

PainFormer: a Vision Foundation Model for Automatic Pain Assessment

Stefanos Gkikas , Raul Fernandez Rojas , and Manolis Tsiknakis 

Abstract—Pain is a manifold condition that impacts a significant percentage of the population. Accurate and reliable pain evaluation for the people suffering is crucial to developing effective and advanced pain management protocols. Automatic pain assessment systems provide continuous monitoring and support decision-making processes, ultimately aiming to alleviate distress and prevent functionality decline. This study introduces *PainFormer*, a vision foundation model based on multi-task learning principles trained simultaneously on 14 tasks/datasets with a total of 10.9 million samples. Functioning as an embedding extractor for various input modalities, the foundation model provides feature representations to the *Embedding-Mixer*, a transformer-based module that performs the final pain assessment. Extensive experiments employing behavioral modalities—including RGB, synthetic thermal, and estimated depth videos—and physiological modalities such as ECG, EMG, GSR, and fNIRS revealed that *PainFormer* effectively extracts high-quality embeddings from diverse input modalities. The proposed framework is evaluated on two pain datasets, *BioVid* and *AI4Pain*, and directly compared to 75 different methodologies documented in the literature. Experiments conducted in unimodal and multimodal settings demonstrate state-of-the-art performances across modalities and pave the way toward general-purpose models for automatic pain assessment. *The foundation model's architecture (code) and weights are available at <https://github.com/GkikasStefanos/PainFormer>.*

Index Terms—Pain recognition, deep learning, transformers, multi-task learning, multimodal, synthetic data, data fusion.

I. INTRODUCTION

PAIN is a fundamental evolutionary function, signaling potential harm to the organism or indicating the initiation of illness. It is an essential component of the body's defense mechanism to maintain its integrity [1]. In addition, the biopsychosocial model of pain suggests that physical symptoms of pain are the culmination of a dynamic interplay among biological, psychological, and social factors [2], which led Williams and Craig [3] to propose that pain is “*a distressing experience associated with actual or potential tissue damage with sensory, emotional, cognitive, and social components*”. Pain can be classified into three types: nociceptive (resulting from tissue damage), neuropathic (originating from nerve damage), or nociplastic (due to a sensitized nervous system). Each type influences diagnostic and treatment choices at various stages [2]. From a time-duration perspective, the two main

categories are acute and chronic, the former persisting or repeating for over three months [4]. Acute pain typically results from visible physiological damage due to injury, surgery, illness, trauma, or painful medical procedures. It generally resolves when the underlying cause is treated or healed; however, if unresolved, it can evolve into a chronic condition beyond the initial acute phase. Postoperative pain, which is a type of acute pain, occurs specifically after surgical interventions and represents a significant concern for both patients and healthcare providers, highlighting the importance of effective pain management strategies to aid recovery and prevent the transition to chronic pain [5]. Chronic pain manifests in various forms related to the temporal dimension, such as chronic-recurrent (*e.g.*, migraine headaches) or chronic-continuous (*e.g.*, low back pain) [6].

Pain is a pervasive medical issue worldwide, impacting up to 30% of the general adult population [2] and between 83% and 93% of senior adults living in residential care facilities [7]. According to the Global Burden of Disease (GBD) study, pain is the leading cause of years lived with disability (YLD) [8], with three of the top contributors being chronic pain conditions: back pain, musculoskeletal disorders, and neck pain [9]. The influence of pain expands beyond individuals to affect society, presenting clinical, economic, and social challenges. The estimated economic and healthcare costs associated with pain due to reduced work productivity in the United States range from \$560 to \$635 billion, exceeding the costs of heart disease, cancer, and diabetes combined [10]. Europe shows similar patterns, with direct healthcare expenses and indirect socioeconomic costs of chronic pain accounting for 3% to 10% of the gross domestic product [11]. Additionally, in Australia, the average annual cost for one of the 15.4% of residents living with chronic pain is between AU\$22,588 and AU\$42,979 when including non-financial costs [12]. Pain extends beyond direct impacts on a patient's life, leading to a range of adverse effects, including opioid use, drug overuse, addiction, deteriorating social relationships, and psychological disorders [13]. Over the past two decades, the use of prescription opioids has significantly increased in the United States, where the rate of overdose deaths has more than quadrupled from 1999 to 2016 [14]. Additionally, the side effects of these opioids, including lethargy, depression, anxiety, and nausea, significantly affect workforce productivity and overall quality of life [15].

