

Improving Minority Stress Detection with Emotions

Jonathan Ivey
 Department of Data Science
 University of Arkansas
 Fayetteville, AR, USA
 jwi001@uark.edu

Susan Gauch
 Department of Computer Science
 University of Arkansas
 Fayetteville, AR, USA
 sgauch@uark.edu

Abstract—Psychological stress detection is an important task for mental healthcare research, but there has been little prior work investigating the effectiveness of psychological stress models on minority individuals, who are especially vulnerable to poor mental health outcomes. In this work, we use the related task of minority stress detection to evaluate the ability of psychological stress models to understand the language of sexual and gender minorities. We find that traditional psychological stress models underperform on minority stress detection, and we propose using emotion-infused models to reduce that performance disparity. We further demonstrate that multi-task psychological stress models outperform the current state-of-the-art for minority stress detection without directly training on minority stress data. We provide explanatory analysis showing that minority communities have different distributions of emotions than the general population and that emotion-infused models improve the performance of stress models on underrepresented groups because of their effectiveness in low-data environments, and we propose that integrating emotions may benefit underrepresented groups in other mental health detection tasks.

Keywords—stress; emotion recognition; natural language processing; social networking.

I. INTRODUCTION

Psychological stress detection from social media posts has been identified as an important task for mental healthcare research [1], but the datasets for this task may not fairly represent all groups, and little prior work has investigated the effectiveness of psychological stress models on minority individuals.

This issue is especially relevant for Sexual and Gender Minority (SGM) people, who are more vulnerable to poor mental health outcomes than the general population. They are at higher risk of mental illnesses and suicide [2]–[4], and social media is often a place where SGM people find peers, seek help, and cope with prejudice [5]–[8].

One way to evaluate the ability of psychological stress models to understand the language of SGM individuals is through the related task of minority stress detection. Like psychological stress detection, minority stress detection uses Natural Language Processing (NLP) techniques to classify social media posts with whether the poster is experiencing stress [9][10]. However, minority stress is a psychosocial stress specific to minority individuals that they experience due to stigmatized social status [11]. An example of minority stress on social media is provided in Figure 1.

This task has an important application in improving the methodology of minority stress studies by circumventing limi-

At school, I have great friends and a good family at home. But I'm a closeted gay. If I ever came out, I know my friends would never talk to me again and my family would disown me. Because of this, I have zero motivation to come out.

I was kicked out of my online video game squad just for being gay (they said gays are pedophiles). My dad was outside my room listening when this happened so he grabbed me by the throat and kicked me out.

Figure 1. Examples of minority stress disclosure on social media from [9].

tations in survey-based self-reporting [12]. The systematic detection of minority stress can also be used to study large-scale health trends on social media that are not feasible to collect survey data on. Additionally, it has applications in automated intervention for those at risk of adverse consequences and screening for comorbid risks, such as cancer, HIV, and reduced cardiovascular health [13].

In this work, we evaluate the effectiveness of psychological stress models at detecting minority stress, and we hypothesize that a lack of diversity in the psychological stress training data causes stress models to overfit and be unable to generalize to minority individuals.

To address this issue, we experiment with the multi-task emotion-infused architectures introduced by [14]. They explored connections between emotions and stress in deep learning models, and they demonstrated that the task of emotion detection, which has more labeled data available, could improve the explainability of stress models.

Initial research found that emotion-infused models did not improve performance on the psychological stress detection task; however, we note that multi-task learning techniques (such as those used for the emotion-infused models) are known to improve generalization [15]. In this work, we explore using multi-task emotion-infused models to improve minority stress detection and highlight their potential for improving the performance of other mental health models on minorities. Our contributions in this work are as follows:

- We conduct experiments to demonstrate that traditional single-task psychological stress detection models under-

perform on minority stress and highlight how this performance difference risks widening preexisting healthcare disparities experienced by minority communities.

- We demonstrate that emotion-infused models reduce the performance gap and exceed State-Of-The-Art (SOTA) performance for the minority stress detection task without training on minority stress data.
- We provide explanatory analysis showing that minority communities have different distributions of emotions than the general population and that emotion-infused models improve the performance of stress models on underrepresented groups because of their effectiveness in low-data environments.

