



MISE: Meta-knowledge Inheritance for Social Media-Based Stressor Estimation

Xin Wang
University of Oxford
Oxford, United Kingdom
xin.wang@eng.ox.ac.uk

Ling Feng
Tsinghua University
Beijing, China
fengling@tsinghua.edu.cn

Huijun Zhang
China Huaneng
Beijing, China
hj_zhang@qny.chng.com.cn

Lei Cao
Beijing Normal University
Beijing, China
caolei@bnu.edu.cn

Kaisheng Zeng
Tsinghua University
Beijing, China
zks19@mails.tsinghua.edu.cn

Qi Li
Beijing Normal University
Beijing, China
liqi2018@bnu.edu.cn

Yang Ding
Tsinghua University
Beijing, China
dingy20@mails.tsinghua.edu.cn

Yi Dai
Tsinghua University
Beijing, China
dai-y21@mails.tsinghua.edu.cn

David Clifton
University of Oxford
Oxford, United Kingdom
Oxford Suzhou Centre for Advanced
Research
Suzhou, China
david.clifton@eng.ox.ac.uk

Abstract

Stress haunts people in modern society, which may cause severe health issues if left unattended. With social media becoming an integral part of daily life, leveraging social media to detect stress has gained increasing attention. While the majority of the work focuses on classifying stress states and stress categories, this study introduces a new task aimed at estimating more specific stressors (like exam, writing paper, etc.) through users' posts on social media. Unfortunately, the diversity of stressors with many different classes but a few examples per class, combined with the consistent arising of new stressors over time, hinders the machine understanding of stressors. To this end, we cast the stressor estimation problem within a practical scenario few-shot learning setting, and propose a novel meta-learning based stressor estimation framework that is enhanced by a meta-knowledge inheritance mechanism. This model can not only learn generic stressor context through meta-learning, but also has a good generalization ability to estimate new stressors with little labeled data. A fundamental breakthrough in our approach lies in the inclusion of the meta-knowledge inheritance mechanism, which equips our model with the ability to prevent catastrophic forgetting when adapting to new stressors. The experimental results show that our model achieves state-of-the-art performance compared with the baselines. Additionally, we construct a social media-based stressor estimation dataset that can help train artificial intelligence models to facilitate human well-being.

CCS Concepts

• **Applied computing** → **Psychology**; • **Human-centered computing** → **Social media**.

Keywords

stressor estimation, social media, meta-knowledge inheritance

ACM Reference Format:

Xin Wang, Ling Feng, Huijun Zhang, Lei Cao, Kaisheng Zeng, Qi Li, Yang Ding, Yi Dai, and David Clifton. 2025. MISE: Meta-knowledge Inheritance for Social Media-Based Stressor Estimation. In *Proceedings of the ACM Web Conference 2025 (WWW '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3696410.3714901>

1 Introduction

With the rapid development of economy and society, people are under unprecedented psychological stress, coming from various aspects of life. As excessive stress without timely relief can negatively affect people's thoughts, feelings, behaviors, and physical and mental health [10], estimating and managing stress have become a big issue in the contemporary society.

Beyond traditional counseling and questionnaires based stress detection methods [3, 6, 12, 17], leveraging social media for stress detection has gained considerably increased attention in recent years. Through analyzing people's free-styled linguistic expressions and social behaviors, it is feasible to automatically and timely detect stress.

So far, the majority of social media-based stress detection work focuses on classifying user's **stress state** (i.e., stressed or non-stressed) [13, 14, 46, 55] and **stress category** (i.e., stress is classified into several broader categories, such as school, work, financial state, etc.) [22, 24, 54].

The aim of this study is to go further and identify users' specific **stressors**. We argue that in order to provide effective treatment for



This work is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.

WWW '25, Sydney, NSW, Australia

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1274-6/25/04

<https://doi.org/10.1145/3696410.3714901>

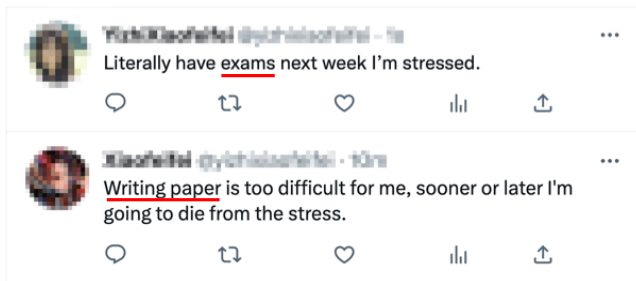


Figure 1: Two posts. The first user’s stress is caused by exams. The second user’s stress is caused by writing paper. While prior studies may classify both posts as belonging to the broader category of school-related stress, our work aims to estimate the specific causes of stress (i.e., exams and writing paper), in order to provide more targeted support for stress relief.

stress relief, it is necessary to have a thorough understanding of the specific causes of stress. Psychological research has demonstrated that stress arises from stressors [25]. Figure 1 shows two posts, from which we can know that the first user’s stress comes from *exams*, and the second user’s stress comes from *writing paper*. Accurately identifying stressors such as *exams* and *writing paper* is crucial in offering targeted support for stress relief.

Nevertheless, building an effective stressor estimation framework is non-trivial, facing two typical challenges. 1) First, users’ stressors are quite diverse, with many long-tailed stressor classes containing only a few samples per class. 2) Second, new stressors incessantly appear as time progresses (e.g., covid-19 is a stressor appearing after 2019). The model must have the ability to learn the latest stressors quickly.

As traditional supervised learning methods require a large amount of data for supervised training to recognize each stressor, they are not suitable for new stressors and long-tailed stressors with scarce data. Hereby, we cast the problem of social media-based stressor estimation within a few-shot learning setting. Few-shot learning intends to train a deep learning model to recognize new classes with only a few labeled training examples, given prior experience with very similar tasks for which we have large training examples available [43, 45]. Previous work has suggested an effective way to acquire knowledge from a few examples via meta-learning [18, 36, 51, 58]. Meta-learning (often described as “learning to learn”) advocates to learn at two levels, each associated with different time scales. The first is to quickly acquire task-specific knowledge within each separate task, and the second is to slowly summarize different task-specific knowledge to get the generic knowledge across the tasks. With the internal representation that is broadly suitable for many tasks, the obtained meta-model can thus quickly adapt to new environments through fine-tuning with a few labeled samples when facing an unseen task.

This motivates us to build a meta-learning based stressor estimation framework that is trained with past data, and can effectively be adopted to estimate new stressors with only a small number of training examples. Inspired by human’s fast learning ability

of entering the task with a large amount of prior knowledge encoded in the brains and DNA [35], we enhance meta-learning with a **meta-knowledge inheritance mechanism**, which defines a novel meta-knowledge inheritance loss and revised overall training objective to inherit knowledge from the prior meta-model without catastrophic forgetting for better model adaption.

In summary, the paper makes the following three contributions.

