

A Self-supervised Framework for Improved Data-Driven Monitoring of Stress via Multi-modal Passive Sensing

Shayan Fazeli*, Lionel Levine*, Mehrab Beikzadeh*, Baharan Mirzasoleiman*,
Bitia Zadeh†, Tara Peris‡, Majid Sarrafzadeh*

*Computer Science Department, UCLA, Los Angeles, US

†Department of Psychology, Chapman University, Los Angeles, US

‡The Jane and Terry Semel Institute for Neuroscience and Human Behavior, UCLA, Los Angeles, US

shayan@cs.ucla.edu, lionel@cs.ucla.edu, mehrabbeikzadeh@cs.ucla.edu, baharan@cs.ucla.edu,
bzadeh@chapman.edu, tperis@mednet.ucla.edu, majid@cs.ucla.edu

Abstract—Recent advances in remote health monitoring systems have significantly benefited patients and played a crucial role in improving their quality of life. However, while physiological health-focused solutions have demonstrated increasing success and maturity, mental health-focused applications have seen comparatively limited success in spite of the fact that stress and anxiety disorders are among the most common issues people deal with in their daily lives. In the hopes of furthering progress in this domain through the development of a more robust analytic framework for the measurement of indicators of mental health, we propose a multi-modal semi-supervised framework for tracking physiological precursors of the stress response. Our methodology enables utilizing multi-modal data of differing domains and resolutions from wearable devices and leveraging them to map short-term episodes to semantically efficient embeddings for a given task. Additionally, we leverage an inter-modality contrastive objective, with the advantages of rendering our framework both modular and scalable. The focus on optimizing both local and global aspects of our embeddings via a hierarchical structure renders transferring knowledge and compatibility with other devices easier to achieve. In our pipeline, a task-specific pooling based on an attention mechanism, which estimates the contribution of each modality on an instance level, computes the final embeddings for observations. This additionally provides a thorough diagnostic insight into the data characteristics and highlights the importance of signals in the broader view of predicting episodes annotated per mental health status. We perform training experiments using a corpus of real-world data on perceived stress, and our results demonstrate the efficacy of the proposed approach in performance improvements¹.

Index Terms—machine learning, eHealth, wireless health, mental health, self-supervised learning, remote health monitoring

I. INTRODUCTION

The rising epidemic of mental health disorders, worsened by the recent COVID-19 pandemic, speaks to the growing need for effective and timely management of mental health disorders.

The pandemic led to an increase in the need for mental health services, while concurrently, given the circumstances surrounding the outbreak, limited access to traditional modalities of care. This necessitated the explosion in the usage of alternative mechanisms to deliver mental health services, mainly through remote formats [1].

For that reason, in spite of increased barriers to access, unprecedented levels of funding have gone into programs to address mental health issues among the general public. For instance, in 2020 alone, the United States government spent around 280 billion dollars on mental health services [2].

Therefore, even as the pandemic, and associated restrictions on in-person activities, have subsided, the gaps it revealed in traditional in-person-based therapeutic services persist, and demand for remote solutions remains high.

While much of the focus of remote mental health services has been around the use of video conferencing, instant messaging, and other modes of communication to facilitate interactions between therapists and patients, the use of mHealth technology and passive monitoring has the potential to be equally impactful at addressing barriers to care and gaps in monitoring. Furthermore, by leveraging personal digital devices equipped with numerous sensors that are capable of monitoring many aspects of an individual’s physiology and lifestyle (e.g., heart rate, activity level), remote health monitoring provides a novel pathway to not only monitor existing indicators of mental health, but also to improve upon our understanding mental health disorders and their impacts on one’s life.

Stress, commonly defined as “physical, mental, or emotional strain or tension,” is a widespread problem with numerous potential causes. According to the American Institute of Stress, 73% of people suffer from acute bouts of stress to a degree of magnitude that impacts their mental well-being. Incidents of Anxiety often manifest similarly to stress, however, it is

¹Codes are available at <https://github.com/shayanfazeli/tabluence>

notable that it is not always immediately tied to a specific triggering or inciting event and may take longer to resolve.

All told, both stress and anxiety problems are very common, to the extent that most adults have been affected by at least one anxiety-related disorder [3], [4]. Anxiety-related disorders can have a significant negative impact on the quality of life, leading to other mental health disorders such as depression, as well as causing physical health problems [5].