II. RELATED WORK

A. Psychological Stress Detection

Psychological stress detection is best studied with physiological data. Prior work has used audio [16], biological markers [17], neuroimaging [18], thermal imaging [19], or combinations of these signals [20] to achieve the most accurate forms of psychological stress detection. However, [1] demonstrated the value and feasibility of detecting psychological stress purely from social media text.

In [14], the authors introduced the use of emotion-infused models for psychological stress detection. These models improved explainability by integrating emotion with multi-task learning or fine-tuning; however, they did not significantly improve the performance of psychological stress detection.

The authors acknowledged demographic imbalances in the psychological stress dataset and noted a lack of language representing minority groups; however, little previous work has explored the performance of psychological stress models on minorities. In this work, we use the minority stress detection task to highlight the limitations of traditional architectures for detecting stress in minority individuals and explore the benefits of the previously introduced emotion-infused architectures for overcoming those limitations.

B. Minority Stress Detection

In [9], the authors introduced the use of NLP techniques for understanding minority stress. They wrote a codebook for identifying minority stress, created the first dataset of social media posts annotated for disclosure of minority stress, and introduced the first machine learning classifier for minority stress on social media. They experimented with using expertly engineered language features in combination with machine learning models to build a classifier. Their models are the current SOTA for minority stress detection.

Building off that work, [10] introduced a proof-of-concept Bidirectional Long Short-Term Memory (BI-LSTM) model to detect minority stress without expertly engineered features. They were the first to use deep learning for this task; however, they found limited results that did not outperform traditional machine learning models in detecting minority stress.

The authors viewed the minority stress detection task in isolation and trained models directly on the minority stress

dataset (which is too small for deep learning models). In this work, we understand minority stress as a subset of psychological stress, and we use this framework to improve minority stress detection by improving models from the related task of psychological stress detection.

C. Pretrained Language Models

Prior work has suggested that domain-specific pretrained language models may benefit mental healthcare tasks. In [21], the authors trained Bidirectional Encoder Representations from Transformers (BERT) and Robustly Optimized BERT Pretraining Approach (RoBERTa) models on a corpus of 13,671,785 sentences from mental health-related subreddits. These models, which the authors named MentalBERT and MentalRoBERTa, demonstrated improved performance on psychological stress detection. However, little previous work uses pretrained language models on minority stress detection. In this work, we experiment with four pretrained language models (BERT, RoBERTa, MentalBERT, and MentalRoBERTa) paired with psychological stress models for minority stress detection.

III. APPROACH

A. Baselines

For our baseline stress model, we use the pretrained BERT language model introduced in [22] followed by an additional dropout layer and dense classification layer. This architecture is the simplest that we evaluate, and it performed the best on the psychological stress detection task when it was introduced in [1]. Later architectures have not provided statistically significant performance improvements. We will refer to this model as Single-Task.

We compare our minority stress models to the current SOTA for minority stress detection established in [9]. This model is a Multilayer Perceptron (MLP) algorithm trained with GloVe word embeddings [23], Linguistic Inquiry and Word Count (LIWC) psycholinguistic categories [24], a gender and orientation hate speech lexicon, n-grams, sentiment classification [25], and stress classification [26]. This series of expertly engineered features draws information from a wide range of data sources that consider lexical and semantic aspects of the text, with a special focus on LGBTQ+ issues.

B. Emotion-Infused Models

We hypothesize that integrating emotions into psychological stress models will account for differences in the emotional expressions of minority individuals and improve the models' ability to generalize to minority stress. To test this claim, we evaluate the emotion-infused models introduced in [14]. When previously evaluated on psychological stress detection, these models did not provide significant performance improvements, but they improved explainability, and they represent key methods for using emotion in mental health tasks. There are three emotion-infused models.

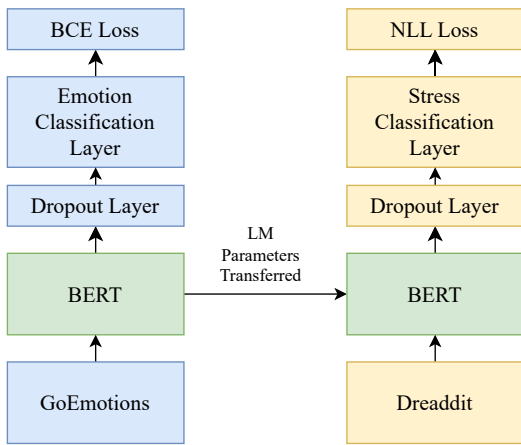


Figure 2. The architecture of Fine-Tune. Components that are used for both tasks are highlighted in green.