- From the **task** perspective, we propose a new task aimed at enhancing human well-being: social media-based *practical scenario few-shot stressor estimation* task. Unlike previous stress classification tasks, this task focuses on estimating specific causes of stress, enabling us to provide more targeted support for stress relief. We define this task with challenges that are practical.
- From the **method** perspective, we introduce a novel meta-learning-based stressor estimation framework, which incorporates a specially designed *meta-knowledge inheritance mechanism*. Our model exhibits strong generalization ability, enabling it to estimate new stressors with a little labeled data and avoid the problem of catastrophic forgetting.
- From the **data** perspective, we create a *stressor-oriented dataset* that contains 4,254 manually annotated posts. Our publicly available data would enable future research facilitating human well-being in different fields.

The performance study on the constructed dataset shows that the proposed stressor estimation framework can achieve over 74.2% F1-score, which significantly outperforms both traditional and few-shot sequence labeling baselines.

As stress-causing health problems have continued to increase all over the world, we hope this work could stimulate further interests in leveraging social media as data sources and artificial intelligence approaches to help address this critical issue¹.

2 Related Work

2.1 Stress Detection on Social Media

Stress State Classification These studies aim to classify whether an individual is stressed or not and the level of stress [37, 57]. Xue et al. [61] proposed a framework for chronic stress detection by aggregating individual tweet’s stress detection results. Saha and Choudhury [37] presented machine learning techniques to assess how the stress of campus population changes following an incident of gun violence. TensiStrength [46] employed a lexical approach and a set of rules to classify direct and indirect expressions of stress or relaxation. Based on this, Gopalakrishna et al. [13] introduced word sense disambiguation by word sense vectors to improve the performance of TensiStrength. Guntuku et al. [14] explored multiple domain adaptation algorithms to adapt user-level Facebook models to Twitter language. Wang et al. [56] studied personalized stress classification step by step from the generic mass level, group level, to final individual level. Turcan et al. [48] explored multi-task learning to co-train stress classification with emotion classification. Alghamdi et al. [1] instructed the large language model to classify stress state based on post content summary.

¹The dataset and code could be applied on <https://github.com/xin-wang18/MISE> (Note that applicants need to sign an agreement about ethics)

Stress Category Classification These studies aim to classify individual’s stressful states in a preset stress category, e.g., study, work, and financial state [65]. Xue et al. [60] investigated a number of features and employed Gaussian Process to classify stress in different categories. Based on this, Lin et al. [23] further introduced image feature and social interaction feature to improve performance. Zhao et al. [65] considered content, posting, interaction, and comment-response features to detect stress category with support vector machine. Lin et al. [24] proposed a multi-task convolutional neural networks model to classify stress category and subject. Li et al. [22] built five stress-related lexicons corresponding to the five stress categories and employed a Chinese natural language processing tool to analyze stress. Cao et al. [4] pre-trained BERT with a stress post classification task and proposed a multi-attention model to detect chronic stress in each category. Wang et al. [54] proposed to classify rarely appeared stress categories with GCN and Mixture of Experts mechanism.

Although these studies have explored states and categories of stress, they have not estimated specific causes of stress, which are critical to subsequent stress relief. Therefore, we further propose to leverage social media for stressor estimation. Moreover, prior studies are scarcely feasible in real world with incessantly emerging new stressors. They rely on a predefined set of categories, being unable to adapt to latest stressors. In this paper, we propose a meta-learning stressor estimation framework. Specifically, The meta-learning process allows the framework to be trained on the past time periods and quickly applied to the latest time period with a few labeled samples.

2.2 Sequence Labeling

Sequence labeling (SL) [39] is the process of efficiently assigning labels to each individual words within a given sequence. A common labeling formats used in this context is BIOES [62], where ‘B’ signifies ‘Begin’, ‘I’ denotes ‘Intermediate’, ‘O’ represents ‘Other’, ‘E’ corresponds to ‘End’, and ‘S’ indicates ‘Single’. In the early stages of research, this task heavily relied on manually engineered features and conditional random fields [20]. However, recent advancements in deep learning have ushered in a revolution in sequence labeling. Modern models like Bidirectional LSTMs [19], Transformers [50], and BERT-based architectures [49] have consistently achieved state-of-the-art performance, transforming the landscape of sequence labeling and extending its applicability to various natural language processing tasks.

Named Entity Recognition (NER) [32] is a common sequence labeling task, aimed at identifying and categorizing entities, such as names of persons, organizations, and locations, within text [31, 33, 40, 41]. However, NER tasks typically involve the identification and further classification of named entities, which differs from the objectives of our application-oriented stressor estimation. Our focus lies in pinpointing the specific causes of stress to enable targeted stress relief, obviating the need for further classification.

Our work focuses on the real-world stressor estimation task, which presents unique challenges beyond traditional SL or NER tasks. 1) In the context of social media, stressors are more diverse, casual, and personalized, encompassing not only named entities but also gerunds (e.g., partying) and phrases (e.g., working overtime).

The lack of strict structure and diverse linguistic expressions in social media posts necessitates the exploration of robust and adaptable natural language processing techniques. 2) Furthermore, our primary concern is to address the issue of catastrophic forgetting that arises when new stressors are learned over time. Handling this continual learning scenario is non-trivial and requires careful consideration in designing our method. This stands in contrast to most few-shot SL or NER tasks, which typically concentrate solely on performance improvements on new classes.

2.3 Meta-learning

In the literature, meta-learning has demonstrated good strength in dealing with few-shot learning problems. It learns the common parts of different tasks, and then adapts the obtained meta-model to new tasks rapidly with a few training examples [58]. Meta-learning approaches generally fall into three categories: model-based, metric-based, and optimization-based.

(1) Model-based meta-learning intends to design a meta-learner model that can update the parameters rapidly with a few training steps [34, 38]. For instance, Santoro et al. [38] trained a Memory-Augmented Neural Network as a meta-learner, which accumulated knowledge about previous tasks in an external memory that can fastly adapt to new tasks. (2) Metric-based meta-learning learn a generalized metric to measure the distance between samples in a feature embedding space [43, 45, 52]. For instance, Prototypical Network [43] averaged the embeddings of the labeled examples as the prototype for every class, and then measured the distance between each test instance and each prototype. (3) Optimization-based meta-learning strives to learn well-initialized model parameters that can generalize better to similar tasks [8, 9, 35, 36]. For instance, MAML [8] trained model’s parameters through double gradient updates, so that several gradient steps with a few labeled samples can achieve good results on new tasks. Among these three approaches, the optimization-based approach is more suitable for our problem. This is because the optimization-based approach is model structure-independent. Unlike the model-based and distance-based approaches, it does not require an additional memory module to store historical knowledge or a specialized relational module to compare the query set with the support set.

In this study, we build a meta-learning-based stressor estimation framework. To avoid catastrophic forgetting [63] when learning new stressors, we expand upon the optimization-based meta-learning by introducing a meta-knowledge inheritance mechanism. Our approach exhibits significant differences from prior meta-learning studies, as we define a novel meta-knowledge inheritance loss and a revised overall training objective for better meta-model adaption.

