

Accounting for Affect in Pain Level Recognition

Md Taufeeq Uddin

University of South Florida, Tampa, FL, USA

MDTAUFEEQ@USF.EDU

Shaun Canavan

University of South Florida, Tampa, FL, USA

SCANAVAN@USF.EDU

Ghada Zamzmi

*University of South Florida, Tampa, FL, USA, current affiliation
National Institutes of Health*

ALZAMZMIGA@NIH.GOV

Abstract

In this work, we address the importance of affect in automated pain assessment and the implications in real-world settings. To achieve this, we curate a new physiological dataset by merging the publicly available bioVid pain and emotion datasets. We then investigate pain level recognition on this dataset simulating participants' naturalistic affective behaviors. Our findings demonstrate that acknowledging affect in pain assessment is essential. We observe degradation in recognition performance when simulating the existence of affect to validate pain assessment models that do not account for it. Conversely, we observe a performance boost in recognition when we account for affect.

Keywords: Pain assessment, affective computing, healthcare

1. Introduction

Automated pain assessment is an essential component in clinical settings to ensure prompt intervention and appropriate treatment. Depending on the context, relying on verbal methods such as pain scales and questionnaires [Williamson and Hoggart \(2005\)](#) for pain assessment may not be objective, reliable, actionable, and scalable [Egede and Valstar \(2017\)](#). Further, babies and uncon-

scious patients are unable to provide verbal responses, while verbal responses from a mentally impaired adult may be unreliable. To tackle this challenge, studies have been performed to automate pain assessment [Walter et al. \(2014\)](#); [Lötsch and Utsch \(2018\)](#).

There has been significant progress in automated pain assessment using physiological data [Gruss et al. \(2015\)](#); [Werner et al. \(2014\)](#), visual data [Werner et al. \(2016\)](#); [Zhao et al. \(2016\)](#), and fusion of them [Egede et al. \(2017\)](#); [Fabiano and Canavan \(2019\)](#) for either detecting pain presence [Lucey et al. \(2011\)](#); [Mansoor et al. \(2018\)](#) or estimating pain intensity [Martinez et al. \(2017\)](#); [Uddin and Canavan \(2020a\)](#); [Salekin et al. \(2020\)](#). These methods achieved good performance when applied to existing publicly available dataset such as UNBC-McMaster Pain dataset [Lucey et al. \(2011\)](#), bioVid pain dataset [Walter et al. \(2013\)](#), and MIntPain dataset [Haque et al. \(2018\)](#).

These previous works are limited as they only compare pain to a neutral state (baseline). This can lead to the failure of these approaches in real-world settings where patients are likely to be in multiple affect states aside from being in pain or neutral. For example, it has been reported that hospitalized patients undergo a wide range of affects including anger, depression, anxiety in addi-

tion to the pain affect [Mirani et al. \(2019\)](#). These affect states can occur as a result of disease, stress, financial concerns, and unsatisfactory treatment [Mirani et al. \(2019\)](#). Patients can also experience different positive affects, such as comfort or joy, as a result of interactions with family/caregivers [Eggins \(2014\)](#). These studies provide evidence that affect is relevant in pain assessment tasks [Gatchel et al. \(2007\)](#), and motivate us to consider affect states while developing pain assessment models. Such models would lead to a reliable and accurate assessment when applied to patients in clinical settings.

In this paper, we propose to account for affect when assessing pain to provide a robust assessment in real-world settings given that the model is aware of other affective states. We curate a new pain-affect dataset we refer to as **bioVid pain-affect** (Section 2). This allows us to build a reliable and robust pain assessment model that considers multiple affect states commonly experienced by patients (e.g., sadness, fear) in clinical settings. Our experimental results on the pain-affect dataset indicate that pain assessment models that do not take affect into consideration are likely to drift and make poor inference when applied to real-world settings. In addition to reliable real-world assessment, the proposed modeling of pain allows to adjust pain score/intensity based on underlying affect of the patient; e.g., it has been shown that there is altered pain sensitivity in patients who experience pre-existing anxiety and depression [Sullivan et al. \(2017\)](#); [Hermesdorf et al. \(2016\)](#). Finally, it can lead to the development of personalized treatment plans as it accounts for underlying states of each patient.

