Pain versus Affect? An Investigation in the Relationship between Observed Emotional States and Self-Reported Pain

Fu-Sheng Tsai* Yi-Ming Weng† Chip-Jin Ng‡ and Chi-Chun Lee§
* Department of Electrical Engineering, National Tsing Hua University, Taiwan
MOST Joint Research Center for AI Technology and All Vista Healthcare, Taiwan
E-mail: fstsai@gapp.nthu.edu.tw, cclee@ee.nthu.edu.tw
† Department of Emergency Medicine, Chang Gung Memorial Hospital, Taiwan
Department of Emergency Medicine, Tao-Yuan General Hospital, Taiwan
Faculty of Medicine, National Yang-Ming University, Taiwan
‡ Department of Emergency Medicine, Chang Gung Memorial Hospital, Taiwan

***Abstract*—Pain is an internal sensation intricately intertwined with individual affect states resulting in a varied expressive behaviors multimodally. Past research have indicated that emotion is an important factor in shaping one's painful experiences and behavioral expressions. In this work, we present a study into understanding the relationship between individual emotional states and self-reported pain-levels. The analyses show that there is a significant correlation between observed valence state of an individual and his/her own self-reported pain-levels. Furthermore, we propose an emotion-enriched multitask network (EEMN) to improve self-reported pain-level recognition by leveraging the rated emotional states using multimodal expressions computed from face and speech. Our framework achieves accuracy of 70.1% and 52.1% in binary and ternary classes classification. The method improves a relative of 6.6% and 13% over previous work on the same dataset. Further, our analyses not only show that an individual's valence state is negatively correlated to the pain-level reported, but also reveal that asking observers to rate valence attribute could be related more to the self-reported pain than to rate directly on the pain intensity itself.**

I. Introduction

Pain is a subjective internal sensation that is closely associated with bodily damages or physical illnesses. There are numerous factors affecting the intensity of one's painful episodes, e.g., memories or experiences of previous pain, social cultural background, expectations and attitudes towards pain, age, gender, race, even emotional or psychological state [1], [2], [3]. Since emotional experiences are often connected to the sensory processing regions in the brain, understanding the relationships between pain and emotion have received increased attention. For example, Carter *et al*. proposes that self-reported pain-levels is related to the depression condition, i.e., subjects who are in the anxiety or depression condition show reduced pain tolerance after being induced into negative affect state [4]. Other study has also pointed out negative emotion state is associated to an increase in the pain felt, whereas positive state lowers pain [5].

While being a subjective internal sensation, quantitative pain assessment is a critical component for providing effective pain management and screening life-threatening patients in clinical applications. While both self-report and observation-based pain assessment instruments have been developed [6], studies have demonstrated the self-reported numerical scale (NRS) remains to be the most valid measure in clinical practices [7], [8], [9]. However, evidences have indicated the measure not only suffers from inconsistent administration in clinical practices, but limits the large-scale applicability of pain assessment for healthcare applications [10].

With the advancements in machine learning and signal processing techniques, researchers have started to develop algorithms to automate the pain-level assessment. For example, Rodriguez *et al*. proposes a long short term memory (LSTM) network to estimate individual pain-level using facial images [11], Tavakolian *et al*. develops a deep convolutional neural network (CNN) architecture that integrates low-level visual descriptors and high-level structural information to estimate pain intensity [12], and Egede *et al*. performs pain-level recognition by using both hand-crafted facial features and deep neural features on images [13]. Furthermore, our recent works have focused on modeling speech acoustic information in a real patient's triage corpus [14]. Tsai *et al*. uses bottleneck LSTM on prosodic cues to perform pain recognition on subset of this triage corpus [15], and Li *et al*. further extends the framework to the entire corpus by incorporating age and gender attributes into variational acoustic network [16]. Most of these past works focus on learning to map behavior cues to self-reported NRS, there is yet a systematic modeling between the relationship of affect states and self-reported pain-levels.