Accurate pain assessment is vital for early diagnosis, disease progression monitoring, and treatment effectiveness evaluation, especially in managing chronic pain [16]. This importance has led to pain being designated as “*the fifth vital sign*” in nursing literature [17]. Furthermore, pain evaluation is essential

Stefanos Gkikas and Manolis Tsiknakis are with the Hellenic Mediterranean University, Department of Electrical and Computer Engineering, Heraklion, Crete 714 10, Greece and with the Institute of Computer Science, Foundation for Research & Technology-Hellas, Heraklion, Crete GR-70013 Greece (email: gkikas@ics.forth.gr; tsiknaki@ics.forth.gr).

Raul Fernandez Rojas is with the University of Canberra, Faculty of Science and Technology, Canberra, ACT 2617, Australia (email: raul.fernandezrojas@canberra.edu.au)

in physiotherapy, where the therapist externally induces stimuli, and understanding the patient's pain level is needed [18]. Objectively evaluating pain is crucial for delivering appropriate care, particularly for vulnerable groups who cannot directly or reliably communicate their pain condition, such as infants, young children, those with mental health conditions, and elderly individuals. Various methodologies are employed to assess pain, such as self-reporting, where individuals describe their pain experiences, which is currently regarded as the gold standard [19]. Pain evaluation methods in clinical settings offer quantifiable indicators of pain, ranging from the Numeric Pain Rating Scale (NPRS) and Visual Analogue Scale (VAS) to quantitative sensory testing methods such as the pressure pain detection threshold (PPDT) [20]. Additionally, behavioral indicators such as facial expressions (like grimacing, open mouth, or lifted eyebrows), vocalizations (such as crying, moaning, or screaming), and body and head movements are important markers [21]. Physiological measurements, including electrocardiography (ECG), electromyography (EMG), galvanic skin responses (GSR), and respiration rate, also provide essential understandings of the physiological manifestations of pain [16], [22]. Furthermore, brain monitoring techniques such as near-infrared spectroscopy (fNIRS) have shown the capability to detect hemodynamic activity changes related to pain stimuli [23], [24].

Over the past two decades, computational science researchers have dedicated their efforts to developing models and algorithms for advancing automatic pain recognition systems [25]. The aim is to accurately identify the presence and intensity of pain by analyzing physiological and behavioral responses. The emergence of deep learning and artificial intelligence (AI) methods has further explored these automatic approaches, aiming at interpreting the complex, multifaceted nature of pain [16]. Multiple studies highlight the potential of automated systems utilizing behavioral or physiological pain assessment modalities [26]. Sario *et al.* [27] demonstrate the viability of accurately recognizing pain via facial expressions, illustrating their value in clinical settings. Multimodal sensing appears especially promising, exhibiting increased accuracy in pain monitoring systems [21]. Additionally, incorporating the temporal dimension of the modalities has been demonstrated to enhance the effectiveness of pain assessment [16]. In recent years, affective computing research has increasingly adopted thermal imaging techniques [28], motivated by literature indicating that stress and cognitive load have notable effects on skin temperature [29]. These effects arise from the autonomic nervous system's (ANS) control over physiological signals such as heart rate, respiration rate, blood perfusion, and body temperature, all of which reflect human emotions and affects [28]. Moreover, muscle contractions can affect facial temperature by transferring heat to the facial skin [30]. Thus, thermal imaging has emerged as a promising technique for capturing transient facial temperature [31]. However, only a few studies have explored thermal imaging in pain research. In [32], researchers reported that facial temperature increases after exposure to a painful stimulus, suggesting that thermal cameras might be effective tools for monitoring pain. In [33], the authors demonstrated that thermal videos achieved similar

