OPEN ACCESS JOURNALS

Indiana Journal of Multidisciplinary Research
Abbreviated Tittle- Ind J Multi Res
ISSN *(Online)*- 2583-3820
Journal Homepage Link- https://indianapublications.com/Journals/IJMR
DOI: https://doi.org/10.5281/zenodo.21560165

INDIANA PUBLICATIONS
PRODUCTIVE AND QUALITY RESEARCH

Research Article

Volume-06|Issue-04|2026

Domain-Adaptive Stressor Identification Using Fine-Tuned BERT and NER Techniques

Likitha A. M[1], Prem Kumar G. Raga[2], Venkat Baba Yemineni[3], Shanthala Kollur[4], Haridas Nayak[5]

[1,2,3,4]RV Institute of Technology and Management, Bangalore, Karnataka, India.
[5]Department of ME, PES University EC Campus, Bengaluru, India.

Article History
Received: 05.05.2026
Accepted: 22.06.2026
Published: 25.07.2026

Citation
Likitha, A. M., Raga, P. K. G., Yemineni, V. B., Kollur, S. & Nayak, H. (2026). Domain-Adaptive Stressor Identification Using Fine-Tuned BERT and NER Techniques. *Indiana Journal of Multidisciplinary Research, 6*(4), 95-99.

Abstract: Stressor identification is essential for understanding psychological distress as ex-pressed in mental health narratives, where language is often implicit, subjective, and varied. This study proposes a domain-adaptive Named Entity Recognition (NER) framework based on a fine-tuned Bidirectional Encoder Representations from Transformers (BERT) model, specifically tailored for stressor entity extraction. Trained on an annotated corpus using the BIO(Beginning-Inside-Outside) tagging scheme, the model achieved an F1-score of 0.72 for the STRESSOR category. Comprehensive evaluations, including confusion matrix analysis, per-class performance metrics, and training loss visualizations demonstrated the model's effectiveness in recognizing nuanced, context-dependent stressor expressions. Theapproach offers amodularpipeline applicable to mental health monitoring, conversational agents, and digital therapeutic tools, contributing to scalable and interpretable AI solutions in psychological well-being.

Keywords: NER, Psychological Stressors, BERT, Mental Health NLP, Domain Adaptation

INTRODUCTION

Mental health disorders are a pressing global concern, with stress-related conditions affecting individuals across diverse age groups. The growing use of digital platforms such as social media, online communities, and chat applications has made them a key medium through which people express emotional struggles. Analyzing such in-formal, user-generated texts can enable timely mental health support by identifying signs of psychological distress.

A critical step in this process is recognizing stressors, specific events, conditions, or experiences that contribute to mental strain. However, stressor detection remains challenging due to the subtle, context-rich, and linguistically diverse nature of person-al expression. Additionally, there is a notable lack of annotated corpora tailored to stressor recognition in mental health contexts. Traditional Named Entity Recognition (NER) systems, which perform well in domains like news and medicine, often falter when applied to informal mental health discourse due to its implicit and subjective nature [1-6].

Transformer-based models such as Bidirectional Encoder Representations from Transformers (BERT) have set new benchmarks in NER by leveraging deep contextual embeddings [7]. However, domain adaptation is essential for these models to effectively capture the nuanced language of psychological stressors [8]. Our study proposes a fine-tuned BERT based architecture specifically designed to extract stressor entities from psychological texts, enabling applications in triaging, therapeutic assistants, and emotion aware digital tools [9].

While previous work has focused on sentence level classification detecting stress or depressive states, token-level sequence labelling for stressor mentions remains under-explored. Emerging models like Mental XLNet and domain-adapted versions of BERT show promise, but are primarily evaluated at sentence or document level [10-11]. Similarly, models like Bio BERT and ClinicalBERT, though successful in biomedical NER, are rarely optimized for psychological data [1-2,12].

Recent advancements in low-resource NER include adapter tuning, prompt-based learning, label attention, and lexicon-guided methods, some of which have shown success in clinical NLP [8,13-14]. Mental health-specific lexicons like MentalLex and model variants such as RoBERTa with label attention have further improved contextual recognition [10]. However, fine-grained, token-level stressor recognition using structured evaluation methods such as confusion matrices and per-class metrics re-mains scarce in the literature.