3 Problem Definition

Given a post $X = \{x_1, x_2, \dots, x_n\}$, where x_i denotes the i -th word in the post. Targeting to identify specific stressors that cause user stress, stressor estimation can be generalized as a sequence labeling problem. To achieve this, our framework needs to learn a mapping function $F(\phi) : X \rightarrow Y$ that takes $X = \{x_1, x_2, \dots, x_n\}$ as the input and output the corresponding label sequence $Y = \{y_1, y_2, \dots, y_n\}$




Figure 2: Schematic diagram of our practical scenario few-shot stressor estimation task, which focuses on the practical scenario that new stressors continue to emerge over time and some long-tailed stressor has scarce labeled data. Specifically, this scenario needs to train models with past data and effectively estimate latest stressor with a few adaption data.

encoded with BIOES (B-Begin, I-Intermediate, O-Other, E-End, S-Single). For instance, given the input sequence “The dual stress of losing job and buying a house make me almost vomit blood!”, we can obtain stressors “losing job” and “buying a house” from the corresponding label sequence “O,O,O,O,B,E,O,B,I,E,O,O,O,O,O,O”.

In this paper, we focus on the practical scenario that new stressors continue to emerge over time and some long-tailed stressor has scarce labeled data. This requires the capability of learning to identify such stressors with just a few labeled samples. We represent the posts from past time periods with $D_p = \{d_1, d_2, \dots, d_m\}$ and the posts from the latest time period with $D_l = \{d_{m+1}\}$. Here, $d_i = \{X_1, X_2, \dots, X_a\}$ represents the post set within the i -th time period, and we define a half-year as a time period².

The aim of our work is to build a stressor estimation framework, which is trained on past time periods D_p and able to quickly adapt to estimate new stressors on latest time period D_l . However, due to the limited labeled support set in D_l , optimizing a satisfactory supervised learning framework is challenging. To address this, we perform meta-learning on D_p to estimate generic knowledge that can be applied to perform better few-shot learning. To facilitate this process, we design a meta-knowledge Inheritance mechanism that inherit prior knowledge to prevent catastrophic forgetting.

4 Methodology

Figure 3 illustrates the overall framework of our meta-learning based stressor estimation enhanced with meta-knowledge inheritance mechanism. Our framework contains three parts: basic estimator, meta-training process, and meta-testing process. The basic estimator is employed to estimate stressors on social media. The meta-training process is designed to get a meta-model which learns the generic stressor contexts. The meta-testing process is designed to test our model for identifying new and unseen stressors with a few labeled samples. In this process, we propose a meta-knowledge inheritance mechanism to avoid catastrophic forgetting through inheriting knowledge from the prior meta-model.

4.1 Basic Estimator

The basic estimator (Figure 3.(a)) is responsible for post encoding and stressor estimation. Since pre-trained language models

²Based on our experimental data, the selection of a half-year time period strikes a balance between capturing historical and new stressors. This interval ensures a substantial amount of historical stressor data, enabling analysis of the system’s ability to learn common stressors. At the same time, it includes recent stressors, which are crucial for evaluating the system’s adaptability to new stressors.

(PLMs) have strong semantic and context learning capabilities, we employ an effective pre-trained language model RoBERTa [26] to encode our post. Specifically, we input post $X = \{x_1, x_2, \dots, x_n\}$ into RoBERTa to obtain the text representation $H = \{h_1, h_2, \dots, h_n\}$.

Constraints are needed to ensure that the predicted label sequences are reasonable, such as ensuring that there will be no prediction of “...B,B...” (i.e., two begin labels are connected in the prediction label sequence). Since Conditional Random Field (CRF) [20] is effective to learn these constraints of the label sequences, we employ CRF to decode:

$$p(Y | H) = \frac{\exp\left(\sum_{i=0}^{\alpha} U(y_i, h_i) + \sum_{i=0}^{\alpha-1} T(y_i, y_{i+1})\right)}{\sum_{y' \in \mathcal{Y}(H)} \exp\left(\sum_{i=0}^{\alpha} U(y'_i, h_i) + \sum_{i=0}^{\alpha-1} T(y'_i, y'_{i+1})\right)}, \quad (1)$$

where U denotes the emission function that represents the probability of predicting y_i . T is the transition matrix which represents the probability that a transition from y_i to y_{i+1} occurs. $\mathcal{Y}(H)$ denotes the set of all possible label sequences for H . Y is the corresponding label sequence to post X . The loss function is defined as:

$$\mathcal{L} = -\log p(Y | H), \quad (2)$$

4.2 Meta-learning Enhanced with Meta-knowledge Inheritance Mechanism

Since new stressors emerge as time progress and their labeled data are always limited, our model is expected to learn to estimate new stressors with a few labeled samples. Therefore, a meta-learning process (Figure 3.(b) and (c)) with our specially designed meta-knowledge inheritance mechanism is deployed.

4.2.1 Meta-Training. The meta-training is deployed on past time periods D_p with relatively large number of labeled data. For each training epoch, a meta-task τ_p is constructed by randomly sampling a time period d_i from the training set D_p , and the K labeled samples in d_i are chosen to serve as the support set S , as well as a part of remainder of d_i are chosen to act as the validation set V .

Inspired by MAML [8], the two-level learning are applied to train a meta-model that can be quickly adapted to unseen new tasks. Let θ denote the model parameter. For each meta task τ_p , we first feed support set S to the model and calculate the loss $\mathcal{L}_\tau(\theta)$ on S to update θ through gradient descent:

$$\theta'_\tau = \theta - \alpha \nabla_\theta \mathcal{L}_\tau(\theta), \quad (3)$$

where α is the learning rate of first level meta-learning. θ'_τ denotes the updated model parameter. This first process is to update the task-specific model, which quickly acquires task-specific knowledge (such as remembering the specific stressors and their context) through minimizing the loss on S . Secondly, we update the meta-model through the loss $\mathcal{L}_\tau(\theta'_\tau)$ on validation set V ,

$$\theta \leftarrow \theta - \beta \nabla_\theta \sum_{\tau \in \mathcal{T}} \mathcal{L}_\tau(\theta'_\tau), \quad (4)$$

where β is the learning rate of the second level meta-learning. This second process is to get the meta-model, which slowly summarizes task-specific knowledge as generic knowledge (such as summarizing generic stressor context) through minimizing the loss on V . After sufficient training over meta-training tasks, the meta-model