2. bioVid Pain-Affect Dataset

We curate a new dataset by merging the publicly available bioVid pain [Walter et al.](#)

[\(2013\)](#) and bioVid emotion [Zhang et al. \(2016\)](#) datasets. The pain dataset contains two sessions of recording participants pain and baseline states. The emotion dataset contains one session of recording participants discrete affects including amusement, anger, disgust, fear, and sadness. The data were collected in the same lab settings using similar equipment. Participants are within the age range of [19, 65] years and in total there are 91 subjects (45 female and 46 male) across both datasets. Further description of the datasets and stimuli for spontaneous pain and affect states can be found in Appendix A, Section 2.

To merge the datasets, we selected the common subjects that appear in both the pain and emotion datasets (82 subjects in total). We then removed 20 participants from the merged dataset, who did not show pain responses during the pain stimulus [Werner et al. \(2017\)](#), resulting in 62 participants (33 female and 29 male). Merging of the datasets resulted in a new age range of [20, 65] with a median age of 36. The merged bioVid pain-affect dataset contains three biopotential signals: skin conductance level or electrodermal activity (**EDA**), electrocardiogram (**ECG**), electromyogram (**EMG**) of trapezius muscle.

3. Pain Assessment Model

3.1. Problem Formulation

Our work aims to recognize pain level while taking into account affect states that exist in real-world settings such as anger and anxiety. Recall that we combine discrete affect states (e.g., sadness, fear) to one affect state named **A**. We hypothesize that the consideration of different affect while assessing pain is necessary to obtain accurate results, and for transparency. To do so, we propose to incorporate affect in PL recognition model as a category (e.g., *A* in our studied dataset) along multiple pain levels (e.g., **LLP** - low

level pain, **HLP** - high level pain in our studied dataset; see Appendix A Section 2 for details). Depending on the context, baseline (**BL**) category could be incorporated in affect category as BL is likely to be a neutral/relaxed affect state, and for some subjects relaxed states could be other positive or negative affect states (e.g., amusement, fear) depending on the situation, environment, or time.

3.2. Feature Representation

To create a feature vector for a given sample, we downsampled the biopotential signals by computing the moving average using a sliding window with 80% overlap. Before downsampling the signals, we used a Savitzky–Golay filter [Press and Teukolsky \(1990\)](#) to remove noise from the raw signals. We also normalized each signal in the range $[0, 1]$ using min-max normalization. We then created a normalized feature vector as $f_m = [c_1, c_2, \dots, c_n]$, where c and n are the downsampled signal frames/components and length of downsampled signal, respectively.

3.3. Case studies

Case -1. We explore the current naive PL recognition (**PLR**) approaches, in which we use only BL and pain data during training, validation and testing data.

Case 0. We simulate real-world settings on the naive PL recognition approach. As we discussed before, a person is likely to experience pain, baseline, and other affects in real-world settings. Hence, we reproduce the PLR model developed in case -1, and tested the model assuming the affect exist in our testing dataset.

Case 5. This case is based on our proposal (Section 3.1). In this case, we acknowledge during the development of PLR model and during testing in real-world settings, that a person would experience affect (e.g., sadness, fear), baseline, and multiple levels of pain. In

our training, validation, and testing datasets, we used affect and pain data. In the affect data, we incorporate the BL state assuming it is part of affect state.

