.52	77.20	71.28	74.08	74.94	74.25	73.09
	0.2	88.19	89.52	89.33	86.92	74.16	72.46	75.21	72.13	74.15	74.27	73.29	73.20
	0.3	90.75	90.15	89.45	88.42	74.58	75.69	75.54	71.91	73.83	75.26	74.61	73.20
	0.4	89.46	89.75	90.35	85.55	77.50	74.90	72.83	72.39	74.61	73.51	75.37	72.68
	0.5	90.31	89.42	89.54	86.72	76.33	73.39	71.60	69.77	74.32	73.37	73.22	71.54

formance of RoBERTa (Liu et al., 2019) and GPT-2 (Radford et al., 2018) varies significantly for all the datasets. For HMC-2019, the performance of DeBERTa (He et al., 2020) does not change much with a change in the noise. The performance of XLNet (Yang et al., 2019) changes sharply for the higher values of λ on HMC-2019, and RHMD datasets.

6 CONCLUSION

This paper proposed a robust meta predictor that will improve the healthcare system by monitoring social media content related to health. The proposed predictor employed a novel training strategy that improved the predictive performance of the language models on 3 datasets. The improved training process prevents the model from quickly over-fitting, which allows it to learn semantic features. The proposed meta predictor exploits the strengths of the 5 models and outperforms existing SotA methods on all benchmark datasets. In future work, the performance of each language model may be enhanced by combining different perturbation and contrastive learning methods.

REFERENCES

Abbas, J., Aman, J., Nurunnabi, M., and Bano, S. (2019). The impact of social media on learning behavior for sustainable education: Evidence of students

from selected universities in pakistan. *Sustainability*, 11(6):1683.

Biddle, R., Joshi, A., Liu, S., Paris, C., and Xu, G. (2020). Leveraging sentiment distributions to distinguish figurative from literal health reports on twitter. In *Proceedings of The Web Conference 2020*, pages 1217–1227.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Goodfellow, I. J., Shlens, J., and Szegedy, C. (2014). Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*.

He, P., Liu, X., Gao, J., and Chen, W. (2020). DeBERTa: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.

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

Huang, J.-Y., Lee, W.-P., and Lee, K.-D. (2022). Predicting adverse drug reactions from social media posts: Data balance, feature selection and deep learning. In *Healthcare*, volume 10, page 618. MDPI.

Iyer, A., Joshi, A., Karimi, S., Sparks, R., and Paris, C. (2019). Figurative usage detection of symptom words to improve personal health mention detection. *arXiv preprint arXiv:1906.05466*.

Jagdale, R. S., Shirsat, V. S., and Deshmukh, S. N. (2019). Sentiment analysis on product reviews using machine learning techniques. In *Cognitive Informatics and Soft Computing: Proceeding of CISC 2017*, pages 639–647. Springer.