This paper addresses that gap by presenting a BERT based NER model fine-tuned on a custom BIO (Beginning-Inside-Outside) tagged dataset, annotated for stressor phrases across categories like academic stress, relationship issues, and workplace

***Corresponding Author:** Shanthala Kollur

anxiety. The pipeline demonstrates strong performance across standard evaluation metrics, offering a robust foundation for scalable, interpretable stressor identification in real-world mental health monitoring systems.

METHODOLOGY

This study proposes a domain-adaptive Named Entity Recognition (NER) approach for extracting psychological stressors from English-language narratives. The model is fine-tuned on a transformer-based architecture, specifically, BERT to learn token-level contextual representations and accurately label stressor entities. The workflow consists of data preprocessing, manual annotation, model training, and evaluation using standard NER metrics.

Dataset and Annotation

The dataset was obtained from a public Kaggle repository [15], consisting of anonymized user-generated sentences reflecting emotional stress. These include experiences related to academic pressure, career anxiety, and interpersonal conflicts.

The dataset was manually annotated following the standard BIO tagging scheme, where each token in the text was labeled as either the beginning of a stressor entity (B-STRESSOR), a continuation of that entity (I-STRESSOR), or outside any entity (O). For instance, in the sentence “I am stressed about exams,” the word *exams* would be tagged as B-STRESSOR. The annotation process was carefully guided by established psychological stress frameworks to ensure conceptual clarity. To maintain consistency and reliability, multiple validation rounds were conducted, helping to refine the labels and achieve strong inter-annotator agreement. The class distribution is provided in **Table 1.**, highlighting the natural imbalance typical of real-world NER datasets.

Table 1. Label Distribution

Label	Token Count
O	**17,300**
B-STRESSOR	**585**
I-STRESSOR	**404**

Model Architecture

BERT-base-uncased model was adopted from the Hugging Face Transformers library. Input sentences were tokenized using Word Piec encoding, preserving sub word information. A token classification head was added on top of the encoder to predict one of the three labels (O, B-STRESSOR, I-STRESSOR) for each token. This configuration enables effective modelling of semantic dependencies and supports robust recognition of stressor phrases across varied sentence structures. A schematic representation of the model pipeline is shown in Fig.1.

Fine-Tuning Strategy

Model fine-tuning was performed using the Hugging Face Trainer API. Only non-padding tokens were considered in loss computation to ensure clean gradient updates. The training parameters used are summarized in **Table 2**.

Table 2. Training Configuration

Hyperparameter	Value
Epochs	3
Batch Size	8
Learning Rate	5e-5
Optimizer	AdamW
Loss Function	Cross-entropy with masking

Evaluation Protocol and Metrics

To assess the performance of the model, evaluation was carried out on a stratified holdout test set using both token-level and entity-level metrics. The evaluation relied on the widely used SeqEval and scikit-learn libraries. Precision was measured as the proportion of predicted entities that correctly matched the annotated ground truth, while recall captured the proportion of true entities that the model successfully identified. The F1-score, serving as the harmonic mean of precision and recall, provided a balanced view of the model’s accuracy. To gain deeper insights into the model’s prediction behavior, a token-level confusion matrix was also generated. This helped in identifying specific error patterns, such as misclassification between B-, I-, and O-labels, or inconsistencies in capturing complete entity spans.

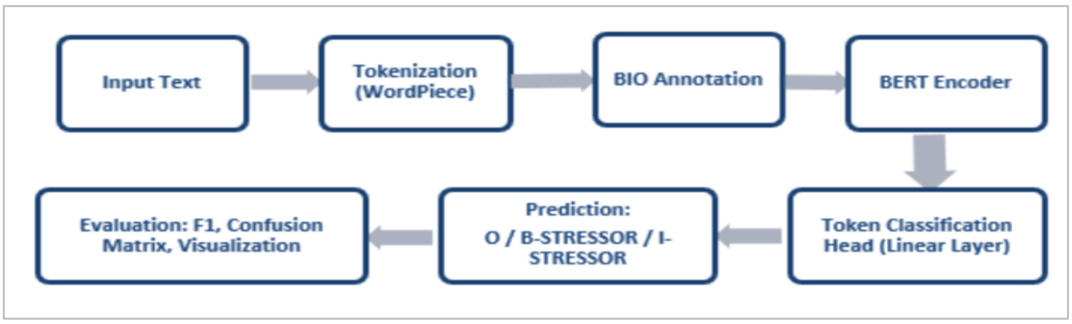

Fig. 1. Schematic of the BERT-based architecture for stressor identification

***Corresponding Author:** Shanthala Kollur

RESULTS AND DISCUSSION

To assess the effectiveness of the BERT-based NER model in recognizing psychological stressor entities token-level and class-wise evaluations were performed in detail. The results demonstrate that the model generalizes well to diverse narrative forms, maintaining both contextual accuracy and span level integrity.