accuracy to RGB videos in recognizing pain in an automatic pain assessment environment. Moreover, our previous study [34] introduced synthetic thermal videos through a deep-learning generative process and evaluated their effectiveness in recognizing pain. The findings showcased that the performances of the synthetic modality are equivalent to the original RGB videos, while the combination of them holds significant potential. Furthermore, depth cameras are employed in emotion recognition research because they enable the extraction and use of features related to head pose and 3D facial landmark analysis [35]. These camera sensors capture body and head dynamics, facilitating the analysis of pixel and depth intensity changes associated with movements and affects [36], [37]. In addition, depth sensors are more invariant and maintain their effectiveness under poor or uneven lighting conditions [38]. The research in [39] demonstrated that depth information combined with RGB enhanced the accuracy of micro-expression analysis and, by extension, emotion recognition tasks, while in [40], the authors utilized the 3D positions of facial feature points to identify basic emotions. In pain research, the adoption of depth modality is limited. The authors in [41] utilized depth map videos to estimate head movements, which enhanced pain recognition performance when combined with facial expressions.

With a new emerging paradigm for building AI systems based on foundation models, there has been a shift towards more adaptable and scalable systems capable of generalizing across various tasks and domains. A foundation model refers to any model trained on extensive datasets, typically through self-supervision at scale, which can thereafter be adapted—for instance, fine-tuned—to a diverse array of downstream tasks. Although foundation models rely on conventional deep learning and transfer learning techniques, their large scale leads to the emergence of new capabilities and improved effectiveness across numerous tasks [42]. A plethora of examples have appeared recently in the literature. For instance, SAM [43] is a foundation model for image segmentation initially trained from scratch on 11 million images. In subsequent studies [44], [45], researchers adapted the SAM model for medical imaging by optimizing it for smaller, specialized datasets. Furthermore, another significant paradigm shift has occurred with the introduction of generalist models [46]—a new type of foundation model trained simultaneously on various tasks under a unified learning policy (usually with supervision). This approach particularly benefits computer vision, where embedding representations can differ significantly across tasks and various vision modalities [47]. In the field of automatic pain assessment, there are approaches that employ pre-trained models. However, they follow the traditional method of pre-training on a particular larger general dataset and fine-tuned for the specific task of pain assessment. Studies like those detailed in [33], [48] are founded on transfer learning techniques from facial recognition datasets, whereas others, such as [34], [49], employ multi-stage pre-training methods that progressively learn facial and emotional feature representations.

The present study introduces a multi-task learning-based vision foundation model, *PainFormer*, for automatic pain assessment. To the best of our knowledge, this represents the first

effort in pain research to develop and implement a foundation system for pain recognition. Our method is motivated by the principles of [46], training simultaneously across various tasks and datasets and founded on the core concept of foundation approaches where representation learning is performed on large-scale corpora and reused to the downstream tasks, in our case, pain assessment. The proposed framework consists of *PainFormer*, the primary model, *Embedding-Mixer*, and *Video-Encoder*. These modules are described in detail in Section III. The key contributions of this research are threefold: (1) the introduction of a foundation model capable of extracting high-quality embeddings regardless of the input modality, (2) the utilization of synthetic thermal videos and estimated depth videos, and (3) an extensive evaluation of diverse behavioral and physiological modalities in unimodal and multimodal settings aiming for effective pain assessment.

II. RELATED WORK

Extensive research has been conducted in the field of automatic pain assessment to explore effective methods for evaluating pain-related conditions. Unimodal methods exploring specific sensory channels, such as vision or contact sensors, have demonstrated high effectiveness while maintaining a degree of simplicity. In addition, numerous multimodal approaches have been developed, aiming to integrate diverse information streams. These usually combine behavioral and physiological data or various physiological data sources alone. While often enhancing performance, these methods introduce the challenge of managing the complexity associated with data fusion.