1) *Fine-Tuning Model (Fine-Tune)*: A visualization of Fine-Tune is shown in Figure 2. In this architecture, we first fine-tune a single-task model for emotion detection. Because that is a multi-label task, the model trains using Binary Cross-Entropy (BCE) loss. Then the language model parameters from that BERT model are transferred to another single-task model that is further fine-tuned for psychological stress detection. Because that is a single-label task, the model trains using Negative Log-Likelihood (NLL) loss. The rationale for this architecture is that the first task would give the BERT language model a better understanding of emotions, and that understanding would enable a more holistic representation of stress.

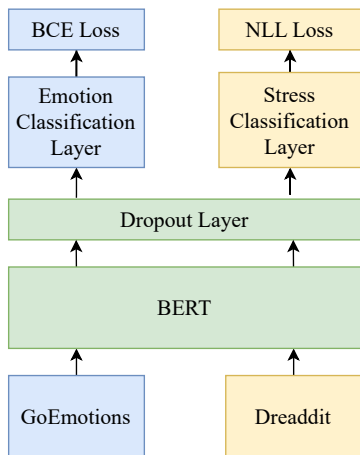


Figure 3. The architecture of Multi^{Alt}. Components that are used for both tasks are highlighted in green.

2) *Alternating Multi-Task Model (Multi^{Alt})*: A visualization of Multi^{Alt} is shown in Figure 3. It follows a similar rationale as Fine-Tune, but instead of training and then transferring a separate language model, it trains a single, shared language

model. During training, it alternates between training for emotion detection and psychological stress detection. Each training batch switches which task it is training for, but these different tasks share the same BERT representation layer. As in the Fine-Tune model, the emotion model trains with BCE loss and the psychological stress model trains with NLL loss.

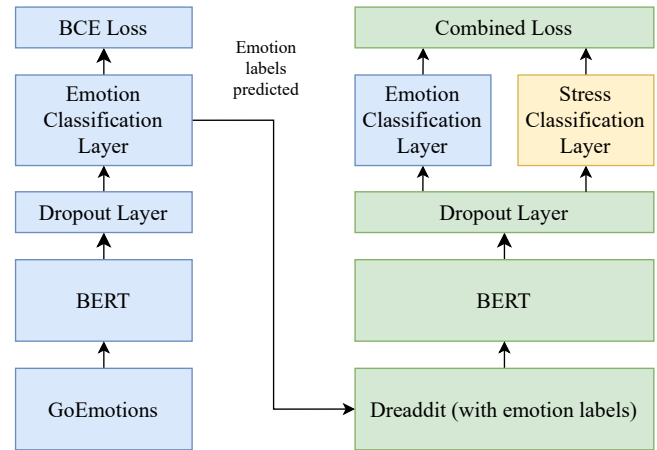


Figure 4. The architecture of Multi. Components that are used for both tasks are highlighted in green.

3) *Classical Multi-Task Model (Multi)*: A visualization of Multi is shown in Figure 4. It differs from the other two models because it uses a classical multi-task architecture that has the same input data for both tasks. However, because the stress data is not labeled with emotions, we first separately train a single-task model for emotion detection and use it to predict emotion labels for the stress data. The multi-task model then uses these emotion labels as targets for training the emotion detection task. In this model the loss is given by $\mathcal{L} = \lambda \mathcal{L}_{stress} + (1 - \lambda) \mathcal{L}_{emotion}$ where \mathcal{L}_{stress} is the NLL loss for psychological stress detection, $\mathcal{L}_{emotion}$ is the BCE loss for emotion detection, and λ is a weight parameter that we tune during model selection. The rationale for this model follows the traditional understanding that the inductive bias from the emotion detection task would improve its generalization.

C. Pretrained Language Models

Single-Task, Fine-Tune, Multi^{Alt}, and Multi were all originally introduced using BERT, but [21] demonstrated that their domain-specific pretrained language models, MentalBERT and MentalRoBERTa, improved performance on the psychological stress detection task. In this work, we evaluate four pretrained models tested in their work (BERT, RoBERTa, MentalBERT, and MentalRoBERTa) paired with each of the previously mentioned stress models.