In contemplating improved means to address mental health challenges generally, and anxiety-related disorders specifically, it is notable that a critical part of modern healthcare involves accurate and efficient tracking of individuals' well-being through time. Examples include tracking athletes and their training trajectories, and patients' rehabilitation exercises [6]–[10]. Compared to physiological health, the mental health domain is less investigated in the context of remote health monitoring. This is largely due to a confluence of reasons. For one, the statistical sufficiency of observations obtained via data-driven approaches is not often intuitively clear (e.g., can one draw a conclusion regarding depression from the number of phone calls?). Another reason is that the data required for enabling the use of artificial intelligence (AI) is often not readily available or exclusive due to privacy and regulatory concerns.

The works in this domain, therefore, have mostly focused on longer-term patient phenotyping (e.g., classifying patients into bipolar disorder vs healthy) [11].

While these high-level labels are useful, they can be limited in their utility, as stress often manifests as an emotional and physiological response of an individual to a triggering event, and can occur to anyone regardless of a formal diagnosis. For example, arguing with someone and being anxious about a deadline are instances of interpersonal and work-related stress that are liable to occur to anyone regardless of the existence of a pre-existing mental health disorder. Furthermore, the highly localized and temporary nature of these shorter-term episodes, which given the scarcity of data, makes making the most out of the available observations critical.

Inspired by the advancements in the domain of self-supervised learning, we propose a multi-modal self-supervised learning framework to learn the context of stress response from continuous physiological readings. This proposed setup addresses the following challenges and concerns regarding data-driven monitoring of stress and anxiety:

- The proposed method is inherently modular with regards to the different modalities of data, and therefore proper data-layer transforms allows leveraging various devices (e.g., smartwatches and wearable sensors different from ours) to learn efficient representations for health monitoring.
- The self-supervised component allows training the network with a higher level of granularity and makes training more efficient. This is especially needed as the amount of labeled data available is often limited and costly to acquire, in contrast to sensor data that is generally trivially available in large quantities.

- The use of the attention mechanism enables a diagnostic view of the system, allowing the researchers to look into the empirical connection between various modes of data for specific monitoring tasks, counteracting the masking effect of many deep-learning frameworks on interpretability.
- In developing this framework, we conducted experiments on real-world data collected on perceived stress and have shown that this approach improves the performance compared to prior work leveraging early-fused embeddings of the same benchmark dataset.

II. RELATED WORKS

Stress and anxiety-related disorders are common mental health challenges. Such disorders can have significant negative impacts on people's lives, including higher chances of depression and suicide as well as associated comorbidities with physical health issues [5]. Unfortunately, in many cases, these issues remain inadequately treated due to challenges ranging from lack of viable access to therapeutic services to associated stigmas with utilization [3], [4]. However, even when an individual decides to seek psychotherapeutic help to alleviate these problems, challenges persist in the diagnosis and effective treatment of their disorder. At the inception of care, the steps to diagnose and monitor often include clinical evaluation and comparing personalized symptoms to standardized criteria, for example, the Diagnostic and Statistical Manual of Mental Disorders (DSM-5), which is commonly used for this matter. Researchers continue to study and improve the practicality and accuracy of guidelines such as DSM-5 [12], [13], but there are challenges in converting aggregated and generalized diagnostic criteria, down to episodic-level incidents of stress and anxiety.

The most obvious approach to doing so leverages biometric data, extracted from wearable sensors embedded in smart devices, that measure a physiological stress response. However, while such data is incredibly valuable, and notably, sensing devices have become increasingly sophisticated at monitoring physiological stress, the resulting analyses are incomplete at best. This owes to the fact that from the standpoint of straightforward correlative analytics, it is known that there is not a direct monotonic correlation between the emotional perception of stress an individual may feel and the manifestation of the underlying physiological stress response. A meta-analysis in the social stress domains, for instance, has recently shown that merely 25% of studies in the domain demonstrated a significant correlation between physiological stress and perceived emotional stress [14].

Given that self-reports of perceived stress often do not contain information on the physiological stress response, understanding the complex relationship between the two becomes a crucial matter [14]. It is also plausible to assume that such complexity also arises from various other confounding factors (e.g., demographics, occupation, and other mental health disorders such as attention-deficit hyperactivity disorder (ADHD) can influence how prone someone is to stress).

This discrepancy has meaningful impacts on the utility of passive detection of stress based largely on physiological indicators. While sensors may be returning accurate readings on physiological stress, if they do not align with the user’s own perceptions of stress, notably if they fail to properly account for moments when a user feels acute emotional distress, then it will demotivate further engagement with a mental health platform.

This hindrance comes in spite of considerable progress that has been made in recent decades regarding the capabilities and efficacy of personal digital devices, including smartwatches, smartphones, and wearable devices. This fact has made such devices attract a lot of research and commercial attention, employing them for various monitoring objectives [15], [16].