Various compelling methods have been introduced, utilizing video modalities, ranging from hand-crafted feature engineering for facial expressions analysis to employing raw videos with sophisticated deep learning architectures. Werner *et al.* [41] analyzed facial expressions by extracting features based on point distances and head poses, which were estimated using depth information. By integrating these into a unified feature vector and applying a random forest classifier, they achieved a 76.60% accuracy rate in the binary pain classification task. Patania *et al.* [50] calculated fiducial points and applied graph neural networks (GNN) to achieve 73.20% accuracy, while the authors in [51] employed vision transformers on raw videos to reach an accuracy of 77.10%. Other researchers have made efforts to compute biosignal-related information directly from videos, aiming to enhance assessment accuracy while minimizing the reliance on contact sensors. Yang *et al.* [52] leveraged raw videos and three-dimensional convolutional neural networks (3D CNNs) to extract facial features and estimate remote photoplethysmography, effectively combining behavioral and physiological cues for enhanced performance. Similarly, the authors in [53] used 3D CNNs to extract pseudo-heart rate from videos, achieving over a 10% increase in accuracy by integrating facial features with heart rate data compared to using videos in isolation.

Multiple research efforts have been reporting exploring the use of biosignals as standalone modalities, recognizing their reliability and immunity to the limited or exaggerated

expressivity issues presented in video-based methods. Thiam *et al.* [54] developed 1D CNNs to analyze and classify raw ECG signals, achieving modest performance, whereas in [41] leveraged domain-specific electrocardiography features—such as the mean of consecutive heartbeat intervals and the square mean root of successive differences—to attain an enhanced performance of 64%. Interestingly, Huang *et al.* [53] achieved a 65% accuracy in pain detection using pseudo heart rate extracted from videos, whereas in [49], authors utilized solely heart rate from ECGs combined with a transformer-based model to achieve an accuracy of 67.04% for the same task. Employing EMG signals and computing a series of time and frequency domain features, the authors in [55] achieved modest performances from this challenging modality. In contrast, Kachele *et al.* [56], utilizing GSR signals—recognized as the most informative modality for pain assessment—extracted statistical features such as skewness, kurtosis, and the temporal slope of the signal. They achieved an 81.90% accuracy in the binary pain classification by applying a random forest classifier.

Considering the multidimensional nature of pain, integrating various modalities within a multimodal system presents a promising yet challenging approach. Combining different information sources can significantly improve the accuracy and sensitivity of pain assessment. While individual modalities yield satisfactory predictive performance, their effective combination can lead to enhanced outcomes [26], albeit with notable complexities in achieving optimal integration. Additionally, leveraging cues from multiple channels could be beneficial and crucial, especially in clinical environments where access to a particular modality might be compromised—for instance, when a patient's rotation obstructs facial visibility. Thiam *et al.* [54] exploited raw ECG, EMG, and GSR signals, forming a 2D representation by concatenating them across the temporal axis and employing a 2D CNN for analysis, which yielded an accuracy of 84.40%. In contrast, study reported in [57] achieved an accuracy of 85.70% using the same modalities but focusing on extracting hand-crafted features, thereby underscoring the significance and value of domain-specific engineering. Zhi *et al.* [58] extracted facial descriptors from videos and a plethora of features from ECG, EMG, and GSR, achieving a notable accuracy of 86% in distinguishing no pain from high pain levels. In [49], the authors enhanced their results by over 5% by integrating raw facial videos with heart rate data, compared to using the video modality individually. Similarly, the research in [59] utilized facial landmarks and a comprehensive set of 131 features from ECG and GSR, including time domain, frequency, and entropy-based metrics, to achieve an improvement of more than 6% over the unimodal video approach.

III. METHODOLOGY

This section describes the architecture of the proposed framework's components. It details the foundation model's multi-task learning-based pre-training, augmentation techniques, and training settings for pre-training and pain assessment tasks. Additionally, it explains the generation of synthetic thermal videos, depth estimation for creating depth videos, and visualization of biosignal modalities used in this study.