Jahanbin, K., Rahmanian, V., et al. (2020). Using twitter

- and web news mining to predict covid-19 outbreak. *Asian Pacific journal of tropical medicine*, 13(8):378.
- Jiang, K., Feng, S., Song, Q., Calix, R. A., Gupta, M., and Bernard, G. R. (2018). Identifying tweets of personal health experience through word embedding and lstm neural network. *BMC bioinformatics*, 19(8):67–74.
- Karisani, P. and Agichtein, E. (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.
- Khan, P. I., Razzak, I., Dengel, A., and Ahmed, S. (2020). Improving personal health mention detection on twitter using permutation based word representation learning. In *International Conference on Neural Information Processing*, pages 776–785. Springer.
- Khan, P. I., Razzak, I., Dengel, A., and Ahmed, S. (2022a). A novel approach to train diverse types of language models for health mention classification of tweets. In *Artificial Neural Networks and Machine Learning–ICANN 2022: 31st International Conference on Artificial Neural Networks, Bristol, UK, September 6–9, 2022, Proceedings, Part II*, pages 136–147. Springer.
- Khan, P. I., Siddiqui, S. A., Razzak, I., Dengel, A., and Ahmed, S. (2022b). Improving health mention classification of social media content using contrastive adversarial training. *IEEE Access*, 10:87900–87910.
- Li, L., Jiang, Y., and Huang, B. (2021). Long-term prediction for temporal propagation of seasonal influenza using transformer-based model. *Journal of biomedical informatics*, 122:103894.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. (2019). Roberta: A robustly optimized bert pre-training approach. *arXiv preprint arXiv:1907.11692*.
- Loshchilov, I. and Hutter, F. (2018). Fixing weight decay regularization in adam. *arXiv preprint arXiv:2011.08042v1*.
- Luo, L., Wang, Y., and Liu, H. (2022). Covid-19 personal health mention detection from tweets using dual convolutional neural network. *Expert Systems with Applications*, 200:117139.
- Luo, X., Gandhi, P., Storey, S., and Huang, K. (2021). A deep language model for symptom extraction from clinical text and its application to extract covid-19 symptoms from social media. *IEEE Journal of Biomedical and Health Informatics*, 26(4):1737–1748.
- Maneriker, P., Stokes, J. W., Lazo, E. G., Carutasu, D., Tajaddodianfar, F., and Gururajan, A. (2021). Urltran: Improving phishing url detection using transformers. In *MILCOM 2021-2021 IEEE Military Communications Conference (MILCOM)*, pages 197–204. IEEE.
- Mozafari, M., Farahbakhsh, R., and Crespi, N. (2020). A bert-based transfer learning approach for hate speech detection in online social media. In *Complex Networks and Their Applications VIII: Volume 1 Proceedings of the Eighth International Conference on Complex Networks and Their Applications COMPLEX NETWORKS 2019 8*, pages 928–940. Springer.
- Naseem, U., Kim, J., Khushi, M., and Dunn, A. G. (2022a). Identification of disease or symptom terms in reddit to improve health mention classification. In *Proceedings of the ACM Web Conference 2022*, pages 2573–2581.
- Naseem, U., Kim, J., Khushi, M., and Dunn, A. G. (2022b). Robust identification of figurative language in personal health mentions on twitter. *IEEE Transactions on Artificial Intelligence*.
- Naseem, U., Lee, B. C., Khushi, M., Kim, J., and Dunn, A. G. (2022c). Benchmarking for public health surveillance tasks on social media with a domain-specific pretrained language model. *arXiv preprint arXiv:2204.04521*.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. (2018). Improving language understanding by generative pre-training.
- Ray, P. and Chakrabarti, A. (2017). Twitter sentiment analysis for product review using lexicon method. In *2017 International Conference on Data Management, Analytics and Innovation (ICDMAI)*, pages 211–216. IEEE.
- Tariq, A., Mehmood, A., Elhadef, M., and Khan, M. U. G. (2022). Adversarial training for fake news classification. *IEEE Access*, 10:82706–82715.
- Wassan, S., Chen, X., Shen, T., Waqar, M., and Jhanjhi, N. (2021). Amazon product sentiment analysis using machine learning techniques. *Revista Argentina de Clínica Psicológica*, 30(1):695.
- Yang, F., Wang, W., Wang, F., Fang, Y., Tang, D., Huang, J., Lu, H., and Yao, J. (2022). scbert as a large-scale pretrained deep language model for cell type annotation of single-cell rna-seq data. *Nature Machine Intelligence*, 4(10):852–866.
- Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R. R., and Le, Q. V. (2019). Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764.
- Zbontar, J., Jing, L., Misra, I., LeCun, Y., and Deny, S. (2021). Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR.
- Zeng, W., Gautam, A., and Huson, D. H. (2023). Mulan-methyl-multiple transformer-based language models for accurate dna methylation prediction. *bioRxiv*, pages 2023–01.