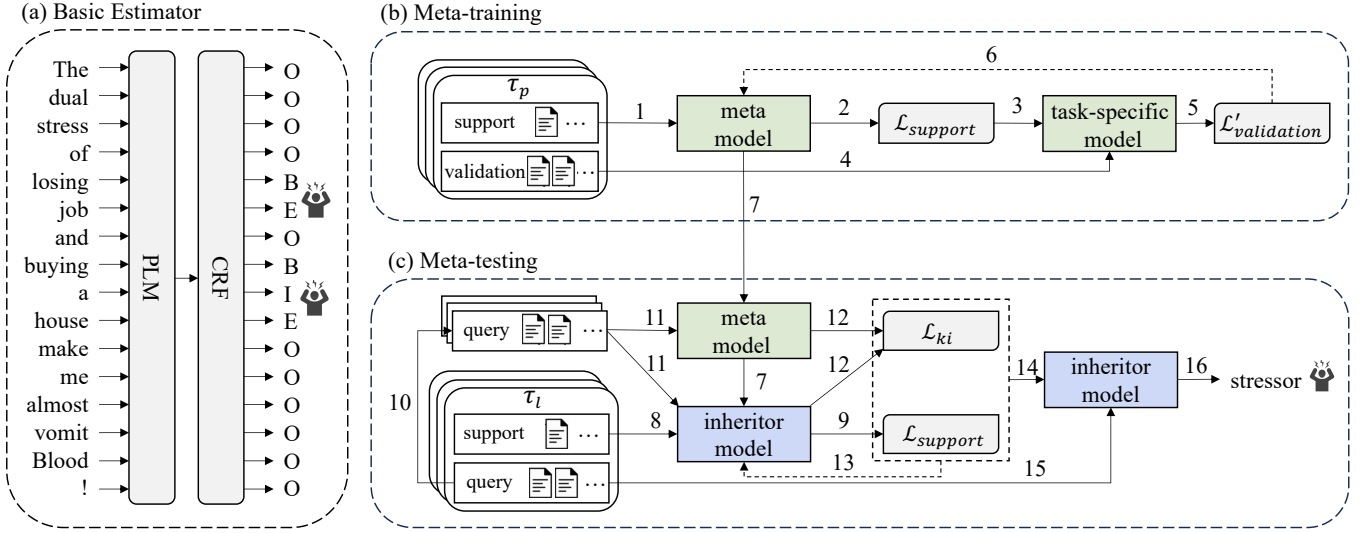


Figure 3: Overview of our proposed framework. (a) The basic estimator. (b) The meta-training process. (c) The meta-testing process with meta-knowledge inheritance mechanism.

can be quickly adapted for stressor estimation over meta-testing tasks.

4.2.2 Meta-Testing. The meta-testing is performed on the latest time period D_l with a few labeled samples. We apply the same epoch-constructed mechanism to test whether our model can indeed adapt quickly to estimate the latest stressor. For each testing epoch, we construct a meta-task τ_l by randomly sampling the support set S and query set Q from D_l . The support set is employed to optimize the meta-model for adapting the new task. The query set is to test the optimized meta-model. The result is defined as the average performance across all testing epochs.

Meta-knowledge Inheritance Mechanism. When adapting the meta-model to a new testing task with a few labeled samples, it will suffer from the problem of **catastrophic forgetting** [11], i.e., the learning of a new task may cause the model to forget the knowledge learned from previous tasks. We argue that prior knowledge is important for estimating stressors in the latest time periods. Because in addition to new stressors and rare stressors, there are also some common stressors (e.g., exam) in the latest time period. Knowledge from previous tasks can help our model effectively estimate these common stressors, so it is essential to avoid catastrophic forgetting.

We propose a meta-knowledge inheritance mechanism to generate an inheritor-model, which consolidates the knowledge from the meta-model when adapting to new testing tasks. Specifically, we define a knowledge inheritance loss for meta-knowledge consolidating. The core idea of the knowledge inheritance loss is to make the inheritor-model imitate the prediction results of the meta-model on the query set Q . These prediction results are represented by soft labels [16], which are formulated as the predicted stressor label probability:

$$P_b^M(e_i, t) = \frac{\exp[f_b(e_i/t)]}{\sum_{c=0}^4 \exp[f_c(e_i/t)]}, \quad (5)$$

where e_i denotes the representation of word x_i in post X_j from query set Q for meta-testing task τ_l . t is the temperature parameter to soften the peaky softmax distribution. $f_c(e_i)$ represents the logit score that e_i achieves on label class c . The knowledge inheritance loss is defined as follows:

$$\mathcal{L}_{ki}(\theta) = t^2 \sum_{X_j} \sum_{e_i} \sum_c P_c^M(e_i, t) \log \left(\frac{P_c^M(e_i, t)}{P_c^I(e_i, t)} \right), \quad (6)$$

where P_c^M and P_c^I denote the predicted distributions of meta-model and inheritor-model, respectively. At the same time, we employ the CRF loss on support set S to make inheritor-model learn task-specific knowledge. Then, we provide the overall learning objective $\mathcal{L}_{total}(\theta)$, which is a summation of knowledge inheritance loss $\mathcal{L}_{ki}(\theta)$ on query set and the CRF loss on support set as follows:

$$\mathcal{L}_{total}(\theta) = (1 - \lambda)\mathcal{L}(\theta) + \lambda\mathcal{L}_{ki}(\theta), \quad (7)$$

where λ denotes a trade-off parameter for the two losses. The inheritor-model for meta-testing tasks is initialized with the meta-model obtained by meta-training, and further updated as follows:

$$\theta^* = \theta - \alpha \nabla_{\theta} \mathcal{L}_{total}(\theta), \quad (8)$$

where α denotes the learning rate. Finally, the inheritor-model is deployed to estimate stressors in the query set of meta-testing tasks. We summarize the meta-training and meta-testing process as shown in Algorithm 1.

Summary. Our approach exhibits significant differences from existing meta-learning studies, as we define a novel meta-knowledge inherit loss and a revised overall training objective for better meta-model adaptation. In other words, equations (5)(6)(7)(8) in our approach are core innovations and differences compared with existing meta-learning approaches.

Algorithm 1 Meta-learning Enhanced with meta-knowledge Inheritance Mechanism**Require:** α, β : the learning rates.

```

while not done the meta-training do
  Sample batch of meta-task  $\tau_p$  from  $D_p$ 
  for Each  $\tau_p$  do
    Sample a support set  $S$  and a validation set  $V$ 
    Evaluate  $\mathcal{L}_\tau(\theta)$  with data  $S$ 
    Compute task-specific parameters with gradient descent:
     $\theta'_\tau = \theta - \alpha \nabla_\theta \mathcal{L}_\tau(\theta)$ 
  end for
  Updating meta-model's parameters with data  $V$  from each
  meta-task:  $\theta \leftarrow \theta - \beta \nabla_\theta \sum_{\tau \in \mathcal{T}} \mathcal{L}_\tau(\theta'_\tau)$ 
end while
while not done the meta-testing do
  Sample batch of meta-task  $\tau_l$  from  $D_l$ 
  for Each  $\tau_l$  do
    Sample the support set  $S$  and the query set  $Q$ 
    Evaluate  $\mathcal{L}(\theta)$  with data  $S$ 
    Inherit knowledge from the meta-model with meta-
    knowledge Inheritance loss:
    
$$P_b^M(e_i, t) = \frac{\exp[f_b(e_i/t)]}{\sum_{c=0}^4 \exp[f_c(e_i/t)]}$$

    
$$\mathcal{L}_{ki}(\theta) = t^2 \sum_{X_j} \sum_{e_i} \sum_c P_c^M(e_i, t) \log\left(\frac{P_c^M(e_i, t)}{P_c^I(e_i, t)}\right)$$

    Calculate the overall loss:  $\mathcal{L}_{\text{total}}(\theta) = (1-\lambda)\mathcal{L}(\theta) + \lambda\mathcal{L}_{ki}(\theta)$ 

    Update the inheritor-model's parameters via optimizing
    Equation:  $\theta^* = \theta - \alpha \nabla_\theta \mathcal{L}_{\text{total}}(\theta)$ 
  end for
end while

```

5 Experiment

5.1 Dataset Construction and Statistics

One significant challenge for this new task is the lack of available dataset. The most relevant datasets we discovered, such as Dreddit [47] for stress state detection and Plscd [54] for stress category classification, only provide labels indicating whether a post expresses stress and its broad category (e.g., work, study, financial state). Unfortunately, these datasets do not include the specific stressor labels that are necessary for our stressor estimation task. To overcome this challenge, we build a social media-based stressor estimation dataset.