Fig. 1: Representation of the main models and their components presented in this study: **(a)** *PainFormer* is organized hierarchically into four stages, integrating *Spectral* and *Self-Attention Layers* to function as the embedding extractor for the inputs; **(b)** The *Spectral Layer*, a primary component of *PainFormer*, applies FFT to compute frequency-related information combined with a learnable filter K to emphasize important frequencies; **(c)** The *Self-Attention Layer*, another primary module of *PainFormer*, facilitates parallel computation of features and their relationship; **(d)** The *Embedding-Mixer*, which combines cross and self-attention mechanisms, serves as the module for final classification of the embeddings used in the pain assessment task; **(e)** The *Video-Encoder*, a compact and efficient module, encodes video representations into a lower dimensional space; **(f)** The *MLP-1* is included in the *Spectral Layer*; **(g)** The *MLP-2*, part of the *Self-Attention Layer*; **(h)** The *MLP-3* structure is incorporated within the *Embedding-Mixer* and *Video-Encoder*.

TABLE I: Number of parameters and FLOPS for the modules.

Module	Params (Millions)	FLOPS (Giga)
<i>PainFormer</i>	19.60	5.82
<i>Embedding-Mixer</i>	9.85	2.94
<i>Video-Encoder</i>	3.37	0.86
Total	32.82	9.62

A. Framework Architecture

The proposed framework comprises three models: the *PainFormer*, the foundation model that functions as an embedding extractor for inputs; the *Embedding-Mixer*, which utilizes these embeddings, whether individually or combined, for the pain classification task; and the *Video-Encoder*, encoding the video presentations into a lower-dimensional latent space, explicitly employed for the multimodal approach detailed subsequently. It should be noted that the framework works in two distinct phases rather than in an end-to-end manner: initially, it extracts embeddings, and subsequently, it utilizes them according to the specific modalities' pipelines and needs. Table I presents the number of parameters and the computational cost of floating-point operations (FLOPS) for the modules and are described in detail in what follows.

1) *PainFormer*: Vision Transformers (ViT) have been successfully implemented across various image analysis tasks,

showcasing the effectiveness of their core self-attention mechanism [60]. Furthermore, recent developments in Vision Multilayer Perceptron (Vision-MLP) models that employ spectral mixing techniques—replacing self-attention layers with Fourier transformation layers—demonstrate that simpler architectures with fewer inductive biases can yield comparable outcomes [61]. Our approach incorporates two principal concepts: the hierarchical Vision Transformers (ViT) [62], which utilize multiple stages of embedding extraction to enhance performance and scalability, and the Fourier transform, which, as demonstrated in [63] effectively mixes information from various tokens. *PainFormer* combines spectral layers, which are implemented using the Fast Fourier Transform (FFT) alongside self-attention layers. Specifically, in the early stages of the model, both spectral and self-attention layers are applied. In contrast, the latter stages exclusively utilize self-attention. Fig. 1(a) illustrates the architecture of the *PainFormer*.

Before discussing the specific components, we provide preliminary information to outline the general concept of the *PainFormer* architecture. Every 2D input image I is divided into n non-overlapping patches, with each patch $\in \mathbb{R}^{n \times h \times w \times 3}$, where h and w define the resolution of each patch and are equal to 16×16 , and the 3 indicates the number of color channels. Each patch is linearly projected into a token with dimension $d = 768$, followed by a positional encoding layer. Applying Discrete Fourier Transform (DFT), to a 1D input sequence of N elements, $x[n]$, where n ranges from 0 to $N - 1$, results

in the following expression:

$$X[k] = \sum_{n=0}^{N-1} x[n] e^{-i2\pi \frac{k}{N} n} = \sum_{n=0}^{N-1} x[n] W_N^{kn}, \quad (1)$$

where i denotes the imaginary unit, and W_N denotes the twiddle factor, $W_N \triangleq e^{-i2\pi/N}$. In this manner, the sequence is transformed into the frequency domain. Conversely, the original input sequence can be reconstructed by applying the inverse Discrete Fourier Transform (IDFT):

$$x[n] = \frac{1}{N} \sum_{k=0}^{N-1} X[k] \cdot e^{i2\pi \frac{k}{N} n}, \quad (2)$$