Token-Level Performance

The model's performance was assessed on a held-out test set following the BIO tagging scheme, focusing specifically on the STRESSOR class, which includes both beginning (B-) and inside (I-) tags. It achieved a precision of 0.68, a recall of 0.76, and an overall F1-score of 0.72. The relatively high recall suggests that the model is adept at detecting stressor mentions, even when they appear in informal, conversational, or complex sentence structures. On the other hand, the moderate precision indicates a cautious prediction strategy, one that avoids excessive false positives. This balance is particularly important in mental health applications, where over identifying stress cues could lead to unnecessary concern or intervention fatigue

Confusion Matrix Analysis

To better understand how the model distinguishes between different types of labels, a normalized confusion matrix was generated (Fig. 2). It provides a clear view of how often each labels, B-STRESSOR, I-STRESSOR, and O was correctly predicted or confused with others. The model performed especially well in identifying non-entity (O) tokens, which make up the majority of the dataset, showing its ability to reliably filter out irrelevant words. However, some confusion was observed between the B-STRESSOR and I-STRESSOR tags. This suggests that while the model recognizes stressor entities, it occasionally struggles with pinpointing the exact beginning and continuation of multi-word phrases, which is a common challenge in sequence labeling tasks.

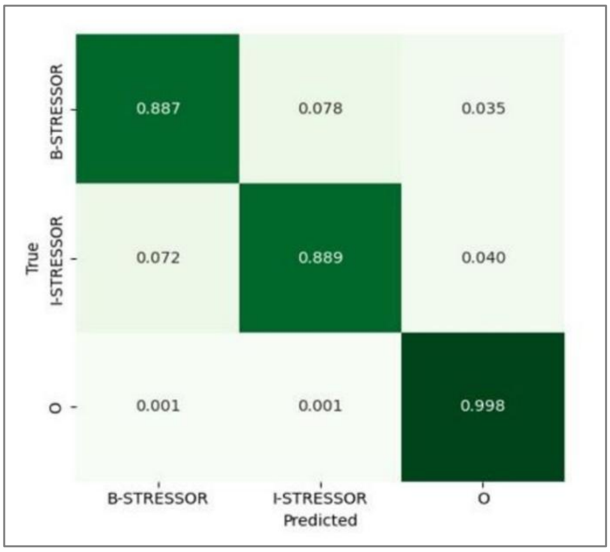

Fig. 2. Normalized confusion matrix for token-level predictions

Per-Class Metrics

A closer look at the per-class performance (shown in Fig. 3) reveals how well the model handles each label type. For B-STRESSOR tokens, the model achieved a high recall of 0.885, indicating its strength in detecting the starting point of stressor entities. However, the slightly lower precision suggests that it occasionally misclassifies the beginning of a stressor. For I-STRESSOR tokens, both precision and recall exceeded 0.86, showing that the model reliably tracks the continuation of stressor phrases across multiple tokens. The O class, which represents non-stressor tokens, achieved an almost perfect F1-score of 0.996, highlighting the model's ability to accurately ignore irrelevant or neutral content. Together, these results confirm the model's capacity to identify complete stressor spans with reasonable accuracy, which is especially important in applications that rely on extracting meaningful segments from mental health narratives.