Case 6. This case is a variation of case 5 in which we remove the baseline data from other affect data. This case allows us to investigate the performance of PLR model in absence of baseline, and in presence of other affect. Also, in this case, we also investigate the impact of a balanced class distribution among classes (e.g., LLP, HLP, A) in the training dataset. Contrary to cases -1, 0, 5, where we preserve the training dataset uniformity in terms of number of examples and class proportion to ensure fair evaluation.

4. Experiments and Analysis

Evaluation setup. To validate the model, we assumed that participants were known during validation. Hence, we randomly split the whole dataset into 70% training and 30% test data, and we ran the experiments 5 times with different seeds. F1 macro was used to report the model performance.

Modalities and Classifiers. To investigate the performance of PLR models in above mentioned cases, we built the models in unimodel and multimodel settings, i.e., we trained and tested PLR model on EDA, ECG, EMG separately, and their combination (EDA + ECG + EMG). We experimented with three well-known classification algorithms: K Nearest Neighbors (**KNN**, $k = 5$) [Devroye et al. \(2013\)](#), Random Forest (**RF**, tree estimator = 750) [Breiman \(2001\)](#), and Extreme Gradient Boosting (**XGB**, tree estimator = 750 and learning rate = 0.1.) [Chen and Guestrin \(2016\)](#).

Case Studies Evaluation. Recall that in case -1 and case 0, we trained the naive PLR model, i.e. model does not know about affect (e.g. anger, sadness). Also, training data size and class proportion were kept uni-




Figure 1: PLR model performance on **known participants** (we assumed participants are likely to know) in investigated cases. From left to right, figures represent results on **EDA**, **ECG**, **EMG**, and **EDA + ECG + EMG** modalities, respectively.

form in case $-1, 0$, and 5 for fair evaluation. When we tested the naive PLR model in case 0 (simulated the existence of affects; see Figure 1, the model showed drift in terms of recognition rate compared to the results obtained in case -1 . This can be explained, in part, due to the PLR model not being aware of the affect data during training. Some evidence behind this issue could be found from Figures 2 and 3 (Appendix A). Since we are missing out on affect data in the naive modeling case, it is extremely hard for the naive PLR model to generalize or adapt to the affect data. More precisely, during inference in simulated natural world, the trained naive PL recognition model seemed to treat the affect data as unknown data.

In case 5 , we take the affect into account in our PLR model. Note that test dataset is exactly same in both case 5 and case 0 . As can be seen in Figure 1, the model showed significant improvement in terms of recognition performance compared to case 0 . This can be explained, in part, due to the affect distribution being known to the proposed PLR model.

Finally, recall that case 6 was designed to investigate the influence of a balanced training dataset in terms of class proportion. Also, to evaluate the proposed model in the absence of BL data in both training and test

dataset. In this case, we are seeing an improving in recognition performance, which can be explained by the evenly distribution of LLP, HLP and A categories in training dataset.

Section 4 in Appendix A presents further experiments, evaluation and analysis of PLR models using numerous settings.

5. Conclusion

This work investigated the necessity of including non-pain affect in pain level recognition (PLR), which is a major limitation in current PLR approaches. Using just a baseline during PLR modeling (case -1 in Section 3.3) produced reasonable results in a controlled environment. However, as evidenced from our experiments (Figures 1 above, and Figures 4-7 in Appendix A), these naive models fail in natural settings as people experience a diverse set of affects apart from neutral and pain. Hence, we conclude that incorporating affect in PLR approaches is crucial. The main limitation of this study is the utilization of only physiological signals. In the future, we will explore behavioral data such as face images, facial expressiveness dynamics Uddin and Canavan (2020b) and facial action units Ekman (1997) along with the combination of physiological and behavioral data.