These monitoring approaches focus primarily on fitness and health-related aspects, resulting in a large body of research and countless commercialized applications. Examples include tracking athletes’ training, detecting falls for the elderly, tracking post-surgery therapeutic and rehabilitation exercises, and posture correction [7], [17]–[19]. While the central focus of health monitoring applications has undoubtedly been on physical health, a wide range of research works has focused on understanding the relationship between observations obtained leveraging digital devices and some aspects of individuals’ mental health status. It is noteworthy that a primary goal in designing smart and automated approaches for mental health monitoring has to do with proposing meaningful passive-sensing tools so that informative observations regarding health status can be made by eliminating or diminishing the need to interfere with users’ daily activities or request repeated active interactions.

As a remote mental health monitoring task, social anxiety was studied in the previous literature, and it was shown that analyzing trajectories obtained via smartphone location services can paint a comprehensive picture concerning individuals’ proneness to it. To do so, the movements and the nature of locations visited (which were obtained by cross-referencing location data with a map API) were taken into consideration, and the hypothesis of whether or not such corpus is informative for recognizing the presence of social anxiety was tested [20]–[23]. Smartphones have also been helpful in developing an understanding of anxiety [24].

Another choice of hardware for gathering data pertinent to health data is application-specific wearable sensors. For instance, wearable electrocardiogram (ECG) sensors were used to recognize perceived anxiety via pattern recognition [25].

Smartwatches have a unique position amongst the wide range of various commonly used digital devices. They are in close contact with the skin and, given their attachment user’s wrist, which is a distal point of a major appendage, make it possible to obtain most measurements (e.g., activity) at higher accuracy, as well as enabling additional measurements such as heart-rate or pulse oximeter. In case of the need for brief questions, interactions, or Ecological Momentary Assessments (EMAs), smartwatches can also be used to issue messages and acquire responses and entries by the user [15], [16],

[26]. Additionally, smartwatches are prevalent, and relying on them as the hardware for health applications provides a better alternative in most cases to application-specific wearable devices in terms of cost, comfort, and user-friendliness. Data-driven analyses leveraging smartwatches’ sensory readings have been successful at the problem of patient classification for bipolar disorder, schizoaffective disorder, and depression [11]. It has also been shown that physiological readings made by basic smartwatch sensors enable efficient modeling of perceived stress response [27].

In the health analytics domain, data and human annotations are often limited. Therefore, dealing with overfitting and memorization is a crucial matter. Additionally, it is beneficial to go beyond the limited number of human annotations available in training efficient inference pipelines. Less reliance on annotations by focusing on unsupervised and self-supervised approaches has received a lot of research attention in recent years [28]–[31]. The core idea in most works in this area is that comparing and contrasting the latent representations of examples that are expected to share certain similarities (e.g., augmented versions of the same image) can benefit the trained weights and help with regularizing the learned decision boundaries [32].

In short, this work is primarily focused on addressing the limitations in the previous literature on remote mental health monitoring. The previous works do not go beyond leveraging scarcely available annotations in training network parameters and mainly rely on data augmentation to improve their performance. They do not focus on encapsulation in embedding different modalities, which can be an obstacle in employing optimal encoders for each modality and can hinder transfer learning. Additionally, they do not focus on the interpretability of the inference pipeline, which is crucial in health-related applications.

To address these challenges, this work proposes a framework for leveraging smartwatch-based sensor-driven data to recognize *perceived stress*, enabling a novel approach to remote mental health monitoring. Our proposed inference pipeline is modular and hierarchical and is composed of modality-specific embedding branches. The final embedding is computed via a task-specific attention-pooling mechanism, which also provides an interpretation of the estimated contribution of each modality’s information to the last embedding. During training, we leverage an inter-modality contrastive objective so as to encourage consistency among the predictions and tune all encoder branches. Figure 1 depicts the overview of our proposed framework. The details of our approach are discussed in the next section.

III. METHODOLOGY

Consider a cohort of P individuals undergoing a study wearing a smartwatch-based remote monitoring system. This wearable setup allows for the collection of sensory readings pertinent to users’ exhibited physiological and activity patterns throughout the day. In our case, the study in question monitors the connection between readings made by the smartwatch,