In this work, we present a computational study to investigate the relationship between observed affect states and self-reported pain-levels. We recruited five evaluators to assess the scale of emotion primitives (activation and valence) and the observed pain-level for each sample in our triage multimodal pain database. Our analysis demonstrates that the observed valence attribute is negatively correlated to the self-reported pain-levels ($\rho = -0.5136$) corroborating past literature [17]. Furthermore, we develop an emotion-enriched multitask network (EEMN), which aims at improving self-reported NRS classification by jointly optimizing with affect states using multimodal behavior data (acoustic and facial expressions). Our EEMN framework achieves a 70.1% unweighted accuracy in classifying extreme set (mild versus severe) and 52.1%

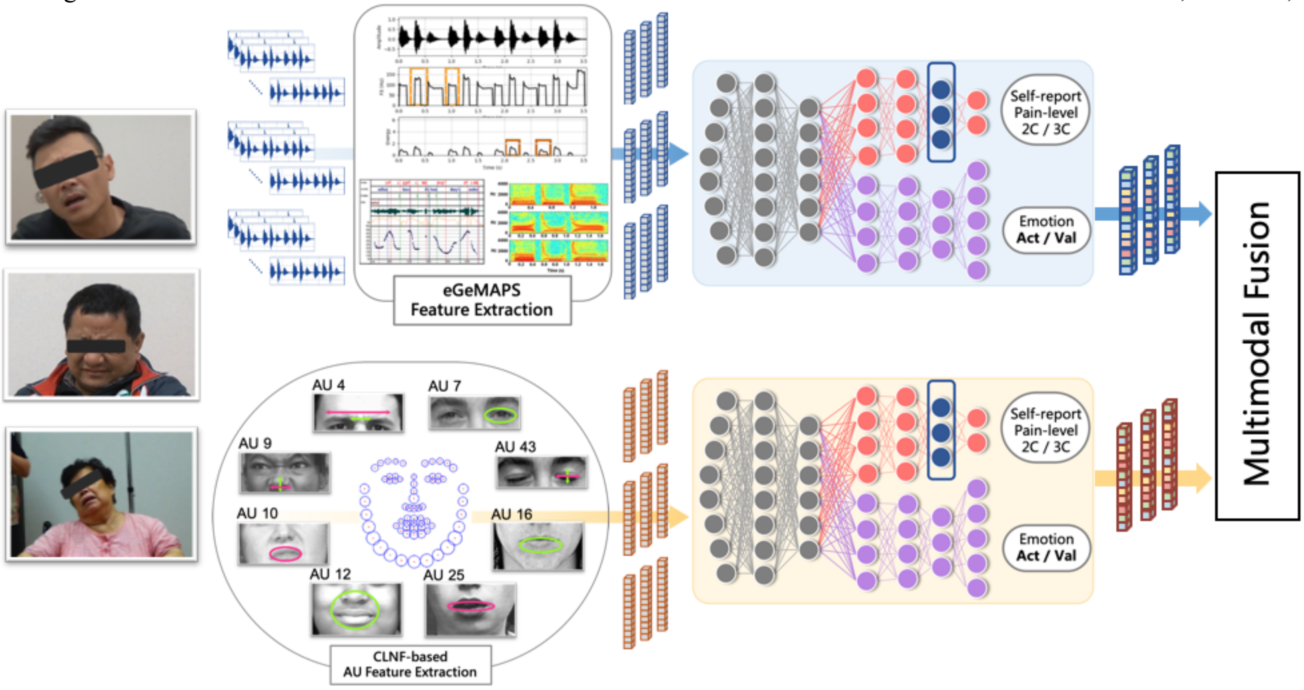

Fig. 1. *It shows the complete architecture of our emotion-enriched multitask network (EEMN) used for automatic pain-level recognition: acoustic and facial feature extraction, training multitask network with NRS pain-levels as main task and affect state as auxiliary task, performing multimodal fusion with embeddings from EEMN using support vector classifier.*

in classifying three-class classification (mild, moderate and severe). Further analysis reveals an interesting insight that in two of the five evaluators, their ratings on the patient's valence attribute are more in concordance with the patient's NRS score than their own observational assessment of patient's pain intensity. The rest of paper is organized as follows: Section 2 describes the database, behavior features and emotion-enriched multitask network. Section 3 includes experimental setups, analyses and results. Section 4 concludes with future work.