The dataset is crawled from Weibo. Since Weibo is one of the most popular social media platforms worldwide, with 249 million daily active users [42]. Specifically, the data labeling process is carried out as follows:

Firstly, we employ the retrieval function to get the stressful posts. The retrieval term is user's self-report, such as "I, stressed". This self-report ensures that the posts contain stress expression [56].

Then, we hire six master students who major in psychology and are active on social media to annotate the stressor. Annotators are asked to read posts carefully and understand the semantics of posts to label stressors. We pay annotators \$0.15 per post. There are situations in which the user only expresses that he/she is stressed

without mentioning the source of the stress. In this case, no annotation is made. Note that the annotators are asked to further remove the following types of noise data:

- Others feel stressed, e.g., "Is it really stressful to chat with me? Why are they so easily nervous?"
- Advertise, e.g., "#Loan# #Internet Loan# #Universal Gold# Can't pay your monthly bills? Repayment stress? Poke me to get the universal gold"
- Stress of the past, e.g., "Every time I pass by here, I think of the time when I was preparing for the postgraduate entrance examination. At that time, I was so stressed."
- Future plan, e.g., "I hope that in the future I will never get married because of age and stress!"
- Lyrics or verses, e.g., "Some days, I just wanna leave the negativity in my head. I just want relief from my stress (Song Lyrics)"

Each post is analyzed by three different annotators. Only if their annotated stressors are consistent, the stressor is used as the annotation result. If their annotated stressors are inconsistent, all different stressors are kept pending further annotation.

Finally, the annotation progress is completed by a psychology expert from the Department of Psychology. The main work of the expert can be summarized in two parts: 1) checking the posts with consistent annotation results and 2) giving an expert opinion as the final annotation for those posts with inconsistent annotation.

In total, we obtain 4,254 labeled posts from June 2018 to June 2022. The average inner-agreement (i.e., Cohen's Kappa) between annotators is 0.71, indicating good agreement. The dataset is divided into different time periods based on half-year intervals. We have seven past time periods, labeled as D_p , spanning from June 2018 to December 2021. Additionally, we have one latest time period, denoted as D_l , covering January 2022 to June 2022.

5.2 Baselines

We compare our method with the following traditional sequence labeling baselines: LSTM-CRF [19], CNN-CRF [30], SED [2], DetIE [49], and CRUP [27]. This paper considers a practical scenario that new stressors continuously emerge over time and some long-tailed stressors have limited labeled data. Thus we employ the following few-shot sequence labeling baselines to further verify the effectiveness: SimBERT [7], Ma21 [29], ConVEx [15], ESD [53], and BDCP [59]. More information about these baselines can be found in Appendix D.

5.3 Main Results

Table 1 reports the performance of all methods. The first part shows the performance of traditional sequence labeling methods. The average performance of these methods is lower than 63%, which indicates that these methods are not good at estimating new stressor with a few labeled samples.

The second part lists the performance of few-shot baselines. We can see that they outperform traditional methods with 5% average improvement, which suggests that the meta-learning or few-shot learning process can significantly enhance the estimation performance in the latest time period.

Table 1: Main Results. The first part lists the performance of traditional sequence labeling baselines. The second part shows the performance of few-shot baselines. MISE denotes our method.

Method	3-shot			5-shot			10-shot		
	Precision	Recall	F1-score	Precision	Recall	F1-score	Precision	Recall	F1-score
LSTM-CRF [19]	0.5684	0.5870	0.5775	0.5674	0.6081	0.5829	0.5778	0.6298	0.5986
CNN-CRF [30]	0.5643	0.6126	0.5862	0.5752	0.6134	0.5958	0.5895	0.6330	0.6014
SED [2]	0.5761	0.6310	0.6023	0.5891	0.6385	0.6056	0.5982	0.6507	0.6311
DEIE [49]	0.5996	0.6333	0.6164	0.6055	0.6503	0.6238	0.6126	0.6854	0.6538
CRUP [27]	0.6007	0.6364	0.6196	0.6120	0.6528	0.6263	0.6230	0.6925	0.6563
SimBERT [7]	0.6086	0.6417	0.6248	0.6359	0.6791	0.6543	0.6538	0.6910	0.6766
Ma21 [29]	0.6274	0.6204	0.6239	0.6463	0.6310	0.6384	0.6742	0.6601	0.6671
ConVEx [15]	0.6216	0.6422	0.6317	0.6457	0.6673	0.6562	0.6731	0.6956	0.6881
ESD [53]	0.6413	0.6612	0.6509	0.6582	0.6897	0.6745	0.6814	0.7020	0.6902
BDCP [59]	0.6360	0.6573	0.6459	0.6498	0.6759	0.6627	0.6734	0.6948	0.6855
MISE (Ours)	0.6825	0.7029	0.6914	0.7053	0.7366	0.7212	0.7332	0.7541	0.7420

Table 2: Ablation Study. M denotes the entire meta-learning process. I denotes the meta-knowledge inheritance mechanism.


Method	3-shot			5-shot			10-shot		
	Precision	Recall	F1-score	Precision	Recall	F1-score	Precision	Recall	F1-score
MISE	0.6825	0.7029	0.6914	0.7053	0.7366	0.7212	0.7332	0.7541	0.7420
w/o M	0.6178	0.6403	0.6289	0.6419	0.6743	0.6570	0.6689	0.6843	0.6785
w/o I	0.6392	0.6713	0.6596	0.6628	0.6954	0.6807	0.6975	0.7182	0.7068

Our method MISE achieves over 74.2% in F1-score, with over 4.0-5.1% improvement compared with the best baseline. It illustrates that our meta-learning method is effective in estimating new stressor with a few labeled samples. We attribute the improvements to the fact that our method is well-designed with the meta-learning process and meta-knowledge inheritance mechanism. Furthermore, our model’s standard deviation is 0.012, 0.013, and 0.019 for 3-shot, 5-shot, and 10-shot settings, indicating good robustness.

There is a general trend of improvement in performance as we move from traditional baselines to few-shot baselines, and then to our MISE method. The improvement in F1-scores from 3-shot to 10-shot learning scenarios suggests that having more support set examples improves the model’s performance, which is a common trend in few-shot problem.

5.4 Ablation Study

To analyze the effect of different components in MISE, we construct ablation experiments as shown in Table 2. The performance of MISE will drop over 6.2% without the entire meta-learning process, which verifies that meta-learning can effectively make our model learn to identify new stressors with a few labeled samples. The average decline in performance will be 3.5% without the meta-knowledge inheritance mechanism. This demonstrates that meta-knowledge inheritance enables our model to effectively avoid catastrophic forgetting and enhance model performance.

**Figure 4: Parameter Study on λ and t .**

5.5 Catastrophic Forgetting Study

To evaluate the effectiveness of MISE in addressing the catastrophic forgetting problem, we conduct an experiment that tests model’s performance on previous tasks after it learns a new task. Firstly, we randomly select 20% of the data from past time periods D_p and utilize 80% of the remaining data to train the model. Next, we adapt the model using a few labeled samples from latest time period D_l . Finally, we test the adapted model using the 20% isolated data. This process is repeated five times to calculate the average performance. As shown in Table 3, MISE outperforms the best baseline over 8.4% in F1-score. It verifies that our proposed meta-knowledge inheritance mechanism can effectively resolve the catastrophic forgetting problem when learning a new task.