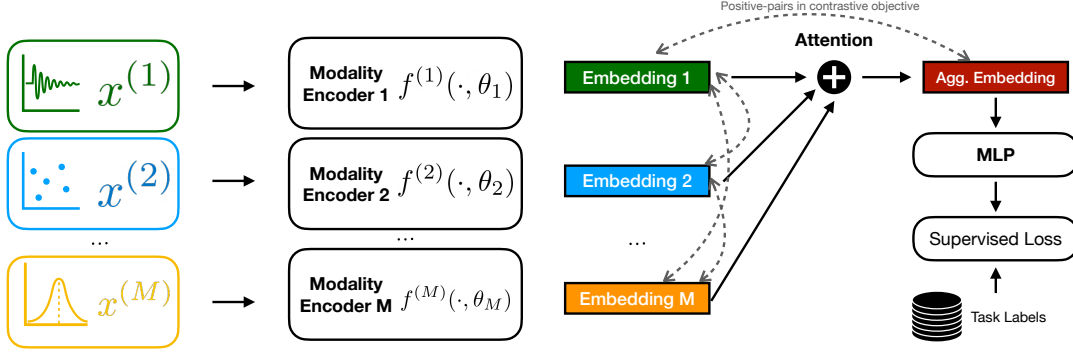


Fig. 1. Our proposed multi-modal self-supervised learning pipeline. Modality-specific data from different distributions are encoded through dedicated encoders and mapped to a shared latent representation space. The aggregated embedding of the segment is then computed by applying attention pooling to the modality-specific representations. A self-supervised contrastive objective aligns this aggregate embedding and the mode-specific representations.

which are mostly related to physiological signals, health and activity status, and short-term *perceived* stress reported by the user. The smartwatch’s readings can thus be grouped into features corresponding to several modalities: $m \in \{1, 2, \dots, M\}$. This grouping depends mainly on the nature of the features, as well as the setup and the interface provided by the smartwatch. From each modality, we have a sequence of observed feature vectors:

$$\mathbf{x}_m^{(p)} = \{x_{m,t}^{(p)}\}_{t \in [T_{m,\max}^{(p)}]} \quad (1)$$

From a user’s timeline, we extract short-term timespans, each of which corresponds to an *episode* e , which is the result of filtering the timeline and restricting it to the episode’s timespan: $e = (t_{\text{start}}, t_{\text{end}})$:

$$\mathbf{x}_{m,e}^{(p)} = \{x_{m,t}^{(p)} \in \mathbf{x}_m^{(p)} | t \in e\} \quad (2)$$

We have a parameterized domain-specific² encoder $f(\cdot; \theta_m)$ for each modality $m \in [M]$, which performs the task of mapping the observed data from this modality to a *shared* semantic space \mathcal{S} :

$$f(\cdot; \theta_m) : \mathcal{X}_m \rightarrow \mathcal{S} \quad \forall m \in [M] \quad (3)$$

Hence, the latent embedding denoted by $z_{m,e}^{(p)}$ can be found as follows:

$$z_{m,e}^{(p)} = f(\mathbf{x}_{m,e}^{(p)}; \theta_m) \quad \forall m \in [M] \quad (4)$$

One could argue that it is plausible to assume that the contributions of observations from different modalities to the final prediction on a specific task follow a non-uniform distribution in most cases. For instance, there is no reason to assume the statistical significance of heart-rate time series is the same as pulse oximeter readings for the task of stress detection. Going one step further, such disparity can manifest itself in the level of *instance* representations as well. To illustrate this further,

²The term *domain* in this manuscript refers to the observation type, for example, a Transformer-based Language Model could efficiently represent data from textual domain, and there could be multiple *modalities* with their observations being text data, each represented by their own specific encoder.

consider a simple case of “missingness” in data or presence of noise. This could mean that even though mode m_1 , for example, is more informative (in expectation) to the task τ , in an instance where the data from this group appears missing or clearly corrupt, the importance of other modalities could change respectively. Therefore, we have designed a *modality importance* head, implemented as a fully connected pipeline, which determines the contribution of each mode by weighing their respective embedding vectors, which were projected to the same semantic space. The first step to this attention-based pooling mechanism involves using the modality importance head and obtaining (yet unnormalized) weight $a_i^{(m)}$ for the latent embedding of modality m ’s information in an instance i :

$$a_i^{(m)} \leftarrow g(\mathbf{z}_i^{(m)}; \psi) \quad \forall m \in [M] \quad (5)$$

This is followed by a softmax operation to make sure that the summation of the predicted contributions maps to unity, in other words, the contribution matrix is right stochastic:

$$\alpha_i^{(m)} = \frac{\exp(a_i^{(m)})}{\sum_{j \in [M]} \exp(a_i^{(j)})} \quad (6)$$

And thus the final aggregated latent is computed using these attention weights: $z = \sum_{i=1}^M \alpha_i^{(m)} \cdot z_m$.