References

- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.
- Ruth Cowen, Maria K Stasiowska, Helen Laycock, and Carsten Bantel. Assessing pain objectively: the use of physiological markers. *Anaesthesia*, 70(7):828–847, 2015.
- Luc Devroye, László Györfi, and Gábor Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.
- Joy Egede, Michel Valstar, and Brais Martinez. Fusing deep learned and hand-crafted features of appearance, shape, and dynamics for automatic pain estimation. In *2017 12th IEEE international conference on automatic face & gesture recognition (FG 2017)*, pages 689–696. IEEE, 2017.
- Joy O Egede and Michel Valstar. Cumulative attributes for pain intensity estimation. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*, pages 146–153, 2017.
- Suzanne Eggins. Hospital humour: Patient initiated humour as resistance to clinical discourse. In *Intersections: Applied linguistics as a meeting place*, volume 43, page 66. Cambridge Scholars Newcastle Upon Tyne, UK, 2014.
- Rosenberg Ekman. *What the face reveals: Basic and applied studies of spontaneous expression using the Facial Action Coding System (FACS)*. Oxford University Press, USA, 1997.
- Diego Fabiano and Shaun Canavan. Emotion recognition using fused physiological signals. In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 42–48. IEEE, 2019.
- Robert J Gatchel, Yuan Bo Peng, Madelon L Peters, Perry N Fuchs, and Dennis C Turk. The biopsychosocial approach to chronic pain: scientific advances and future directions. *Psychological bulletin*, 133(4):581, 2007.
- Sascha Gruss, Roi Treister, Philipp Werner, Harald C Traue, Stephen Crawcour, Adriano Andrade, and Steffen Walter. Pain intensity recognition rates via biopotential feature patterns with support vector machines. *PloS one*, 10(10), 2015.
- Mohammad A Haque, Ruben B Bautista, Fatemeh Noroozi, Kaustubh Kulkarni, Christian B Laursen, Ramin Irani, Marco Bellantonio, Sergio Escalera, Golamreza Anbarjafari, Kamal Nasrollahi, et al. Deep multimodal pain recognition: a database and comparison of spatio-temporal visual modalities. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 250–257. IEEE, 2018.
- Marco Hermesdorf, Klaus Berger, Bernhard T Baune, Jürgen Wellmann, Ruth Ruscheweyh, and Heike Wersching. Pain sensitivity in patients with major depression: differential effect of pain sensitivity measures, somatic cofactors, and disease characteristics. *The Journal of Pain*, 17(5):606–616, 2016.
- Li Hu and Gian Domenico Iannetti. Painful issues in pain prediction. *Trends in neurosciences*, 39(4):212–220, 2016.
- Mahdi Khezri, Mohammad Firoozabadi, and Ahmad Reza Sharafat. Reliable emo-