Table 3: Catastrophic Forgetting Study. All methods are tested on the past time periods after adaptation to latest time period.

Method	Precision	Recall	F1-score
LSTM-CRF [19]	0.6674	0.6643	0.6665
CNN-CRF [30]	0.6740	0.6628	0.6687
SED [2]	0.6921	0.6747	0.6854
DeIE [49]	0.7137	0.7141	0.7139
CRUP [27]	0.7110	0.7163	0.7135
SimBERT [7]	0.6942	0.6748	0.6836
Ma21 [29]	0.7039	0.6902	0.6965
ConVEx [15]	0.7217	0.7089	0.7141
ESD [53]	0.7175	0.7055	0.7106
BDCP [59]	0.7146	0.7074	0.7113
MISE (ours)	0.7952	0.8015	0.7983

5.6 Parameter Study

To analyze the sensitivity of MISE, we report the performance with a variety of hyperparameters. Specifically, we analyze parameters λ and t , corresponding to the trade-off for the overall loss and the re-scaling temperature for the meta-knowledge inheritance mechanism, respectively. As shown in Figure 4, $\lambda = 0.2$ and $t = 5$ are the best parameter settings. In addition, the performance of MISE keeps good with different λ and t , which illustrates that MISE is relatively insensitive and robust to the changes of the settings.

5.7 Case Study

We provide insights based on real-world cases, as demonstrated in Table 4. In the first example, our framework correctly identified the user’s stressor as “*monkeypox*”. It is worth noting that “*monkeypox*” is a new stressor for the framework. This demonstrates that our meta-learning framework is capable of estimating the latest stressors with few samples. In the second example, our framework correctly identified the stressor as “*decoration*”, which is an infrequently appeared stressor in the training data. In contrast, few-shot baselines fail to estimate the stressor “*decoration*”, because these models forget the stressor “*decoration*” when adapting to the latest time period with few samples. This highlights the importance of our proposed meta-knowledge inheritance mechanism.

6 Discussion

6.1 Potential Applications

Our work has the potential to make a significant impact on society by improving the well-being of individuals.

1. Potential applications in the field of psychological diagnosis. For instance, if a teenager experiences prolonged stress, they may be at an increased risk of self-harm or suicide. In such cases, our work can identify their stressor and alert their parents, enabling them to provide target care and support to prevent potential harm.
2. New ideas and support for psychological research. Traditional psychology experiments and conclusions are often based on tens of hired subjects [21], with very low population sampling rates. Our social media-based approach can help them study psychology on a

Table 4: Case Study. There are two real-world posts. MISE successfully estimates new stressor and forgettable stressor.

<p>Post: I feel stressed most of the time. <u>Monkeypox</u> has made life not go well for anyone lately.</p> <p>Ground Truth: Monkeypox MISE: Monkeypox ✓</p>
<p>Post: I want to buy an iPad Pro... but I have to save money for the <u>decoration</u> of my new home. I am very stressed.</p> <p>Ground Truth: decoration MISE: decoration ✓</p>

larger population scale. As a foundation, it can advance the field of psychology.

6.2 Limitation

While much of the existing social media-related research assumes that users’ expressions on social media are proactive and authentic [64], it is important to acknowledge that there are situations where users’ posts may not fully reflect their innermost thoughts. For instance, a user’s real stressor is “interview”, but he/she might post: “This bad weather makes me stressed.” and not mention about “interview”. In such cases, the effectiveness may be compromised. Due to dataset limitations, we had to construct and experiment on a single social media platform, which may introduce platform bias.

7 Conclusion

In this paper, we discover and propose a new social media-based stress-related task, i.e., stressor estimation. To address the practical problem of estimating the latest stressor with limited data, we propose a meta-learning framework that is trained on past data and can be adapted to the latest stressor estimation. We expand upon optimization-based meta-learning by introducing a meta-knowledge inheritance mechanism to avoid catastrophic forgetting when learning new stressors. Our approach defines a novel meta-knowledge inheritance loss and a revised training objective, which distinguishes it from prior meta-learning studies. Experimental results demonstrate that our framework is significantly effective when compared to all state-of-the-art baselines. In addition, a well-labeled social media stressor estimation dataset is proposed.

Acknowledgement

This work was supported by the Pandemic Sciences Institute at the University of Oxford; the National Institute for Health Research (NIHR) Oxford Biomedical Research Centre (BRC); an NIHR Research Professorship; a Royal Academy of Engineering Research Chair; the Zhipu AI; the Wellcome Trust funded VITAL project (grant 204904/Z/16/Z); the EPSRC (grant EP/W031744/1); the NSFC (62306039); the Fundamental Research Funds for the Central Universities (310422127); and the InnoHK Hong Kong Centre for Cerebrocardiovascular Engineering (COCHE).