We leverage the cosine similarity $\phi(\cdot, \cdot)$ to measure the compatibility between the latent representation of each mode and the aggregate representation z_i .

$$\phi(\mathbf{u}, \mathbf{v}) = \frac{h(\mathbf{u})^T \cdot h(\mathbf{v})}{\|h(\mathbf{u})\|_2 \cdot \|h(\mathbf{v})\|} \quad (7)$$

In other words, we use the aggregated embedding z as an anchor and define a contrastive objective to leverage the distances and inconsistency between the latent embeddings:

$$\mathcal{L}_{\text{cl}} = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \frac{1}{|\mathcal{M}|} \sum_{m \in \mathcal{M}} -\log \frac{\exp(\phi(\mathbf{z}_i^{(m)}, \mathbf{z}_i)/\tau)}{\sum_{j \in \mathcal{B}, j \neq i} \exp(\phi(\mathbf{z}_i^{(m)}, \mathbf{z}_j)/\tau)} \quad (8)$$

We have experimented with \mathcal{L}_{cl} in the following training schemes:

- *Pre-training*: Pre-training the model parameters by optimizing \mathcal{L}_{cl} through a long training sequence. Afterward, start with the resulting weights as the initial point for the supervised fine-tuning of the model with the cross-entropy objective:

$$\mathcal{L}_{\text{cross-entropy}} = - \sum_{c \in \mathcal{C}} y_c \ln p_c \quad (9)$$

In the equation above, \mathcal{C} is the set of all classes (e.g., in our experiments, the two categories of stressed and non-stressed for each episode), and p_c is the predicted probability of class c for an observation, computed by passing representations through a final projection and Softmax layer.

- *Regularization*: Use $\lambda_{\text{reg}} \cdot \mathcal{L}_{cl}$ as a regularization term in the overall loss, and train the model by optimizing this loss simultaneously as the supervised learning objective.

There are several points worth remarking upon with regard to the comparison of these two training schemes. To begin with, deciding whether pre-training is going to lead to better generalization performance versus the regularization-based approach depends on model complexity, availability of data, and the challenges of the specific task that one is targeting. That being said, the regularization approach is expected to be considerably faster than the two-stage pre-training and fine-tuning method, and in our experiments on the task of predicting stress labels, it led to better test performance as well.

IV. EXPERIMENTS

A. Data

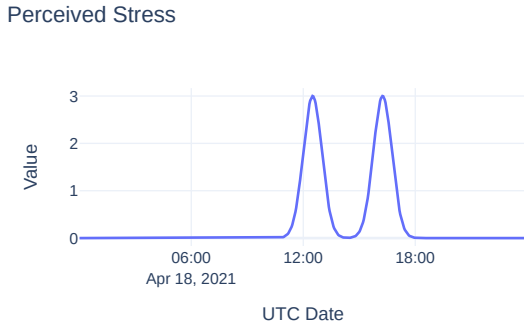


Fig. 2. A sample portion of a continuous supervision signal generated based on user inputs, from which episode stress labels can be sampled

The cohort in this study consists of 14 college students who are ex or active-duty members of the United States military. The status of these individuals, both as military members as well as students, renders them an interesting cohort for our stress study, given that the individuals from both groups are known to be relatively more prone to experiencing stress.

The smartwatch in this study was Garmin vivoactive 4S. Nevertheless, it is noteworthy that there is no component in the proposed methodology that limits the solution to this



Fig. 3. The average contribution of the four modalities to the final episode embeddings.

TABLE I
NUMBER OF PARTICIPANTS PER CATEGORY BASED ON THEIR DUTY STATUS AND SERVICE BRANCH

	Duty Status
National Guard	2
Active Duty	3
Veteran	9
	Service Branch
Airforce	2
Marine	3
Navy	4
Army	5

smartwatch. The feature groups and various modalities in our configuration are shown in Table II.

B. Labeling

The focus of this study has been on making predictions on *perceived* stress, for which the participants agreed to indicate the episodes in which they felt stressed and provide us with the intensity and timespans of these episodes.

For each record input to our system by an individual, we created a softened (via a Gaussian function) time-series per the following steps:

- The peak (corresponding to the *mean* of this Gaussian function) is set to the given timestamp, or the midpoint of the timespan $((t_{\text{start}} + t_{\text{end}})/2)$.

TABLE II
DATA MODALITIES AND THE FEATURES THEY ARE FOCUSED ON

Modality	Focus
Daily	heart-rate readings, number of floors climbed, BMR kilocalories, distance traveled, activity levels, aggregated HRV measures
Pulse Ox	SPO2
Respiration	Respiration rate
Stress	HRV-based readings

- The standard deviation of 30 minutes (scaled proportional to the length of time-span, if a time span of over one hour is provided).
- The magnitude of the peak point corresponds to the indicated for the episode: $\{0, 1, 2, 3\}$ for $\{\text{None, Low, Medium, High}\}$, respectively.