- tion recognition system based on dynamic adaptive fusion of forehead biopotentials and physiological signals. *Computer methods and programs in biomedicine*, 122(2): 149–164, 2015.
- Jörn Lötsch and Alfred Ultsch. Machine learning in pain research. *Pain*, 159(4): 623, 2018.
- Patrick Lucey, Jeffrey F Cohn, Kenneth M Prkachin, Patricia E Solomon, and Iain Matthews. Painful data: The unbc-mcmaster shoulder pain expression archive database. In *Face and Gesture 2011*, pages 57–64. IEEE, 2011.
- Zara Mansoor, Mustansar Ali Ghazanfar, Syed Muhammad Anwar, Ahmed S Alfa-keeh, and Khaled H Alyoubi. Pain prediction in humans using human brain activity data. In *Companion Proceedings of the The Web Conference 2018*, pages 359–364, 2018.
- Daniel Lopez Martinez, Ognjen Rudovic, and Rosalind Picard. Personalized automatic estimation of self-reported pain intensity from facial expressions. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2318–2327. IEEE, 2017.
- Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- Shahid H Mirani, Dharmoon Areja, Syeda Suman Gilani, Amber Tahir, Murk Pathan, and Suman Bhatti. Frequency of depression and anxiety symptoms in surgical hospitalized patients. *Cureus*, 11(2), 2019.
- W. Press and S. Teukolsky. Savitzky-golay smoothing filters. *Computers in Physics*, 4(6):669–672, 1990.
- Md Sirajus Salekin, Ghada Zamzmi, Dmitry Goldgof, Rangachar Kasturi, Thao Ho, and Yu Sun. First investigation into the use of deep learning for continuous assessment of neonatal postoperative pain. *arXiv preprint arXiv:2003.10601*, 2020.
- Erin Sullivan, Jordin Shelley, Evan Rainey, Monica Bennett, Purvi Prajapati, Mark B Powers, Michael Foreman, and Ann Marie Warren. The association between post-traumatic stress symptoms, depression, and length of hospital stay following traumatic injury. *General hospital psychiatry*, 46:49–54, 2017.
- Md Taufeeq Uddin and Shaun Canavan. Multimodal multilevel fusion for sequential protective behavior detection and pain estimation. *IEEE Automatic Face and Gesture Recognition (FG)*, 2020a.
- Md Taufeeq Uddin and Shaun Canavan. Quantified facial temporal-expressiveness dynamics for affect analysis, 2020b.
- Steffen Walter, Sascha Gruss, Hagen Ehleiter, Junwen Tan, Harald C Traue, Philipp Werner, Ayoub Al-Hamadi, Stephen Crawcour, Adriano O Andrade, and Gustavo Moreira da Silva. The biovid heat pain database data for the advancement and systematic validation of an automated pain recognition system. In *2013 IEEE international conference on cybernetics (CYBCO)*, pages 128–131. IEEE, 2013.
- Steffen Walter, Sascha Gruss, Kerstin Limbrecht-Ecklundt, Harald C Traue, Philipp Werner, Ayoub Al-Hamadi, Nicolai Diniz, Gustavo Moreira da Silva, and Adriano O Andrade. Automatic pain

quantification using autonomic parameters. *Psychology & Neuroscience*, 7(3): 363–380, 2014.

Philipp Werner, Ayoub Al-Hamadi, Robert Niese, Steffen Walter, Sascha Gruss, and Harald C Traue. Automatic pain recognition from video and biomedical signals. In *2014 22nd International Conference on Pattern Recognition*, pages 4582–4587. IEEE, 2014.

Philipp Werner, Ayoub Al-Hamadi, Kerstin Limbrecht-Ecklundt, Steffen Walter, Sascha Gruss, and Harald C Traue. Automatic pain assessment with facial activity descriptors. *IEEE Transactions on Affective Computing*, 8(3):286–299, 2016.

Philipp Werner, Ayoub Al-Hamadi, and Steffen Walter. Analysis of facial expressiveness during experimentally induced heat pain. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, pages 176–180. IEEE, 2017.

Amelia Williamson and Barbara Hoggart. Pain: a review of three commonly used pain rating scales. *Journal of clinical nursing*, 14(7):798–804, 2005.

Lin Zhang, Steffen Walter, Xueyao Ma, Philipp Werner, Ayoub Al-Hamadi, Harald C Traue, and Sascha Gruss. “biovid emo db”: A multimodal database for emotion analyses validated by subjective ratings. In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 1–6. IEEE, 2016.

Rui Zhao, Quan Gan, Shangfei Wang, and Qiang Ji. Facial expression intensity estimation using ordinal information. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3466–3474, 2016.

Appendix A

Section 2: Further bioVid Pain-Affect Dataset Details

Stimulus. To elicit spontaneous pain, in the bioVid pain dataset, self-calibrated heat was applied as the stimulus. There are five pain levels in the pain dataset including baseline (BL) - pain level (**PL**) = 0, pain threshold - $PL = 1$, pain tolerance level - $PL = 4$; after collecting data from the participants, two intermediate pain levels were selected: $PL = 2$ and $PL = 3$. In the bioVid emotion dataset, video clips from movies were used to elicit spontaneous discrete emotions including amusement (Am), anger (An), disgust (Di), fear (F) and sadness (S).