References

- [1] Zeyad Alghamdi, Tharindu Kumarage, Garima Agrawal, Huan Liu, and H Russell Bernard. 2024. Less is More: Stress Detection through Condensed Social Media Contents. In *European Conference on Social Media*, Vol. 11, 13–22.
- [2] Giannis Bekoulis, Johannes Deleu, Thomas Demeester, and Chris Develder. 2019. Sub-event detection from twitter streams as a sequence labeling problem. In *NAACL Association for Computational Linguistics*, Minneapolis, Minnesota, 745–750. doi:10.18653/v1/N19-1081
- [3] Edward S Bordin. 1955. Psychological counseling. (1955).
- [4] Lei Cao, Huijun Zhang, Ningyun Li, Xin Wang, Wisong Ri, and Ling Feng. 2021. Category-Aware Chronic Stress Detection on Microblogs. *IEEE Journal of Biomedical and Health Informatics* (2021).
- [5] Stevie Chancellor, Michael L Birnbaum, Eric D Caine, Vincent MB Silenzio, and Munmun De Choudhury. 2019. A taxonomy of ethical tensions in inferring mental health states from social media. In *Proceedings of the conference on fairness, accountability, and transparency*, 79–88.
- [6] Sheldon Cohen, Tom Kamarck, and Robin Mermelstein. 1983. A global measure of perceived stress. *Journal of health and social behavior* (1983), 385–396.
- [7] 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)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. doi:10.18653/v1/N19-1423
- [8] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*. PMLR, 1126–1135.
- [9] Chelsea Finn, Kelvin Xu, and Sergey Levine. 2018. Probabilistic model-agnostic meta-learning. *Advances in neural information processing systems* 31 (2018).
- [10] Mental Health Foundation. 2022. Stress: are we coping? <https://www.mentalhealth.org.uk/publications/stress-are-we-coping>.
- [11] Robert French. 1993. Catastrophic interference in connectionist networks: Can It Be predicted, can It be prevented?. In *Advances in Neural Information Processing Systems*, J. Cowan, G. Tesauro, and J. Alspector (Eds.), Vol. 6. Morgan-Kaufmann. <https://proceedings.neurips.cc/paper/1993/file/28267ab848bcf807b2ed53c3a8f8c8a-Paper.pdf>
- [12] Sherryl H Goodman, Daniel R Sewell, and Ruth C Jampol. 1984. On going to the counselor: Contributions of life stress and social supports to the decision to seek psychological counseling. *Journal of Counseling Psychology* 31, 3 (1984), 306.
- [13] Reshmi Gopalakrishna Pillai, Mike Thelwall, and Constantin Orasan. 2018. Detection of stress and relaxation magnitudes for tweets. In *WWW*, 1677–1684.
- [14] Sharath Chandra Guntuku, Anneke Buffone, and Kokil Jaidka. 2019. Understanding and measuring psychological stress using social media. In *AAAI*, 214–225.
- [15] Matthew Henderson, Ivan Vulic, and et al. 2021. ConVEx: Data-Efficient and Few-Shot Slot Labeling. In *NAACL Association for Computational Linguistics*, Online, 3375–3389. doi:10.18653/v1/2021.naacl-main.264
- [16] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. 2014. Distilling the knowledge in a neural network. In *Deep Learning and Representation Learning Workshop*. NIPS.
- [17] Thomas H Holmes and Richard H Rahe. 1967. The social readjustment rating scale. *Journal of psychosomatic research* (1967).
- [18] Timothy Hospedales, Andreas Antoniou, Paul Micaelli, and Amos Storkey. 2021. Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence* 44, 9 (2021), 5149–5169.
- [19] Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991* (2015).
- [20] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001)*, Williams College, Williamstown, MA, USA, June 28 - July 1, 2001, Carla E. Brodley and Andrea Pohorecký Danyluk (Eds.). Morgan Kaufmann, 282–289.
- [21] Richard S Lazarus. 1966. Psychological stress and the coping process. (1966).
- [22] Qi Li, Yuanyuan Xue, Liang Zhao, Jia Jia, et al. 2016. Analyzing and identifying teens' stressful periods and stressor events from a microblog. *JBHI* 21, 5 (2016), 1434–1448.
- [23] Huijie Lin, Jia Jia, Quan Guo, et al. 2014. Psychological stress detection from cross-media microblog data using deep sparse neural network. In *ICME*. IEEE, 1–6.
- [24] Huijie Lin, Jia Jia, Liqiang Nie, and Guangyao Shen. 2016. What Does Social Media Say about Your Stress?. In *IJCAI*, 3775–3781.
- [25] Nan Lin and Walter M Ensel. 1989. Life stress and health: Stressors and resources. *American sociological review* (1989), 382–399.
- [26] 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. *arXiv preprint arXiv:1907.11692* (2019).
- [27] Yanhe Liu, Peng Wang, Wenjun Ke, Guozheng Li, Xiye Chen, Jiteng Zhao, and Ziyu Shang. 2024. Unify Named Entity Recognition Scenarios via Contrastive Real-Time Updating Prototype. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38, 14035–14043.
- [28] Ilya Loshchilov and Frank Hutter. 2018. Fixing Weight Decay Regularization in Adam. <https://openreview.net/forum?id=rk6qdGgCZ>
- [29] Jianqiang Ma, Zeyu Yan, Chang Li, and Yang Zhang. 2021. Frustratingly simple few-shot slot tagging. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 1028–1033.
- [30] Xuezhe Ma and Eduard Hovy. 2016. End-to-end Sequence Labeling via Bidirectional LSTM-CNNs-CRF. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, 1064–1074. doi:10.18653/v1/P16-1101
- [31] Alireza Mansouri, Lilly Suriani Affendey, and Ali Mamat. 2008. Named entity recognition approaches. *International Journal of Computer Science and Network Security* 8, 2 (2008), 339–344.
- [32] Mónica Marrero, Julián Urbano, Sonia Sánchez-Cuadrado, Jorge Morato, and Juan Miguel Gómez-Berbis. 2013. Named entity recognition: fallacies, challenges and opportunities. *Computer Standards & Interfaces* 35, 5 (2013), 482–489.
- [33] Sudha Morwal, Nusrat Jahan, and Deepti Chopra. 2012. Named entity recognition using hidden Markov model (HMM). *International Journal on Natural Language Computing (IJNLC) Vol 1* (2012).
- [34] Tsensuren Munkhdalai and Hong Yu. 2017. Meta networks. In *International Conference on Machine Learning*. PMLR, 2554–2563.
- [35] Alex Nichol, Joshua Achiam, and John Schulman. 2018. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999* (2018).
- [36] Sachin Ravi and Hugo Larochelle. 2016. Optimization as a model for few-shot learning. (2016).
- [37] Koustuv Saha and Munmun De Choudhury. 2017. Modeling stress with social media around incidents of gun violence on college campuses. *HCI* 1, CSCW (2017), 1–27.
- [38] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. 2016. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*. PMLR, 1842–1850.
- [39] Burr Settles and Mark Craven. 2008. An analysis of active learning strategies for sequence labeling tasks. In *proceedings of the 2008 conference on empirical methods in natural language processing*, 1070–1079.
- [40] Khaled Shaalan and Hafsa Raza. 2007. Person name entity recognition for Arabic. In *Proceedings of the 2007 workshop on computational approaches to semitic languages: common issues and resources*, 17–24.
- [41] Khaled Shaalan and Hafsa Raza. 2008. Arabic named entity recognition from diverse text types. In *Advances in Natural Language Processing: 6th International Conference, GoTAL 2008 Gothenburg, Sweden, August 25-27, 2008 Proceedings*. Springer, 440–451.
- [42] Sina. 2021. Annual Report 2021. <http://ir.weibo.com/financial-information/annual-reports>.
- [43] Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems* 30 (2017).
- [44] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* 15, 1 (2014), 1929–1958.
- [45] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. 2018. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1199–1208.
- [46] Mike Thelwall. 2017. TensiStrength: Stress and relaxation magnitude detection for social media texts. *Information Processing & Management* (2017), 106–121.
- [47] Elsbeth Turcan and Kathleen McKeown. 2019. Dreddit: A Reddit Dataset for Stress Analysis in Social Media. *EMNLP-IJCNLP 2019* (2019), 97.
- [48] Elsbeth Turcan, Smaranda Muresan, et al. 2021. Emotion-Infused Models for Explainable Psychological Stress Detection. In *NAACL*, 2895–2909.
- [49] Michael Vasilkovsky, Anton Alekseev, Valentin Malykh, Ilya Shenbin, Elena Tutubalina, Dmitriy Salikhov, Mikhail Stepnov, Andrey Chertok, and Sergey Nikolenko. 2022. DetIE: Multilingual Open Information Extraction Inspired by Object Detection. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*.
- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [51] Ricardo Vilalta and Youssef Drissi. 2002. A perspective view and survey of meta-learning. *Artificial intelligence review* 18, 2 (2002), 77–95.
- [52] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. 2016. Matching networks for one shot learning. *Advances in neural information processing systems* 29 (2016).
- [53] Peiyi Wang, Runxin Xu, Tianyu Liu, Qingyu Zhou, Yunbo Cao, Baobao Chang, and Zhifang Sui. 2022. An Enhanced Span-based Decomposition Method for Few-Shot Sequence Labeling. In *NAACL Association for Computational Linguistics*, Seattle, United States, 5012–5024. doi:10.18653/v1/2022.naacl-main.369