IV. DATA

For psychological stress detection, we use Dreaddit [1], a dataset of 3,553 segments of Reddit posts from communities where stress is commonly disclosed. Each segment was labeled with whether the poster is expressing stress using

TABLE I
MINORITY STRESS PERFORMANCE

Model	BERT		RoBERTa		MentalBERT		MentalRoBERTa	
	F1	Accuracy	F1	Accuracy	F1	Accuracy	F1	Accuracy
Single-task	69.85	70.88	74.88	75.09	73.49	73.56	73.33	73.95
Fine-Tune	69.47	70.31	70.87	71.45	68.71	70.50	72.24	72.61
Multi ^{Alt}	70.95	71.45	70.60	70.69	73.58	74.52	71.88	72.41
Multi	75.55	68.58	75.16	69.35	75.58	72.99	78.53	74.52

Results of the models evaluated on minority stress detection with different pretrained language models. The best result under each metric is bolded.

TABLE II
PSYCHOLOGICAL STRESS PERFORMANCE

Model	BERT		RoBERTa		MentalBERT		MentalRoBERTa	
	F1	Accuracy	F1	Accuracy	F1	Accuracy	F1	Accuracy
Single-task	77.70	77.95	78.18	78.37	76.03	76.88	79.35	79.39
Fine-Tune	75.85	75.67	77.48	77.53	77.30	77.43	76.37	76.78
Multi ^{Alt}	78.76	78.97	77.00	77.25	80.80	80.89	79.42	79.59
Multi	77.53	77.36	79.00	79.16	77.90	78.27	79.86	79.91

Results of the models evaluated on psychological stress detection with different pretrained language models. The best result under each metric is bolded.

TABLE III
STRESS LABEL DISTRIBUTIONS

Dataset	Split	Stress	Non-Stress
Dreaddit	Training	1,110	1,012
	Development	374	342
	Testing	374	341
MStress	Development	72	103
	Testing	72	103

Label distributions for the training, development, and testing sets of Dreaddit and MStress.

crowdsourced annotation, requiring a majority vote from five annotators. We use this dataset to train the stress models and evaluate their performance on psychological stress detection.

To evaluate how the models generalize to minority stress detection, we use an existing dataset of 350 Reddit posts collected from LGBTQ+ communities by [9]. These posts were manually labeled by the authors with whether they contain the disclosure of minority stress using a codebook built based on Meyer’s Minority Stress model [11]. In this paper, we will refer to this dataset as MStress. Table III provides more information about the label distributions of Dreaddit and MStress.

Finally, to train the emotion-infused models, we use both Dreaddit and the GoEmotions dataset [27]. GoEmotions consists of 58,009 Reddit comments labeled by crowd workers with one of more than 27 emotions (or neutral). Based on the findings of [14] we use a relabeling of this dataset created with agglomerative clustering to cluster the original labels into the Ekman 6 basic emotions (anger, disgust, fear, joy, sadness, surprise, neutral) [28].

V. EXPERIMENTAL SETUP

Due to the scarcity of minority stress data, we do not train our models directly on MStress. Instead, all the models are trained on Dreaddit and GoEmotions (for the emotion-

infused models) with minibatch gradient descent using the Adam optimizer [29].

We test on Dreaddit to evaluate their psychological stress detection, and we test on MStress to evaluate their minority stress detection. For Dreaddit and GoEmotions, we use 60% of the data for training, 20% for hyperparameter tuning, and 20% for testing. We choose F1 score to evaluate the performance of our models because of its ability to account for class imbalances.

We replicate [14]’s training and hyperparameter tuning processes with the same parameter ranges; however, for our primary tests, we use 50% of MStress for hyperparameter tuning instead of using Dreaddit. We make this change to find the peak performance of the models for minority stress detection and highlight the associated reductions in psychological stress detection. We run each of these experiments three times with three different random seeds, and we report the mean of the three runs.

VI. RESULTS

We report the results of our primary tests when evaluated on minority stress in Table I, and we report the results of our primary tests when evaluated on psychological stress in Table II.

A. Single-Task Models

We find that the traditional models underperform on minority stress detection. The Single-Task models achieve F1 scores between 2.54 and 7.85 points lower on minority stress detection than on psychological stress detection. The MentalRoBERTa model performs best on psychological stress detection out of the Single-Task models with an F1 of 79.35, but that score drops to 73.33 when evaluated on minority stress detection.