II. Research Methodology

A. Triage Pain-Level Multimodal Database

The Triage Pain-Level Multimodal Database was collected at the Emergency Department of Chang Gung Memorial Hospital during triage sessions for on-boarding patients [14]. The data includes audio-video recordings, vital sign, and other clinically-related outcomes. We include patients with symptoms of chest, abdominal, lower-back, limbic pain and headaches. Each triage session lasts approximately 30 seconds where the nurse asks each patient about the pain location, the NRS scale of pain intensity (0-10, where 10 means the worst pain ever), and a brief description on the type of pain felt. In this work, we use a total of 323 samples from 184 unique patients, i.e., the same setting as the most recent work on this corpus [16]. We categorize the NRS into three commonly-used pain-levels, i.e., mild (0-3), moderate (4-6), and severe (7-10).

1) Observer Annotations: In this work, in order to investigate the relationship between the observed emotional states and the self-reported pain-levels, we recruit five naive raters (2 males, 3 females) to annotate the emotional states and also the perceived pain-levels of these patients by viewing the recorded sample audio-visually. Since the patient's self-reported pain-levels were also recorded, we first mask the NRS answering portion of each video sample to ensure that the raters would not be influenced by the patient's own answer. Raters annotate the affective states, i.e., activation and valence using a 5-level discretized scale between -2 to 2, and also indicate their perception on the patient's pain-level as one of the four classes, i.e., no pain (0), mild (1-3), moderate (4-6), and severe (7-10).

TABLE I
Inter-evaluator agreement from five evaluators using entropy-based metric.

	Act	*Val*	*OV*
E1 ♂	0.0921	0.0256	0.0712
E2 ♂	0.0187	0.1007	0.0156
E3 ♀	0.0300	0.0260	**0.1867**
E4 ♀	0.0339	0.0404	0.0323
E5 ♀	0.1339	0.0298	0.0516

The inter-evaluator agreements obtained between the five annotators on activation (Act), valence (Val), and observed (OV) pain-levels are shown in Table I. In this work, we adopt the use of entropy-based metric proposed by Steidl *et al*. in measuring the inter-evaluator agreement [18]. The idea is to measure the amount of uncertainty when including the annotator as a new rater into the current rating pool. Specially, assume a probability distribution p with evaluations gathered from n existing raters. When a new rater is added, we estimate a new distribution $\bar{p}$. The inter-evaluator agreement of this rater with others is computed using the following equation:

$$S_{ent} = H(\bar{p}) - H(p) = -(\sum \bar{p} \cdot log\bar{p} + \sum p \cdot logp) \quad (1)$$

The lower the S_{ent} indicates a higher agreement level. In general, we notice that our evaluators tend to be in agreement with each other across the three different annotations. Only for the third annotator, when rating the observed pain-level, her rating shows a much larger disagreement with other annotators.

B. Emotion-Enriched Multitask Network (EEMN)

In this work, we further design a multitask network that learns an improved multimodal pain-affect representation to perform pain-level classification. Figure 1 depicts our Emotion-Enriched Multitask Network (EEMN) that models acoustic and facial expressions by using pain as the main task and emotion state as an auxiliary task. In the following section, we will detail the audio-video feature extraction, our proposed network and final pain classification.

1) Acoustic Features: We extract acoustic features on the patient's speaking portion for each recording sample using the eGeMAPS feature set [19]. It includes a set of 88 features covering a various statistical properties of spectral, cepstral, prosodic, and voice quality information extracted using the openSMILE toolkit [20].

2) Facial Low-level Descriptors: Facial action units (AU) have been known to be related to pain-level [21], and in this work, we design 25 facial action unit-inspired descriptors per frame characterizing eyes, mouth, eyebrows, and nose movement (details in [14]). These features are extracted based on the tracked 68 facial landmarks using method of constrained local neural fields (CLNF) [22]. We further compute 15 different statistical functionals on these extracted facial LLDs to generate a feature vector for each sample. The list of functionals includes maximum, minimum, mean, median, standard deviation, 1^{st} percentile, 99^{th} percentile, $99^{th} - 1^{st}$ percentile, skewness, kurtosis, minimum position, maximum position, lower quartile, upper quartile, interquartile range.