- [54] Xin Wang, Lei Cao, Huijun Zhang, Ling Feng, Yang Ding, and Ningyun Li. 2022. A Meta-learning based Stress Category Detection Framework on Social Media. In *Proceedings of the ACM Web Conference 2022*. 2925–2935.
- [55] Xin Wang, Boyan Gao, Yi Dai, Lei Cao, Liang Zhao, Yibo Yang, and David Clifton. 2024. Cognition Chain for Explainable Psychological Stress Detection on Social Media. *arXiv preprint arXiv:2412.14009* (2024).
- [56] Xin Wang, Huijun Zhang, Lei Cao, and Ling Feng. 2020. Leverage Social Media for Personalized Stress Detection. In *ACM MM*. 2710–2718.
- [57] Xin Wang, Huijun Zhang, Lei Cao, Kaisheng Zeng, Qi Li, Ningyun Li, and Ling Feng. 2023. Contrastive Learning of Stress-specific Word Embedding for Social Media based Stress Detection. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 5137–5149.
- [58] Lilian Weng. 2018. Meta-Learning: Learning to Learn Fast. <https://lilianweng.github.io/lil-log/2018/11/30/meta-learning.html>.
- [59] Xiaojun Xue, Chunxia Zhang, Tianxiang Xu, and Zhendong Niu. 2024. Robust Few-Shot Named Entity Recognition with Boundary Discrimination and Correlation Purification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 19341–19349.
- [60] Yuanyuan Xue, Qi Li, Li Jin, Ling Feng, David A Clifton, and Gari D Clifford. 2014. Detecting adolescent psychological pressures from micro-blog. In *HIS*. Springer, 83–94.
- [61] Yuanyuan Xue, Qi Li, Liang Zhao, Jia Jia, Ling Feng, Feng Yu, and David A Clifton. 2016. Analysis of teens' chronic stress on micro-blog. In *International Conference on Web Information Systems Engineering*. Springer, 121–136.
- [62] Jie Yang, Shuailong Liang, and Yue Zhang. 2018. Design Challenges and Misconceptions in Neural Sequence Labeling. In *Proceedings of the 27th International Conference on Computational Linguistics*. 3879–3889.
- [63] Pauching Yap, Hippolyt Ritter, and David Barber. 2021. Addressing catastrophic forgetting in few-shot problems. In *International Conference on Machine Learning*. PMLR, 11909–11919.
- [64] Reza Zafarani, Mohammad Ali Abbasi, and Huan Liu. 2014. *Social media mining: an introduction*. Cambridge University Press.
- [65] Liang Zhao, Qi Li, Yuanyuan Xue, Jia Jia, and Ling Feng. 2016. A systematic exploration of the micro-blog feature space for teens stress detection. *Health information science and systems* 4, 1 (2016), 1–12.

Table 5: Past Time Periods Performance Study. All methods are tested under traditional supervised scenario.

Method	Precision	Recall	F1-score
LSTM-CRF [19]	0.7652	0.7578	0.7615
CNN-CRF [30]	0.7719	0.7616	0.7667
SED [2]	0.7835	0.7573	0.7702
DeIE [49]	0.8016	0.8034	0.8023
CRUP [27]	0.8028	0.8069	0.8054
SimBERT [7]	0.7904	0.7752	0.7781
Ma21 [29]	0.7761	0.7894	0.7827
ConVEx [15]	0.7993	0.8018	0.8006
ESD [53]	0.8025	0.8065	0.8047
BDCP [59]	0.8021	0.8053	0.8040
MISE (ours)	0.8182	0.8235	0.8206

A Experiment Details

For each meta-task of meta-learning, we sample 3, 5, or 10 posts as the support set (i.e., 3-shot, 5-shot, or 10-shot) and 15 posts as the validation or query set. The first level learning rate α is set to $2e-5$. The second level learning rate β is $5e-5$. The max training epoch is set to 5000. We report the average performance from 50 random testing epochs. AdamW [28] is adopted as our optimizer. 10% drop out [44] is deployed to avoid overfitting. The temperature parameter t is set to 5 and the trade-off parameter λ is set to 0.2. Experiments are run on a Linux server with RTX 2080 GPUs. Pytorch 1.71 is used to construct the models. We employ transformers 4.18 to load RoBERTa [26] without pooling layer as PLM. Our model has 101.77M parameters. The total computational time is 2.2 hours. As we want to estimate the specific stressor with clear boundaries, we employ token-level precision, token-level recall, and token-level F1 score to measure the performance.

B Past Time Periods Performance Study

To verify MISE’s effectiveness in estimating stressors from past time periods (supervised scenario), we design an experiment that trains and tests the models on the past time periods’ data. Specifically, we conduct five-fold cross-validation on data from past time periods D_p . The results are shown in Table 5. MISE achieves 82.0% F1-score, with over 1.5% improvement compared with the best baseline. It illustrates that our framework can not only work better on the estimating latest stressors with a few labeled data but also be effective on estimating stressors from past time periods.

C Ethical Considerations

Since this work is an intersection study of human psychology and data mining, the ethical issues must be carefully deliberated.

Privacy. All the data used in this paper is publicly available on social media. We acknowledge that the psychological experiments

may potentially impact subjects [5], but our work is just an observation on public posts and there are no subjects.

Data protection. The dataset will be anonymized before being shared. We only provide post text without any personal information. Applicants must sign an agreement before they get the dataset. This agreement guarantees: 1) they will never attempt to identify or contact any user in the dataset; 2) they will never make or use the dataset for commercial purposes; etc.

D Baselines

We compare our method with the following traditional sequence labeling baselines:

- LSTM-CRF [19]. A classical sequence labeling method with bidirectional LSTM and Conditional Random Field.
- CNN-CRF [30]. Another classical sequence labeling with bidirectional LSTM, CNN, and CRF.
- SED [2]. It employs attention mechanism to get text representation and LSTM to decode the representation.
- DeIE [49]. It freezes part of BERT and trains the rest part combined with fully connected layer.
- CRUP [27]. It enhances the representations through contrastive learning for sequence labeling.

These traditional baselines are not specifically designed for few-shot learning. For fair comparison with our meta-learning model, all these baselines also only take K samples (i.e., support set in meta-task) to learn to estimate stressors in the latest time period. Specifically, we utilize all the past time periods data D_p in the training stage. Then the trained model is fine-tuned with K samples from D_l . Lastly, we test the fine-tuned model on the latest time period D_l (except the K samples).

We consider a practical scenario that new stressors continue to emerge over time and some long-tailed stressor has scarce labeled data. Thus we employ the following few-shot sequence labeling baselines to further verify the effectiveness:

- SimBERT [7]. It employs BERT to get token representation and predict labels through finding the most similar labeled token in the support set.
- Ma21 [29]. It casts sequence labeling as machine reading comprehension problem and generates the labeling types as questions.
- ConVEx [15]. It proposes a pairwise cloze task to pre-training model with Reddit data and then fine-tune the pre-trained model with a few labeled samples.
- ESD [53]. It formulates the sequence labeling problem as classification of each span and combines the idea of Prototypical Network [43] for classification.
- BDCP [59]. It introduces entity boundary discriminative module to provide a highly distinguishing boundary representation space for labeling.