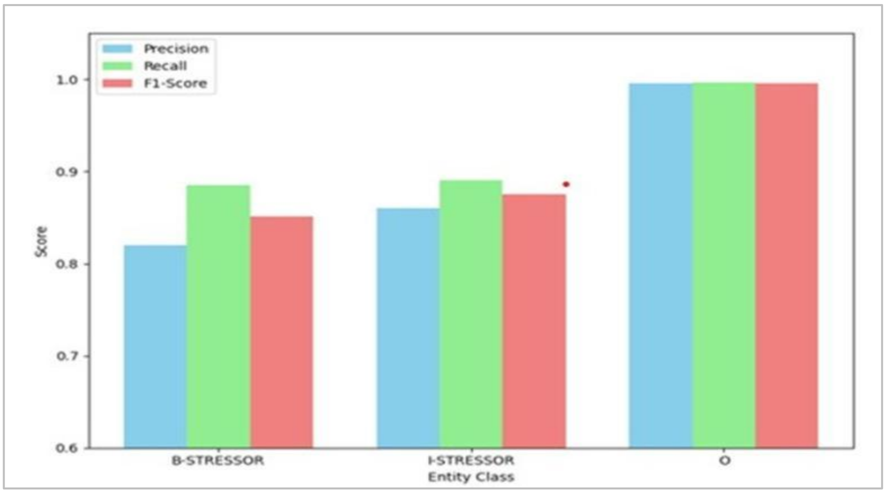

Fig. 3. Per Class Performance Metrices

***Corresponding Author:** Shanthala Kollur

Error Analysis

While the model showed strong overall performance, a detailed error analysis revealed three recurring challenges. First, false positives were observed where emotionally intense but contextually neutral words such as "tired" or "angry"—were incorrectly identified as stressor entities. This likely stems from their semantic closeness to actual stress cues, leading the model to overextend its boundaries. Second, false negatives occurred in instances where the stressor was expressed indirectly or spread across multiple words. For example, phrases like "expectations at home" were sometimes missed entirely, suggesting the model struggles with abstract or figurative expressions of stress. Lastly, boundary errors were common in longer entity spans, where the model occasionally confused the beginning (B-tag) and continuation (I-tag) labels. This was evident in cases like "college applications," where each word was tagged as a separate entity rather than a unified phrase. Addressing these errors in future work may require deeper contextual modeling, such as incorporating semantic role labeling or discourse-level features.

Training Dynamics

Understanding training dynamics is crucial for assessing a model's convergence behaviour, generalization ability, and robustness. In this study, the domain adaptive BERT based NER model was fine-tuned for three epochs on a curated dataset of stressor-annotated narratives using the Hugging Face Trainer API. Training loss was monitored throughout to evaluate learning progress.

As shown in Fig. 4, the loss declined smoothly and consistently, converging to a final value of 0.0477. This steady reduction indicates effective learning from the training data without signs of instability, divergence, or oscillation. The absence of anomalies suggests that the optimization process was stable and that the hyperparameter learning rate (5e-5) and batch size (8) were appropriately chosen to ensure both computational efficiency and convergence reliability. Importantly, no signs of overfitting were observed, and performance on the test set improved steadily across epochs. This underscores the advantage of domain-adaptive fine-tuning and the strength of pre-trained contextual embeddings, which together helped the model generalize well to unseen examples while requiring minimal training epochs. The observed convergence trend lends credibility to the model's downstream performance across token-level and class-wise evaluations. As such, the training loss curve serves as a reliable diagnostic tool and supports the validity of later results such as confusion matrices and per-class F1 scores.