3) Affect-Pain Multitask Architecture: We use an emotion-enriched multitask network (EEMN) to learn a unified affect-pain feature representation for pain-level classification. EEMN architecture includes a typical hard-sharing layers at the beginning and branches out with task-specific layers toward the end [23]. In this work, the main task is the pain-level classification where the auxiliary task is the emotional state recognition. We train a separate speech-EEMN and face-EEMN. The embedding prior to the pain-level softmax layer can be seen as behavior representation that captures emotion-related information suitable for pain-level recognition.

The concatenation of vocal- and facial-embeddings derived from the modality-specific EEMN is used as the multimodal input to a linear-kernel support vector machine for final NRS pain-level recognition.

III. Experimental Setups and Results

In this work, we set up two different experiments:

- **Exp I**: a correlation analysis between observed emotional states and self-reported pain-level
- **Exp II**: NRS pain-level recognition tasks using EEMN

Exp I is designed to investigate the relationship between emotional states and varying pain severity. Exp II is setup to demonstrate the effectiveness of the EEMN to improve pain-level recognition by integrating emotion-related information.

TABLE II
Correlations of the observed emotional attributes with the self-reported pain-levels.

		Activation (Act)		***Valence (Val)***	
		$Sp.\ \rho$	p	$Sp.\ \rho$	p
2-Class	E1 ♂	0.3376	< 0.001	−0.3948	< 0.001
	E2 ♂	−0.1079	> 0.001	−0.4433	< 0.001
	E3 ♀	−0.1566	> 0.001	−0.5472	< 0.001
	E4 ♀	0.0192	> 0.001	−0.5336	< 0.001
	E5 ♀	0.0887	> 0.001	−0.4752	< 0.001
	Mean	0.0208	> 0.001	**−0.5136**	< 0.001
3-Class	E1 ♂	0.2847	< 0.001	−0.3312	< 0.001
	E2 ♂	−0.094	> 0.001	−0.3771	< 0.001
	E3 ♀	−0.1356	> 0.001	−0.4539	< 0.001
	E4 ♀	0.011	> 0.001	−0.4744	< 0.001
	E5 ♀	0.069	> 0.001	−0.3979	< 0.001
	Mean	0.012	> 0.001	**−0.4376**	< 0.001

A. Exp I: Emotional States Analyses

We present spearman correlations computed between each annotator's (E*) rated valence (Val) and activation (Act) values and the patient's self-reported pain-level (2-class: mild (0) and severe (1), or 3-class: mild (0), moderate (1), and severe (2)). We further take the mean of the five annotators rated emotion attributes to perform similar correlation computations.

1) Experimental Results: Table II summarizes the correlations of the observed emotional states with the self-reported pain-levels. Generally, we observe there is a significantly negative correlation between the valence state and the severity of the pain-level (mean: $\rho = -0.5136$ and $\rho = -0.4376$ for 2-class and 3-class respectively). In fact, we observe a consistent negative correlation between valence and pain intensity across all of the five annotators. This finding suggesting that a more negative emotion observed would correlate with a higher pain-level reported corroborates with past literatures [17], [24], [25]. This strong relationship does not hold for the activation. On average, activation is not correlated with the NRS score; the only significant correlation found is in the first evaluator, where the more activated the patient is, the more severe the pain experienced.

B. Exp II: NRS Pain-Level Recognition

In Exp II, we perform two different recognition setups: 1) binary classification between the extreme pain-levels, i.e., mild vs. severe pain, and 2) three-class classification, i.e., mild vs. moderate vs. severe. Our EEMN network architecture contains 2 fully-connected layers to be our hard sharing layers, and the number of task-specific layers for the main and auxiliary task are selected between 2 to 6. The output layer (24 nodes for speech and 32 nodes for face) is extracted as patients' feature representation with emotion-related information. All the batch size are specified as 300 and the learning rate is chosen between the range of 0.01 to 0.0001 using Adam optimizer.

1) Comparison Models: The following is the list of comparison models:

- **STN**: Model trained on the single main task of pain-level recognition without emotion attribute.