Figure 5. Performance of the Single Task and Multi models trained with different proportions (10%, 25%, 50%, 75%, and 100%) of the original training set and evaluated on the psychological stress data.

B. Emotion-Infused Models

Though Fine-Tune and Multi^{Alt} do not provide significant performance improvements, we find that the Multi models improve minority stress detection performance over baselines in all cases. Multi combined with MentalRoBERTa achieves an F1 of 78.53 on minority stress detection. While this result is still lower than the best psychological stress models, it closes the gap significantly and outperforms 13 out of the 16 psychological stress models. This result demonstrates that using the Multi architecture improves the ability of psychological stress models to generalize to minority stress detection.

Importantly, our best Multi models are not trained directly on minority stress data, but they outperform the current SOTA MLP for minority stress detection. The MLP proposed by [9], which is trained directly on minority stress data, achieved an F1 of 75, and our best Multi model achieves an F1 score of 78.53.

C. Pretrained Language Models

We find that in most cases, domain-specific language models such as MentalBERT and MentalRoBERTa provide marginal improvements in both psychological stress detection and minority stress detection compared to the standard BERT and RoBERTa models. This result confirms prior work demonstrating that the MentalBERT and MentalRoBERTa language models perform better on a variety of mental healthcare tasks [21].

D. Discussion

These findings have important implications for the use of stress models in research and healthcare applications. First, traditional Single-Task models perform worse on minority stress detection than psychological stress detection and risk reinforcing preexisting mental healthcare disparities for SGM individuals. Second, the Multi architecture creates models that can generalize well to minority stress detection and significantly reduce the performance gap. Finally, minority stress researchers can benefit from using psychological stress detection data to surpass the current SOTA without directly training on minority stress data.

In the next section, we support these conclusions with an analysis of the Single-Task and Multi models for psychological stress detection with reduced training sets.

VII. ANALYSIS

A. Data Reduction Analysis

We propose that the disparity in performance of the baseline models between minority stress and psychological stress is due to overfitting on the psychological stress data. The Single-Task models gain too much sample-specific information and, as a result, are struggling to perform well on out-of-sample stress disclosures like minority stress.

Multi-task learning techniques improve generalization by using domain information contained in related tasks as an inductive bias [15]. We hypothesize that this improvement in generalization explains why the Multi models have improved performance on minority stress detection compared to baselines.

To further support this reasoning, we experiment with reducing the size of the Dreddit training set for the psychological stress detection task. This reduction in the training set simulates the data scarcity that is present for minority stress detection.

We perform the same experimental setup as described in Section V, but we use psychological stress data for our hyperparameter tuning and change the size of the training set to be either 10%, 25%, 50%, 75%, or 100% of the original training set. We perform these experiments with the Single-Task and Multi models paired with each of the four language models (BERT, RoBERTa, MentalBERT, and MentalRoBERTa).

We report our results in Figure 5. We find that while Single-Task and Multi achieve equivalent performance with the full training set, reduced training sets reduce the performance of Single-Task models much more significantly than Multi models. We see that at 100% the Single-Task models all have F1 scores near 80, but at 50% they drop to be between 77.38 and 78.13. By comparison, the Multi models have F1 scores of at least 80 with only 50% of the training data, but they do not significantly improve as the training size increases.

This finding demonstrates that the improved generalizability of the multi-task architecture of the models makes them more effective in low-data environments. It consequently explains why Multi models are more effective at minority stress detection: the training set has a limited amount of minority stress, so detecting it is a low-data environment.