Modalities. The bioVid emotion dataset contains three biopotential signals: EDA, ECG, EMG of trapezius muscle, and videos of participants’ frontal face. In addition to the above mentioned modalities, bioVid pain dataset contains EMG signal collected from corrugator and zygomaticus muscles. In merged bioVid pain-affect dataset, we only selected the common biopotential signals (e.g., EDA, ECG, EMG from trapezius muscle) and videos. In this work, we focus only on the biopotential signals as studies in healthcare Khezri et al. (2015); Cowen et al. (2015) indicated that biopotential signals are major objective indicators of pain and other affect.

Sample length. The bioVid pain dataset contains 5.5 seconds long subsamples (from the raw data) along with raw data, while the bioVid emotion dataset contains raw data (range = [32, 245] seconds long) only. In this work, we extracted 5.5 seconds long subsamples from the raw emotion data following the pain dataset to have same length samples for pain, baseline, and affect in our merged dataset.

Discrete categories. As pointed out by Werner et al. Werner et al. (2017), the vi-

sual responses of PL 1 and PL 2 in bioVid pain dataset were difficult to distinguish due to the similarity in terms of responses and some subjects not showing any affective responses. Further, PL 3 and PL 4 had similar responses to a certain extent. Based on the those findings, in this new merged dataset, we merged PL 1 and PL 2 to one level named low level pain (LLP), and PL 3 and PL 4 to another pain level named high level pain (HLP). Since our primary goal is to investigate the influence of affect in PL recognition approaches we merged all discrete affects (e.g., Am, An, Di) to one specific affect class named Affect (A). Hence, in the new bioVid pain-affect dataset, we have four discrete categories: BL, LLP, HLP, and A.

Section 4: Further Experiments

Evaluation setup. Since the major goal of this study is to investigate the limitation of previous PLR approach, and evaluate the effectiveness of proposed approach, we performed two additional types of model validations: i) **evaluation on unknown / new participants**: we used leave-n-participant-out validation, where $n = 15$ in our experiments to evaluate all studied cases on participants who are unknown during training; ii) **person-specific evaluation**: in this evaluation, we trained and tested on each participant under the assumption that participants are different in terms of eliciting affective behavior, hence, for a given participant’s data, we split the data into 70% training and 30% test data, and we ran the experiments 5 times with different seeds. We used F1 macro (macro-average F1 measure) to report the model performance given its robustness towards balanced and unbalanced dataset.

Generalization to unknown, known participants, and personalization. The following experiments were performed to evaluate whether we could build generaliz-

able PLR model that can generalize to everyone as people are likely to be different in terms of eliciting affects. Comparing Figures 1 to 4, and Figures 6 to 7, we can see that when we assumed participants are known we are seeing improvement in performance, this could be due to the subjective variability. In Figure 3, we can see that it is easier to distinguish between B, LLP, HLP, and A in person-specific case, however, when sampled from multiple people (demography-specific, general), it is harder to distinguish.

An interesting pattern can be observed is that as we are increasing the number of subjects in samples, it is getting more and more difficult to distinguish between affective states, which in turn likely to makes generalizations to unknown participants harder. We also looked into demographic information as context to evaluate the influence of gender, age and demographic group ¹. The results suggest that demographic information showed some improvement in PL recognition performance in known participants setting as presented in Figures 7 and 6.

Figure 5 shows that the PL recognition model gave encouraging results for some participants, however, for others even though model were personalized (trained and tested on each participant), it is challenging to distinguish between multiple states. There could be multiple explanations for this. First, it is possible those subjects did not feel affect or pain during experiments. Secondly, they may have been going through some other affective states that was not captured in the data.

Lastly, the physiological data was not different across the affective states for those

participants (i.e. pain and affect looked similar). Considering this, we need to further look into the data to make any concluding remarks. Finally, in all evaluations (unimodal and multimodal features with / without incorporation of context), we observed improvement when we accounted for affect in PLR models compared to cases when affect were not accounted.