TABLE III
IT SUMMARIZES THE UNWEIGHTED AVERAGED RECALL (UAR) OBTAINED IN OUR PROPOSED PAIN-LEVEL RECOGNITION EXPERIMENT. STN INDICATES REPRESENTATION DERIVED FROM FEED-FORWARD NEURAL NETWORK WITHOUT USING EMOTIONAL ATTRIBUTES AS AUXILIARY TASK. EEMN IS OUR PROPOSED EMOTION-ENRICHED NETWORK, AND THE LETTERS AFTER EEMN INDICATES EITHER LEARNING WITH VALENCE (VAL) OR ACTIVATION (ACT) BY USING SPEECH (S) OR FACE (F) BEHAVIORS.

	Speech				***Face***			***Multimodal***				
	STN	VAE	EEMN Act	EEMN Val	STN	EEMN Act	EEMN Val	STN	EEMN Act(S)+Act(F)	EEMN Act(S)+Val(F)	EEMN Val(S)+Act(F)	EEMN Val(S)+Val(F)
2-Class												
Mild	65.2	66.0	70.4	65.2	71.3	76.5	73.0	64.3	75.7	75.7	69.5	67.0
Severe	60.8	61.0	51.0	66.7	51.0	52.9	57.8	62.7	52.9	56.9	70.6	62.7
UAR	63.0	63.5	60.7	65.9	61.1	64.7	65.4	63.5	64.3	66.3	**70.1**	64.9
3-Class												
Mild	42.6	43.7	33.0	42.6	52.2	64.3	44.3	37.4	55.6	52.1	53.9	51.3
Moderate	40.5	35.4	38.7	54.7	42.4	41.5	41.5	47.2	37.7	38.7	61.3	58.5
Severe	36.3	38.1	41.2	45.0	17.6	38.2	42.2	39.2	43.1	37.2	41.2	37.2
UAR	39.8	39.1	37.6	47.4	37.4	**48.0**	42.7	41.2	45.5	42.7	**52.1**	49.0

- **VAE** [16] : Model trained on encoding acoustic LLDs using variational learning with Maximum-Mean Discrepancy criterion
- **EEMN-X**: Model trained on main task of pain-level and auxiliary task of emotion X-attribute using a multitask network, where X can be either activation (Act) or valence (Val) attributes.

These feature representations are fed into our final classifier of linear-kernel support vector machine. The evaluation processes are done via leave-one-patient-out cross-validation. Unweighted average recall (UAR) is used as evaluation metric.

2) Experimental Results: Table III summarizes our self-reported pain-level classification results. Our proposed emotion-enriched multitask network achieves the best recognition rates of 70.1% and 52.1% in 2-class and 3-class classification in our multimodal fusion model of EEMN-Act (F) and EEMN-Val (S). It obtains a relative improvement of 6.6% and 10.9% over the baseline multimodal STN without integrating the information about individual affective states. This results also surpass the previous work on the same database using the variational autoencoder approach [16], specifically improves 6.6% and 13% in 2-class and 3-class recognition task.

Generally, we observe that integrating valence as an auxiliary emotion recognition task, it would improve the overall pain-level recognition rates for both speech and face modality. However, an interesting observation that we notice is that in the more challenging 3-class pain-level recognition, using activation as auxiliary task help the face modality more than using valence attribute (48.0% versus 42.7%). In fact our best model is a fusion between valence-enriched acoustic network with activation-enriched facial network, i.e., EEMN-Val (S) and EEMN-Act (F). This may due to the fact that valence attributes provides additional information in the acoustic features given that acoustic behaviors alone generally capture activation well already; activation on the other hand provides additional modeling capacity for facial modality, where facial expressions alone are known to capture valence dimensions reliably.

TABLE IV
IT SUMMARIZES THE UNWEIGHTED AVERAGED RECALL (UAR) OBTAINED FROM OBSERVED PAIN-LEVELS (OV), VALENCE RATINGS (VAL), AND ACTIVATION RATINGS (ACT) TO THE SELF-REPORTED NRS SCORE.

	OV	***Act***	***Val***
E1 ♂	47.2	44.2	44.5
E2 ♂	46.6	31.6	**47.0**
E3 ♀	63.3	33.9	49.9
E4 ♀	47.9	33.6	**52.0**
E5 ♀	53.6	33.6	47.6
Mean	51.7	35.4	48.2

3) Additional Analyses: With the strong relationship between the perceived valence annotation and the self-reported pain intensity, we further compute the concordance rate of each annotator's observed pain-levels (OV) to the patient's self-reported NRS pain-levels. By binning the OVs into the same three classes and emotion states also as three classes (i.e., class 0: (-2, -1), class 1: 0, class 2 (1, 2) for activation, and class 0: (1, 2), class 1: 0, class 2 (-1, -2) for valence), we can use UAR to measure the concordance rate of each observed annotation to the NRS pain-levels.