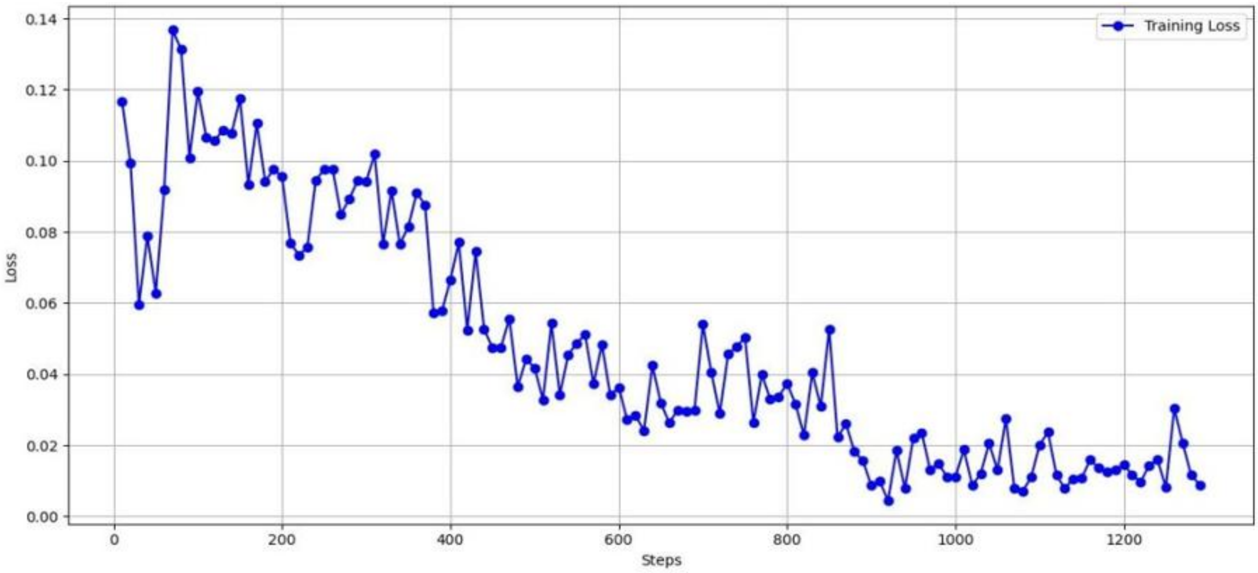

Fig. 4. Training Loss Curve of the Fine-Tuned BERT-Based NER Model

CONCLUSION

This study introduces a domain-adaptive BERT-based Named Entity Recognition approach for identifying psychological stressors in user generated mental health narratives. Trained on a curated dataset using BIO tagging, the model achieves a strong F1-score of 0.72 for stressor recognition.

The model's high recall and span level precision demonstrate its ability to capture nuanced, context dependent stressor mentions embedded in informal, emotionally expressive text. Confusion matrix and class-wise evaluations further confirm its dis-criminative reliability.

The approach presents a modular pipeline that can support applications in mental health analytics, intelligent chatbots, and digital therapeutic interventions.

***Corresponding Author:** Shanthala Kollur

REFERENCES

1. Lee, J., Yoon, W., Kim, S., *et al.* (2020). BioBERT: A pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics, 36*(4), 1234–1240.
2. Alsentzer, E., Murphy, J. R., Boag, W., Weng, W.-H., Jin, D., Naumann, T., & McDermott, M. (2019). Publicly available clinical BERT embeddings. *Proceedings of the 2nd Clinical Natural Language Processing Workshop (ACL)*, 72–78.
3. Nischall, K. S. N. S., Sai, G. N., Mathew, C., Gowda, G. C. S., & Chandrakala, B. M. (2020). A survey on recognition of handwritten ZIP codes in a postal sorting system. *International Research Journal of Engineering and Technology (IRJET), 7*(3).
4. Chandrakala, B. M., & Linga Reddy, S. C. (2019). Proxy re-encryption using MLBC (modified lattice-based cryptography). *International Conference on Recent Advances in Energy-efficient Computing and Communication (ICRAECC)*, 1–5.
5. Ramesh, S., Tiwari, A., Choubey, P., Kashyap, S., Khose, S., Lakara, K., Singh, N., & Verma, U. (2021). BERT based transformers lead the way in extraction of health information from social media. *Proceedings of SMM4H Workshop*, 33–38.
6. Litake, O., Sabane, M., Patil, P., Ranade, A., & Joshi, R. (2022). Mono vs multilingual BERT: A case study in Hindi and Marathi named entity recognition. *arXiv preprint*, 2203.12907.
7. Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT*, 4171–4186.
8. Lample, G., Ballesteros, M., Subramanian, S., Kawakami, K., & Dyer, C. (2016). Neural architectures for named entity recognition. *Proceedings of NAACL-HLT*, 260–270.
9. Tang, B., Cao, H., Feng, X., *et al.* (2022). Domain-specific named entity recognition with Adapter-BERT in clinical narratives. *Journal of the American Medical Informatics Association, 29*(6), 1048– 1057.
10. Ji, S., Zhang, T., Ansari, L., Fu, J., Tiwari, P., & Cambria, E. (2022). MentalBERT: Publicly available pre-trained language models for mental healthcare. *Proceedings of the 13th Language Resources and Evaluation Conference (LREC)*, 7184–7190.

***Corresponding Author:** Shanthala Kollur