Based on our experiments, we can opine that the effective and reliable way to solve the problem is to incorporate affect in PL recognition approaches. As we can see in Figure 3, pain level recognition when accounting for affect is a challenging task due to the similarity of data patterns of pain and non-pain affect, as well as the variability of data based on people differences [Hu and Iannetti \(2016\)](#). The proposed formulation still performed comparatively better due to its ability to adapt to affect dataset. Hence, even though incorporation of affect could make the modeling challenging, it would make the modeling sound, reliable, and realistic.

1. Demographic group is constructed via grouping participants based on their age, and gender. Female aged in between [20,30), [30,50), and [50,65) are grouped and labeled as F_1 , F_2 , F_3 , respectively. Male in similar age categories are grouped and labeled as M_4 , M_5 , M_6




Figure 2: Distribution of physiological signals depicting affective states. Here, standardized (zero-mean and unit variance) data were sampled from one participant. We can see that a subject could be in BL, LLP, HLP and A states in different moment of time in natural world.




Figure 3: Latent space representation of physiological data depicting affective states, using UMAP (uniform manifold approximation and projection) algorithm [McInnes et al. \(2018\)](#). From left to right, figures depict **person-specific**, **demography-specific** and **general** scenarios. A grid of values were used to tune the major parameters (nearest neighbors = [2, 5, 10, 20, 50, 100, 200], and minimum distance = [0.0, 0.1, 0.25, 0.5, 0.8, 0.99]) of UMAP. In person-specific scenario, data were sampled from one participant; demography-specific scenario, data were sampled from one specific demographic group (see footnote 1); general scenario, data were sampled from all 62 participants.




Figure 4: PLR model performance on **unknown** participants in investigated cases. From left to right, figures represent results on **EDA**, **ECG**, **EMG**, and **EDA + ECG + EMG** modalities, respectively.

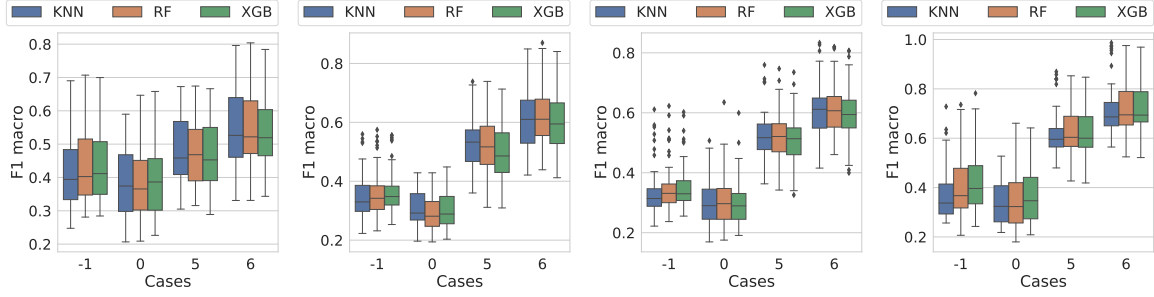


Figure 5: Personalized PLR model performance in investigated cases. From left to right, figures represent results on **EDA**, **ECG**, **EMG**, and **EDA + ECG + EMG** modalities, respectively. Each point in each box-whisker plot indicates F1-macro computed on each participant.




Figure 6: PLR model performance on **known participants** (we assumed participants are likely to know) in investigated cases. From left to right, figures represent results on **demographic information with EDA**, **ECG**, **EMG**, and **EDA + ECG + EMG** modalities, respectively. Demographic information were incorporated in feature vector as context.




Figure 7: PLR model performance on **unknown participants** in investigated cases. From left to right, figures represent results on **demographic information with EDA**, **ECG**, **EMG**, and **EDA + ECG + EMG** modalities, respectively. Demographic information were incorporated in feature vector as context.