where $x[n]$ is the initial time-domain sequence. Furthermore, Eq. (1) can be adapted for two-dimensional inputs, $x[m, n]$, with $0 \leq m \leq M-1$ and $0 \leq n \leq N-1$, characterized by:

$$X[u, v] = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} x[m, n] \cdot e^{-i2\pi \left(\frac{um}{M} + \frac{vn}{N} \right)}, \quad (3)$$

where $X[u, v]$ represents the two-dimensional frequency representation of the spatial-domain input $x[m, n]$. *PainFormer* performs the aforementioned processes using specific modules, detailed as follows:

a) Spectral Layer: For the tokens x from image I , a 2D FFT is applied across the spatial dimensions to transform x into the frequency domain:

$$X = \mathcal{F}[x] \in \mathbb{C}^{h \times w \times d}. \quad (4)$$

After applying the FFT to extract the various frequency components of the image, we employ a learnable filter, $K \in \mathbb{C}^{h \times w \times d}$, which acts as a gate to regulate the significance of each frequency component. This modulation of the spectrum allows for identifying and learning features such as lines and edges. Specifically:

$$\tilde{X} = K \odot X, \quad (5)$$

where $\odot$ defines the element-wise multiplication. Afterward, the inverse Fast Fourier Transform (IFFT) is applied, which converts the spectral space back into the physical space:

$$x \leftarrow \mathcal{F}^{-1}[\tilde{X}], \quad (6)$$

where the physical space is referred to as the spatial domain in this case. The final component of a spectrum layer is an MLP module, which enables efficient channel mixing communication:

$$\Phi(x) = W_2 \cdot \text{GELU}(\text{DWConv}(W_1 \cdot x + b_1)) + b_2, \quad (7)$$

where GELU refers to the Gaussian Error Linear Unit activation function, and DWConv denotes a depthwise convolution layer. In addition, layer normalization is employed before and after the FFT and IFFT processes, refer to Fig. 1(b).

TABLE II: Details of the *PainFormer*'s architecture.

Stage	# Spectral Layers	# Self-Attention Layers	# Self-Attention Heads	Dimension d
1	2	1	2	64
2	2	2	4	128
3	–	12	10	320
4	–	3	16	160

d : token dimensions

b) Self-Attention Layer: The mechanism employed in this layer is the classic self-attention mechanism used in transformers, where for every token sequence X , the attention is defined as follows:

$$\text{Att}(X) := \text{softmax} \left(\frac{XW_q(XW_k)^T}{\sqrt{d}} \right) XW_v, \quad (8)$$

where $\text{Att} : \mathbb{R}^{N \times d} \rightarrow \mathbb{R}^{N \times d}$, and $N := hw$. Also, $W_q, W_k, W_v \in \mathbb{R}^{d \times d}$ are the query, key, and value matrices. A layer normalization before and after the attention mechanism similar to the spectrum layer are also applied in the self-attention layer. Furthermore, the MLP module in this layer described as:

$$\Phi(x) = W_2 \cdot \text{GELU}(W_1 \cdot x + b_1) + b_2. \quad (9)$$

The architecture of the layer is depicted in Fig. 1(c).

c) Stages: A stage-based architecture has been developed to produce a hierarchical representation. *PainFormer* consists of four stages. At the end of each one, a single-layer 2D CNN reduces the number of tokens by downsampling the resolution by a factor of 2. In addition, each stage employs a unique combination of spectral and self-attention layers, varying numbers of heads in the self-attention layers, and different dimensions for the extracted tokens. Table II provides the corresponding details.

2) Embedding-Mixer: This model is a transformer-based network that incorporates cross- and self-attention mechanisms. As in other studies [64], it introduces an asymmetry in attention computation through cross-attention, which involves fewer latent variables. This approach reduces computational complexity and enhances the model's efficiency. Cross-attention operates similarly to self-attention as in Eq. (8). However, the dimensions for W_q, W_k , and W_v are $n \times d$ instead of $d \times d$, where $n < d$ and here, n is set to 256. The *Embedding-Mixer* consists of 2 layers, each containing 1 cross-attention module and 2 self-attention modules. In addition, the number of heads for the cross and self-attention is 1 and 8, correspondingly. The output embedding has a length of 512 and is utilized for the final classification task, refer to Fig. 1(d).