The summation of these Gaussian signals comprises the signal used as the primary supervision objective. The way the labels are computed is by looking at the end-point of each episode, and its *stress* label is marked `True` if the value of this signal on that point is larger than a threshold of 0.5, and `False` otherwise.

C. Modeling

Our inference model is composed of a specific encoder for each modality. In our case, each encoder is defined based on an initial mapping and normalization (via fully connected layers) followed by a bi-directional recurrent neural network (RNN) in long short-term memory (LSTM) configuration. Specifically, the data from each modality was first projected to a 32-dimensional vector via a multi-layer perceptron (MLP) with one hidden layer. The output was then forwarded to the modality-specific bi-LSTM with the hidden-layer neuron count of 64. The last stage for representing each modality was another fully-connected projection layer, generating a 32-dimensional vector per modality, which were used as modality representations in our framework.

Note that the overall pipeline does not have any constraint on the local modality encoders as long as they share the final semantic space to which they project that modality’s observations. Given that we were mainly dealing with time-series data, we used RNNs to model each branch. Nonetheless, modalities from substantially different domains and their encoders (e.g., Transformer-based Language Model for textual data) can also fit into the same system.

D. Results

Focusing on our real-world perceived stress corpus, we conducted experiments under the main settings of 1) supervised training baseline, 2) pre-training the contrastive objective and fine-tuning via supervised objective, and 3) training the supervised objective and simultaneously optimizing a scaled version of the contrastive term as a regularizing loss.

We observed that leveraging more features and following a late-fusion protocol for combining modality representations did lead to an improved generalization performance over the supervised setup proposed in [27], which combined the features at the beginning of the pipeline. In the case of our cohort, training with contrastive regularization led to the best generalization on the unseen test data, and the results are shown in Table III. Note that, in general, it is hard to say which self-supervised setup (pre-training versus regularization) is best, as it could depend on other factors, including model complexity, optimization, data availability, and task

TABLE III
PERFORMANCE COMPARISON FOR THE TRAINED PIPELINE UNDER DIFFERENT LEARNING SETUPS

Method	Test Accuracy
Early-fusion + Supervised Training [27]	64%
Late-fusion + Supervised Training	66%
Late-fusion + Contrastive Pre-training + Fine-tuning	70%
Late-fusion + Supervised Training + Contrastive Regularization	73%

difficulty. That being said, our approach allows learning high-quality representations by optimizing the modality-contrastive objective via both of these setups.

Additionally, we focused on interpretability as well and leveraged the task-specific attention mechanism in our pipeline, which pools the representations from different modalities, to study the utility and contribution of observations from each feature group. This enables the network to dynamically assign weights to each modality’s latent representation (in the shared space) as it processes each instance, allowing us to study their contribution both per instance and in expectation for performing the desired task. In Figure 3, we have shown the results on this matter for the contrastive regularization setup³. The results indicate that even though the contributions of the different modalities follow a non-uniform distribution as expected, none of them were ignored by the model and they all play a part in the final predictions.

V. DISCUSSION

A. Broader Impact

In the context of remote health monitoring, there are several factors addressing which is of paramount importance. In what follows, we elaborate upon these factors and how the solution proposed in this work attempts to address them:

- *Affordability and Compatibility*: For the scalability of a proposed remote health monitoring framework, focusing on widely available devices that are sold at affordable price renders it easier to deploy the system. In this work, we focused on basic physiological signals for which reading sensors are available in most commercially available smartwatches. Nonetheless, the proposed methodology has no intrinsic limitation regarding the modalities used; thus, additional data available in often more expensive devices (e.g., galvanic skin response) can also be utilized in the same methodology, and the main requirement is providing a modality-specific encoder fit for the data domain. Furthermore, this framework offers a more encapsulated view in representing different modalities as the observation from each can be embedded by a dedicated encoder first, and the contrastive objective encourages each local branch to optimize its parameters towards the given task as well. This has clear advantages in terms of transferring knowledge as well, an example of

³The label `heart` in Figure 3 corresponds to the `daily` modality’s information, given that its main focus is heart-rate.

which could be initializing each branch separately via pre-trained weights so as to prepare a better starting point for the model and optimization.