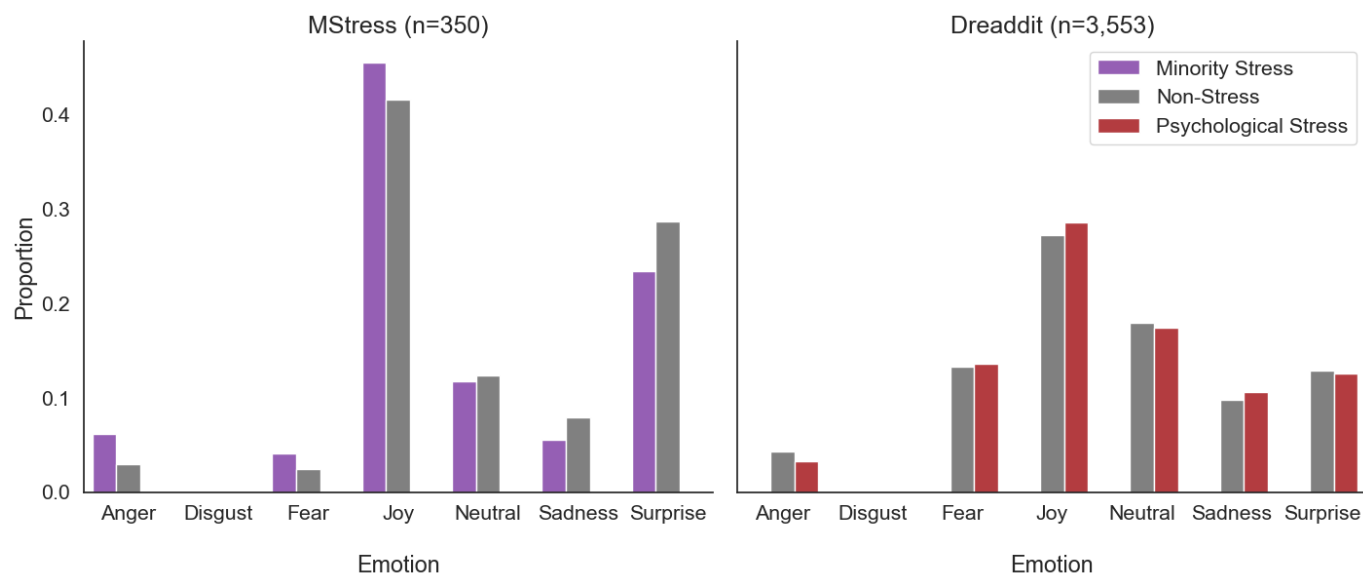


Figure 6. Distributions of predicted emotion labels in MStress and Dreddit. Note that posts can have multiple emotion labels.

This finding also suggests that multi-task emotion-infused architectures may improve stress detection for other underrepresented groups, and further work should be done to explore using emotions to create equitable mental health models.

B. Emotion Distributions

To provide additional support for the importance of emotion analysis for supporting underrepresented groups, we examine the predicted emotion distributions of both MStress and Dreddit (shown in Figure 6). These emotion labels were created using a single-task MentalRoBERTa model with a macro F1 of 61.13.

From these distributions, we see that emotions do not significantly vary based on stress status. Posts in MStress marked as minority stress have a similar emotion distribution as posts marked as non-stress, and posts in Dreddit marked as psychological stress have a similar emotion distribution as posts marked as non-stress.

However, we see a significant difference in the emotion distributions of posts from MStress compared to Dreddit. This suggests that posts from minority communities exhibit different emotions, and that difference may affect single-task models' ability to understand mental health conditions in minorities. Other work has found that minority stress mediates emotion regulation, leading to dysregulation and emotion suppression [30], which supports this finding.

This difference provides further explanation for the underperformance of single-task stress models on minorities. They are trained on a distribution of emotional expressions that are not representative of minority communities. This finding further suggests that multi-task, emotion-infused architectures may make more equitable models for other mental health tasks.

VIII. CONCLUSION

In this work, we find that traditional single-task stress models underperform on minority stress detection and are at risk of widening the healthcare gap for SGM individuals. We also find that risk can be reduced with the use of a multi-task architecture that integrates the task of emotion detection. Our experiments show that architecture performs well on psychological stress detection and outperforms the SOTA for minority stress detection without training on minority stress data. Finally, we provide explanatory analysis demonstrating the Multi model's superior performance in low-data environments, and we highlight how differences in emotion expressions in minority communities make them vulnerable to reduced effectiveness in mental health modeling. Our analysis suggests that integrating emotions may be effective for improving the performance of mental health models on underrepresented groups, and future work should explore using emotions to create equitable models for other mental health tasks.

IX. LIMITATIONS

Our data was collected from Reddit, which is not necessarily representative of the general population. While our work focuses on generalizing stress models to work effectively on minority stress, it does not evaluate their ability to generalize to other social media platforms or beyond social media.

Additionally, our models focus on detecting minority stress for sexual and gender minorities, but we did not explore the detection of minority stress for other underrepresented groups such as racial and ethnic minorities. Other minority groups may disclose minority stress differently, and future work should evaluate the performance of current minority stress models on other minority groups.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Award # OIA-1946391, "Data Analytics that are Robust and Trusted (DART)".