3) Video-Encoder: The architecture of the specific module resembles the *Embedding-Mixer*; however, for efficiency, it comprises just 1 layer featuring a 1 cross-attention module with 1 head. The number of latent variables, n , is set to 256, while the output embedding length is 40. This module is employed only within the particular framework for one of the multimodal approaches presented, where video embeddings are integrated with GSR embeddings. The module is illustrated in Fig. 1(e). Note that the *Spectral Layer*, *Self-Attention Layer*,




Fig. 2: Examples of different vision modalities in frame samples: (a) RGB frame, (b) synthetic thermal frame, and (c) depth estimation frame.




Fig. 3: Examples of different visual representations for biosignals: (a) waveform, (b) spectrogram-angle, (c) spectrogram-phase, and (d) spectrogram-PSD.

Embedding-Mixer, and *Video-Encoder* all incorporate MLP layers, as depicted in Fig. 1(f, g, h).

B. Synthetic Thermal & Depth Videos

In this study, we integrate thermal and depth vision modalities, in addition to RGB videos, into the pain assessment pipelines. For the thermal modality we employ the thermal videos from our previous work [34], which introduced an image-to-image translation (I2I) approach using a conditional generative adversarial network (cGAN). The network was developed and trained to translate the data distribution from the RGB to the thermal domain, enabling the generation of synthetic thermal representations from new RGB videos. For the depth videos we are using the “*Depth Anything*” method [65], a foundational model for monocular depth estimation (MDE) that employs a vision transformer-based encoder-decoder architecture and semi-supervised learning. Fig. 2 presents a frame sample from the RGB, synthetic thermal, and depth modalities.

C. Biosignal Visualization

Since the foundation model presented in this study is vision-based, a 2D representation of physiological modalities needs to be employed – representing signals as images has also been successfully demonstrated in other affective-related studies [66]–[68]. Four different visualizations are explored: (1) *waveform* diagrams, which depict the shape and form of a signal as it progresses over time, representing amplitude, frequency, and phase; (2) *spectrogram-angle*, which illustrates the phase angles of the frequencies; (3) *spectrogram-phase*, which shows phase information with unwrapping to address discontinuities; and (4) *spectrogram-PSD*, which represents the power spectral density, exhibiting how power is distributed across frequencies over time. Fig. 3 presents an example of each of these four visualizations.

TABLE III: Datasets utilized for the multi-task learning-based pre-training process of the *PainFormer*.

Dataset	# samples	# classes	Modality
VGGFace2 [69]	3.31M	9,131	Facial Images
SpeakingFaces RGB [73]⊙	0.76M	142	Facial Images
SpeakingFaces Thermal [73]⊙	0.76M	142	Facial Images
DigiFace-1M [70]	0.72M	10,000	Facial Images
DigiFace-1M [70]	0.50M	100,000	Facial Images
AffectNet [71]	0.40M	8	Facial Images
SFace [74]	1.84M	10,341	Facial Images
CACIA-WebFace [75]	0.50M	10,575	Facial Images
RAF-DB basic [72]	15,000	7	Facial Images
RAF-DB compound [72]	4,000	11	Facial Images
Compound FEE-DB [76]	6,000	26	Facial Images
EEG-BST-SZ [77]⊙	1.50M	2	EEG signals
Silent-EMG [78]⊙	0.19M	8	EMG signals
ECG HBC Dataset [79]⊙	0.45M	5	ECG signals
Total: 14 datasets–tasks	10.9M		

EEG: electroencephalogram EMG: electromyography ECG ⊙: The datasets were also used for the I2I process described in III-B, in addition to the training of the *PainFormer* ⊙: The samples were transformed into spectrograms, equally divided into three parts for each spectrogram type, before being employed.