- *Ease of use*: Optimizing a remote health monitoring with regards to minimizing the amount of required user interaction makes it easier for individuals to use the system. This is why passive monitoring techniques are receiving more attention in the eHealth domain.
- *Interpretation*: In all automated healthcare applications of machine learning, any insight and interpretation into what parts of the observation a model mostly focused on in determining the final decision, is crucial and can help experts better validate the system as well. In this work, we incorporated a task-specific attention mechanism for pooling the representations from different modalities, which helps determine the weights assigned to each modality (per instance and in expectation) to perform the task efficiently.
- *Limited Data*: The data availability for eHealth applications is often limited due to the difficulty and costs of conducting large-scale studies, the exclusivity of data, and privacy reasons. It is, therefore, important to try to maximize the use of data in training inference pipelines. This work combines label smoothing with inter-mode self-supervision objectives to go beyond self-reported supervision objectives.

B. Limitations

It is crucial to discuss the limitations of this work given the sensitive nature of dealing with health as its objective. In this work, we relied on self-reported entries to decide the supervision signals for individual timelines. This has the issue of being prone to human error, as one might not accurately recall the time and extent to which one has felt stress. Additionally, reports on the intensity of the felt stress are also subject to noise. Another challenge is the small size of our dataset. A primary reason behind our self-supervision component in this work was alleviating the negative impacts associated with the aforementioned limitations.

C. Conclusion

We proposed a remote health monitoring solution that is modular and multi-modal, thus, allowing the use of various encoders best suited for each modality. We proposed an instance-level attention mechanism to tune the contribution of each modality to the final representation and provide insight into the expected importance of each modality for the task at hand. We conducted experiments with the proposed method to recognize perceived stress in short-term episodes and empirically demonstrated its superior performance over supervised training.

REFERENCES

- [1] K. Tsamakidis, A. S. Triantafyllis, D. Tsiptsios, E. Spartalis, C. Mueller, C. Tsamakidis, S. Chaidou, D. A. Spandidos, L. Fotis, M. Economou *et al.*, "Covid-19 related stress exacerbates common physical and mental pathologies and affects treatment," *Experimental and therapeutic medicine*, vol. 20, no. 1, pp. 159–162, 2020.
- [2] "Reducing the economic burden of unmet mental health needs," Jun 2022. [Online]. Available: <https://www.whitehouse.gov/cea/written-materials/2022/05/31/reducing-the-economic-burden-of-unmet-mental-health-needs/>
- [3] [Online]. Available: <https://www.nlm.nih.gov/health/statistics/any-anxiety-disorder>
- [4] 2020. [Online]. Available: <https://www.apa.org/news/press/releases/stress/2020/sia-mental-health-crisis.pdf>
- [5] 2018. [Online]. Available: <https://www.mayoclinic.org/diseases-conditions/anxiety/symptoms-causes/syc-20350961>
- [6] D. R. Seshadri, R. T. Li, J. E. Voos, J. R. Rowbottom, C. M. Alfes, C. A. Zorman, and C. K. Drummond, "Wearable sensors for monitoring the physiological and biochemical profile of the athlete," *NPJ digital medicine*, vol. 2, no. 1, pp. 1–16, 2019.
- [7] M. Gwak, S. Fazeli, G. Ershadi, M. Sarrafzadeh, M. Ghodsi, A. Aminian, and J. A. Schlechter, "Extra: exercise tracking and analysis platform for remote-monitoring of knee rehabilitation," in *2019 IEEE 16th International Conference on Wearable and Implantable Body Sensor Networks (BSN)*. IEEE, 2019, pp. 1–4.
- [8] F. R. Valente, M. de Paiva Guimarães, E. J. R. Cirilo, and D. R. C. Dias, "A multi-agent body tracking application framework applied to physical and neurofunctional rehabilitation," in *International Conference on Computational Science and Its Applications*. Springer, 2022, pp. 459–472.
- [9] M. Gwak, "Internet of things (iot)-enabled health monitoring systems: Implementation and validation," Ph.D. dissertation, University of California, Los Angeles, 2022.
- [10] Z. Zhao, D. Que, D. Jiang, and G. Shi, "Design of intelligent monitoring system for knee joint," in *Journal of Physics: Conference Series*, vol. 1449, no. 1. IOP Publishing, 2020, p. 012106.
- [11] S. Hassantabar, J. Zhang, H. Yin, and N. K. Jha, "Mhdeep: Mental health disorder detection system based on wearable sensors and artificial neural networks," *ACM Transactions on Embedded Computing Systems (TECS)*, 2022.
- [12] N. M. Batelaan, J. Spijker, R. de Graaf, and P. Cuijpers, "Mixed anxiety depression should not be included in dsm-5," *The Journal of nervous and mental disease*, vol. 200, no. 6, pp. 495–498, 2012.
- [13] M. Zimmerman, J. Martin, H. Clark, P. McGonigal, L. Harris, and C. G. Holst, "Measuring anxiety in depressed patients: A comparison of the hamilton anxiety rating scale and the dsm-5 anxious distress specifier interview," *Journal of psychiatric research*, vol. 93, pp. 59–63, 2017.
- [14] J. Campbell and U. Ehler, "Acute psychosocial stress: Does the emotional stress response correspond with physiological responses?" *Psychoneuroendocrinology*, vol. 37, pp. 1111–34, 01 2012.
- [15] A. Ng, M. Rddy, and S.M. Schueller, "Veterans' perspectives on fitbit use in treatment for post-traumatic stress disorder: An interview study," *MIR Mental Health*, 2018.
- [16] Texas A&M Engineering, "Veterans take wearable ptsd monitor on test ride," 2017. [Online]. Available: <https://today.tamu.edu/2017/06/05/veterans-take-wearable-ptsd-monitor-on-test-ride/>
- [17] S. Darabi, B. Moatamed, W. Huang, M. Gwak, C. Metoyer, M. Linn, and M. Sarrafzadeh, "Heart rate compression & time reduction method for hrv monitoring in athletes," in *2017 IEEE Healthcare Innovations and Point of Care Technologies (HI-POCT)*. IEEE, 2017, pp. 152–155.
- [18] T. Vilarinho, B. Farshchian, D. G. Bajer, O. H. Dahl, I. Egge, S. S. Hegdal, A. Lønes, J. N. Slettevold, and S. M. Weggelsen, "A combined smartphone and smartwatch fall detection system," in *2015 IEEE international conference on computer and information technology; ubiquitous computing and communications; dependable, autonomic and secure computing; pervasive intelligence and computing*. IEEE, 2015, pp. 1443–1448.
- [19] D. J. Wile, R. Ranaway, and Z. H. Kiss, "Smart watch accelerometry for analysis and diagnosis of tremor," *Journal of neuroscience methods*, vol. 230, pp. 1–4, 2014.
- [20] M. Boukhechba, P. Chow, K. Fua, B. A. Teachman, L. E. Barnes *et al.*, "Predicting social anxiety from global positioning system traces of college students: feasibility study," *JMIR mental health*, vol. 5, no. 3, p. e10101, 2018.
- [21] P. I. Chow, K. Fua, Y. Huang, W. Bonelli, H. Xiong, L. E. Barnes, and B. A. Teachman, "Using mobile sensing to test clinical models of depression, social anxiety, state affect, and social isolation among college students," *Journal of medical Internet research*, vol. 19, no. 3, p. e6820, 2017.