Table IV summarizes the results. The baseline chance concordance rate is 33%. We notice that the observed pain-levels (OV) while is generally most related to the NRS (mean UAR is 51.7%) among the three attributes, its UAR is actually worse than our best performing recognition model (fusion of EEMN-Val (S) and EEMN-Act (F) obtains UAR of 52.1%). It implies the possibility that by using computational methods may indeed be able to model subject internal feelings, such as pain, surpass human's observation. We also observe an interesting fact that for E2 and E4, their valence attributes reveal more about the patient's NRS than their rated observed pain-levels (E2: 47.0% vs. 46.6%, E4: 52.0% vs. 47.9%). This results seem to indicate an intriguing annotator's perceptual mechanism in rating pain versus emotion though these two are supposedly distinct internal constructs.

IV. Conclusions

In this work, we present a computational investigation into understanding the relationship between the self-reported pain versus the observed emotion state. Our analyses indicate that perceived valence is correlated negatively to the self-reported pain (NRS) intensity. We further propose a multi-task learning framework (EEMN) to improve automated NRS recognition by using observed emotion states as an auxiliary task. We also observe that the use of automated method can surpass pain-level ratings done by humans observations. Finally, our analysis indicates a possibility that by asking observer to evaluate valence ratings potentially carry more information about the patient's internal pain sensation than directly asking to rate the pain. This may further help in mitigating the inconsistency in the administration of NRS during triage, where the triage nurse would often rate how painful the patient is through observation when it becomes difficult to solicit answers from the patients.

To our knowledge, this is one of the first study into the relationship between pain and emotion states in a large real patients audio-video recordings. Aside from continuing to advance our technical framework in improving the recognition rates of NRS, we would also investigate scientifically the underlying mechanism of one's painful sensation and emotion states in modulating our multimodal behavior expressions.

References

[1] R. A. Bryant, "Memory for pain and affect in chronic pain patients," *Pain*, vol. 54, no. 3, pp. 347–351, 1993.

[2] J. Strong, R. Ashton, and D. Chant, "The measurement of attitudes towards and beliefs about pain," *Pain*, vol. 48, no. 2, pp. 227–236, 1992.

[3] L. D. Wandner, C. D. Scipio, A. T. Hirsh, C. A. Torres, and M. E. Robinson, "The perception of pain in others: how gender, race, and age influence pain expectations," *The Journal of Pain*, vol. 13, no. 3, pp. 220–227, 2012.

[4] L. E. Carter, D. W. McNeil, K. E. Vowles, J. T. Sorrell, C. L. Turk, B. J. Ries, and D. R. Hopko, "Effects of emotion on pain reports, tolerance and physiology," *Pain Research and Management*, vol. 7, no. 1, pp. 21–30, 2002.

[5] M. C. Bushnell, M. Čeko, and L. A. Low, "Cognitive and emotional control of pain and its disruption in chronic pain," *Nature Reviews Neuroscience*, vol. 14, no. 7, p. 502, 2013.

[6] S. S. Ngu, M. P. Tan, P. Subramanian, R. A. Rahman, S. Kamaruzzaman, A.-V. Chin, K. M. Tan, and P. J. Poi, "Pain assessment using self-reported, nurse-reported, and observational pain assessment tools among older individuals with cognitive impairment," *Pain management nursing*, vol. 16, no. 4, pp. 595–601, 2015.

[7] J. Younger, R. McCue, and S. Mackey, "Pain outcomes: a brief review of instruments and techniques," *Current pain and headache reports*, vol. 13, no. 1, pp. 39–43, 2009.

[8] S. B. McMahon, M. Koltzenburg, I. Tracey, and D. Turk, *Wall & Melzack's Textbook of Pain: Expert Consult-Online and Print*. Elsevier Health Sciences, 2013.