REFERENCES

- [1] E. Turcan and K. McKeown, "Dreaddit: A Reddit dataset for stress analysis in social media," in *Proceedings of the Tenth International Workshop on Health Text Mining and Information Analysis (LOUHI 2019)*, Hong Kong: Association for Computational Linguistics, Nov. 2019, pp. 97–107. DOI: 10.18653/v1/D19-6213.
- [2] R. M. Díaz, G. Ayala, E. Bein, J. Henne, and B. V. Marin, "The impact of homophobia, poverty, and racism on the mental health of gay and bisexual Latino men: Findings from 3 US cities.," *American Journal of Public Health*, vol. 91, no. 6, pp. 927–932, Jun. 2001, ISSN: 0090-0036.
- [3] G. Remafedi, S. French, M. Story, M. D. Resnick, and R. Blum, "The relationship between suicide risk and sexual orientation: Results of a population-based study.," *American Journal of Public Health*, vol. 88, no. 1, pp. 57–60, Jan. 1998, ISSN: 0090-0036.
- [4] E. A. Tebbe and B. Moradi, "Suicide risk in trans populations: An application of minority stress theory," eng, *Journal of Counseling Psychology*, vol. 63, no. 5, pp. 520–533, Oct. 2016, ISSN: 0022-0167. DOI: 10.1037/cou0000152.
- [5] E. Formby, *Exploring LGBT Spaces and Communities: Contrasting Identities, Belongings and Wellbeing*, en. Routledge, 2017, Google-Books-ID: GEccswEACAAJ, ISBN: 978-1-138-81400-4.
- [6] S. Tropiano, "'A Safe and Supportive Environment': LGBTQ Youth and Social Media," en, in *Queer Youth and Media Cultures*, C. Pullen, Ed., London: Palgrave Macmillan UK, 2014, pp. 46–62, ISBN: 978-1-349-48056-2 978-1-137-38355-6. DOI: 10.1057/9781137383556_4.
- [7] L. B. McInroy and S. L. Craig, "'It's like a safe haven fantasy world': Online fandom communities and the identity development activities of sexual and gender minority youth," *Psychology of Popular Media*, vol. 9, pp. 236–246, 2020, Place: US Publisher: Educational Publishing Foundation, ISSN: 2689-6575. DOI: 10.1037/ppm0000234.
- [8] N. Woznicki, A. S. Arriaga, N. A. Caporale-Berkowitz, and M. C. Parent, "Parasocial relationships and depression among LGBQ emerging adults living with their parents during COVID-19: The potential for online support," *Psychology of Sexual Orientation and Gender Diversity*, vol. 8, pp. 228–237, 2021, Place: US Publisher: Educational Publishing Foundation, ISSN: 2329-0390. DOI: 10.1037/sgd0000458.
- [9] K. Saha *et al.*, "The Language of LGBTQ+ Minority Stress Experiences on Social Media," *Proceedings of the ACM on human-computer interaction*, vol. 3, no. CSCW, p. 89, Nov. 2019, ISSN: 2573-0142. DOI: 10.1145/3361108.
- [10] C. J. Cascalheira *et al.*, "Classifying Minority Stress Disclosure on Social Media with Bidirectional Long Short-Term Memory," en, *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 16, pp. 1373–1377, May 2022, ISSN: 2334-0770. DOI: 10.1609/icwsm.v16i1.19390.
- [11] I. H. Meyer, "Prejudice, Social Stress, and Mental Health in Lesbian, Gay, and Bisexual Populations: Conceptual Issues and Research Evidence," *Psychological bulletin*, vol. 129, no. 5, pp. 674–697, Sep. 2003, ISSN: 0033-2909. DOI: 10.1037/0033-2909.129.5.674.
- [12] P. P. Heppner, B. E. Wampold, J. Owen, K. T. Wang, and M. N. Thompson, *Research design in counseling*. Boston, MA: Cengage Learning, 2016, ISBN: 978-1-305-08731-6.
- [13] A. Flentje, N. C. Heck, J. M. Brennan, and I. H. Meyer, "The relationship between minority stress and biological outcomes: A systematic review," eng, *Journal of Behavioral Medicine*, vol. 43, no. 5, pp. 673–694, Oct. 2020, ISSN: 1573-3521. DOI: 10.1007/s10865-019-00120-6.
- [14] E. Turcan, S. Muresan, and K. McKeown, "Emotion-infused models for explainable psychological stress detection," in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Online: Association for Computational Linguistics, Jun. 2021, pp. 2895–2909. DOI: 10.18653/v1/2021.naacl-main.230.