- [22] Y. Huang, H. Xiong, K. Leach, Y. Zhang, P. Chow, K. Fua, B. A. Teachman, and L. E. Barnes, "Assessing social anxiety using gps trajectories and point-of-interest data," in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2016, pp. 898–903.
- [23] J. Gong, Y. Huang, P. I. Chow, K. Fua, M. S. Gerber, B. A. Teachman, and L. E. Barnes, "Understanding behavioral dynamics of social anxiety among college students through smartphone sensors," *Information Fusion*, vol. 49, pp. 57–68, 2019.
- [24] L. M. Levine, M. Gwak, K. Kärkkäinen, S. Fazeli, B. Zadeh, T. Peris, A. S. Young, and M. Sarrafzadeh, "Anxiety detection leveraging mobile passive sensing," in *EAI International Conference on Body Area Networks*. Springer, 2020, pp. 212–225.
- [25] Z. King, J. Moskowitz, L. Wakschlag, and N. Alshurafa, "Predicting perceived stress through mirco-emas and a flexible wearable ecg device," in *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*, 2018, pp. 106–109.
- [26] J. Chum, M.S. Zielinski, M. Bhatt, D. Chung, S.Yeung and Z. Samaan, "Acceptability of the fitbit in behavioural activation therapy for depression: a qualitative study." *Evidence-Based Mental health*, 2017.
- [27] S. Fazeli, L. Levine, M. Beikzadeh, B. Mirzasoleiman, B. Zadeh, T. Peris, and M. Sarrafzadeh, "Passive monitoring of physiological precursors of stress leveraging smartwatch data," in *2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 2022, pp. 2893–2899.
- [28] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [29] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9912–9924, 2020.
- [30] X. Chen and K. He, "Exploring simple siamese representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15 750–15 758.
- [31] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar *et al.*, "Bootstrap your own latent-a new approach to self-supervised learning," *Advances in neural information processing systems*, vol. 33, pp. 21 271–21 284, 2020.
- [32] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," *arXiv preprint arXiv:1710.09412*, 2017.