[9] E. Castarlenas, R. de la Vega, M. P. Jensen, and J. Miró, "Self-report measures of hand pain intensity: current evidence and recommendations," *Hand clinics*, vol. 32, no. 1, pp. 11–19, 2016.

[10] M. J. Sullivan, S. R. Bishop, and J. Pivik, "The pain catastrophizing scale: development and validation." *Psychological assessment*, vol. 7, no. 4, p. 524, 1995.

[11] P. Rodriguez, G. Cucurull, J. Gonzàlez, J. M. Gonfaus, K. Nasrollahi, T. B. Moeslund, and F. X. Roca, "Deep pain: Exploiting long short-term memory networks for facial expression classification," *IEEE transactions on cybernetics*, no. 99, pp. 1–11, 2017.

[12] M. Tavakolian and A. Hadid, "Deep binary representation of facial expressions: A novel framework for automatic pain intensity recognition," in *2018 25th IEEE International Conference on Image Processing (ICIP)*. IEEE, 2018, pp. 1952–1956.

[13] J. Egede, M. Valstar, and B. 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)*. IEEE, 2017, pp. 689–696.

[14] F.-S. Tsai, Y.-L. Hsu, W.-C. Chen, Y.-M. Weng, C.-J. Ng, and C.-C. Lee, "Toward development and evaluation of pain level-rating scale for emergency triage based on vocal characteristics and facial expressions." in *INTERSPEECH*, 2016, pp. 92–96.

[15] F.-S. Tsai, Y.-M. Weng, C.-J. Ng, and C.-C. Lee, "Embedding stacked bottleneck vocal features in a lstm architecture for automatic pain level classification during emergency triage," in *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 2017, pp. 313–318.

[16] J.-L. Li, Y.-M. Weng, C.-J. Ng, and C.-C. Lee, "Learning conditional acoustic latent representation with gender and age attributes for automatic pain level recognition," *Proc. Interspeech 2018*, pp. 3438–3442, 2018.

[17] P. Rainville, Q. V. H. Bao, and P. Chrétien, "Pain-related emotions modulate experimental pain perception and autonomic responses," *Pain*, vol. 118, no. 3, pp. 306–318, 2005.

[18] S. Steidl, M. Levit, A. Batliner, E. Noth, and H. Niemann, "" of all things the measure is man" automatic classification of emotions and inter-labeler consistency [speech-based emotion recognition]," in *Proceedings.(ICASSP'05). IEEE International Conference on Acoustics, Speech, and Signal Processing, 2005.*, vol. 1. IEEE, 2005, pp. I–317.

[19] F. Eyben, K. R. Scherer, B. W. Schuller, J. Sundberg, E. André, C. Busso, L. Y. Devillers, J. Epps, P. Laukka, S. S. Narayanan *et al.*, "The geneva minimalistic acoustic parameter set (gemaps) for voice research and affective computing," *IEEE Transactions on Affective Computing*, vol. 7, no. 2, pp. 190–202, 2016.

[20] F. Eyben, M. Wöllmer, and B. Schuller, "Opensmile: the munich versatile and fast open-source audio feature extractor," in *Proceedings of the 18th ACM international conference on Multimedia*. ACM, 2010, pp. 1459–1462.

[21] P. Lucey, J. F. Cohn, I. Matthews, S. Lucey, S. Sridharan, J. Howlett, and K. M. Prkachin, "Automatically detecting pain in video through facial action units," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 41, no. 3, pp. 664–674, 2011.

[22] T. Baltrusaitis, P. Robinson, and L.-P. Morency, "Constrained local neural fields for robust facial landmark detection in the wild," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2013, pp. 354–361.

[23] N. Jaques, S. Taylor, E. Nosakhare, A. Sano, and R. Picard, "Multi-task learning for predicting health, stress, and happiness," in *NIPS Workshop on Machine Learning for Healthcare*, 2016.

[24] K. Wiech and I. Tracey, "The influence of negative emotions on pain: behavioral effects and neural mechanisms," *Neuroimage*, vol. 47, no. 3, pp. 987–994, 2009.

[25] J. L. Rhudy and M. W. Meagher, "Fear and anxiety: divergent effects on human pain thresholds," *Pain*, vol. 84, no. 1, pp. 65–75, 2000.