- [15] R. Caruana, "Multitask Learning," en, *Machine Learning*, vol. 28, no. 1, pp. 41–75, Jul. 1997, ISSN: 1573-0565. DOI: 10.1023/A:1007379606734.
- [16] X. Zuo, T. Li, and P. Fung, "A multilingual natural stress emotion database," in *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey: European Language Resources Association (ELRA), May 2012, pp. 1174–1178.
- [17] A. P. Allen, P. J. Kennedy, J. F. Cryan, T. G. Dinan, and G. Clarke, "Biological and psychological markers of stress in humans: Focus on the Trier Social Stress Test," eng, *Neuroscience and Biobehavioral Reviews*, vol. 38, pp. 94–124, Jan. 2014, ISSN: 1873-7528. DOI: 10.1016/j.neubiorev.2013.11.005.
- [18] F. Al-Shargie *et al.*, "Mental stress assessment using simultaneous measurement of EEG and fNIRS," *Biomedical Optics Express*, vol. 7, no. 10, pp. 3882–3898, Sep. 2016, ISSN: 2156-7085. DOI: 10.1364/BOE.7.003882.
- [19] S. Kumar *et al.*, *StressNet: Detecting Stress in Thermal Videos*, arXiv:2011.09540 [cs, eess], Nov. 2020. DOI: 10.48550/arXiv.2011.09540.
- [20] M. Jaiswal *et al.*, "MuSE: A multimodal dataset of stressed emotion," English, in *Proceedings of the Twelfth Language Resources and Evaluation Conference*, Marseille, France: European Language Resources Association, May 2020, pp. 1499–1510, ISBN: 979-10-95546-34-4.
- [21] S. Ji *et al.*, "MentalBERT: Publicly available pretrained language models for mental healthcare," in *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, Marseille, France: European Language Resources Association, Jun. 2022, pp. 7184–7190.
- [22] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "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)*, Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 4171–4186. DOI: 10.18653/v1/N19-1423.
- [23] J. Pennington, R. Socher, and C. Manning, "GloVe: Global vectors for word representation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar: Association for Computational Linguistics, Oct. 2014, pp. 1532–1543. DOI: 10.3115/v1/D14-1162.
- [24] Y. R. Tausczik and J. W. Pennebaker, "The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods," en, *Journal of Language and Social Psychology*, vol. 29, no. 1, pp. 24–54, Mar. 2010, ISSN: 0261-927X, 1552-6526. DOI: 10.1177/0261927X09351676.
- [25] C. Manning *et al.*, "The Stanford CoreNLP natural language processing toolkit," in *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, Baltimore, Maryland: Association for Computational Linguistics, Jun. 2014, pp. 55–60. DOI: 10.3115/v1/P14-5010.

- [26] K. Saha and M. De Choudhury, "Modeling Stress with Social Media Around Incidents of Gun Violence on College Campuses," *Proceedings of the ACM on Human-Computer Interaction*, vol. 1, no. CSCW, 92:1–92:27, Dec. 2017. DOI: 10.1145/3134727.
- [27] D. Demszky *et al.*, "GoEmotions: A dataset of fine-grained emotions," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Online: Association for Computational Linguistics, Jul. 2020, pp. 4040–4054. DOI: 10.18653/v1/2020.acl-main.372.
- [28] P. Ekman, "Are there basic emotions?" eng, *Psychological Review*, vol. 99, no. 3, pp. 550–553, Jul. 1992, ISSN: 0033-295X. DOI: 10.1037/0033-295x.99.3.550.
- [29] D. P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, 2017. arXiv: 1412.6980 [cs.LG].
- [30] A. Singh, A. Dandona, V. Sharma, and S. Z. H. Zaidi, "Minority stress in emotion suppression and mental distress among sexual and gender minorities: A systematic review," en, *Ann. Neurosci.*, vol. 30, no. 1, pp. 54–69, Jan. 2023.