D. Foundation Training

PainFormer, the proposed foundation model, functions as an embedding extractor, as previously described. To accomplish this, it has been trained extensively across 14 datasets comprising 10.9 million samples; refer to Table III for details. The training data encompass a range of human-related datasets, including facial recognition datasets like *VGGFace2* [69] and *DigiFace-1M* [70] and basic and compound emotion recognition datasets such as *AffectNet* [71] and *RAF-DB* [72], respectively. Additionally, biosignal datasets, EEG, EMG, and ECG-based, have also been utilized. Regarding the training procedure, *PainFormer* has been trained using a multi-task learning approach, where each dataset corresponds to a distinct supervised task. From an architectural perspective, the model remains consistent with the definition in III-A1. However, it now includes auxiliary classifiers for task-specific purposes. Each classifier is a compact, single-layer, fully connected network with an ELU (Exponential Linear Unit) activation function. The training objective is to learn from all 14 datasets/tasks simultaneously. The process is described as follows:

$$L_{total} = \sum_{i=1}^{14} [e^{w_i} L_{S_i} + w_i], \quad (10)$$

where L_{S_i} represents the loss associated with each specific dataset/task, and w_i denotes the learned weights that enable the learning process by minimizing the aggregate loss L_{total} , which encompasses all the individual losses. The foundation model was trained in this manner for 200 epochs.

E. Augmentation & Regularization Methods

Several augmentation and regularization methods were employed in the *PainFormer*’s pre-training and for the downstream task of pain assessment. *TrivialAugment* [80] and *AugMix* [81]

TABLE IV: Training details of the proposed framework.

Task	Optimizer	LR	LR decay	Weight decay	Warmup epochs	Cooldown epochs	Batch size
MTL	<i>AdamW</i>	<i>2e-5</i>	<i>cosine</i>	0.1	5	10	126 [Ⓢ]
Pain	<i>AdamW</i>	<i>2e-5</i>	<i>cosine</i>	0.1	10	10	32

MTL: multi-task learning for pre-training the foundation model Pain: pain assessment task LR: learning rate [Ⓢ]: batch size is proportionally distributed across the 14 tasks

were employed for the foundation training. Additionally, a custom augmentation technique involving adjustments to brightness, contrast, saturation, and image cropping was implemented. The pre-training process also incorporated random noise from a Gaussian distribution. Furthermore, a method was developed to mask out random square sections of input images. *DropPath* [82] and *Label Smoothing* [83] were employed to regularize during pre-training. Two augmentation techniques have been integrated within the framework of the pain assessment task. The first, called *Basic*, combines polarity inversion with noise addition. This approach transforms the original input embeddings by reversing the polarity of data elements and introducing random noise sourced from a Gaussian distribution, thereby creating variability and perturbations. The second technique, *Masking*, applies zero-valued masks to the embeddings, nullifying segments of the vectors. The size and placement of these masks are randomly determined, covering 10% to 20% of the embedding's total dimensions. For regularization, the *DropOut* [84] and *Label Smoothing* [83] techniques were utilized. Table IV provides additional details for the two training procedures.

F. Dataset Details

To assess the performance and robustness of our proposed framework, we conducted experiments on two distinct pain datasets, namely *BioVid* [85] and *AI4Pain* [86]. These datasets provide a diverse and robust basis for assessing the effectiveness of our model in pain assessment.

1) *BioVid Heat Pain Database*: The particular dataset is widely known and established in the field of pain research. It includes facial videos, electrocardiograms, electromyograms, and galvanic skin response measurements from eighty-seven ($n=87$) healthy individuals (44 males and 43 females, aged between 20 and 65). The experimental setup involved using a thermode to induce pain in the participants' right arm. Before data collection, each participant's pain and tolerance thresholds were determined. These thresholds marked the minimum and maximum levels of pain, with two additional intermediate levels, leading to a total of five distinct pain intensities: No Pain (NP), Mild Pain (P_1), Moderate Pain (P_2), Severe Pain (P_3), and Very Severe Pain (P_4). The temperature settings used for inducing pain varied within the range from P_1 to P_4 but never exceeded 50.5°C . Each participant experienced 20 pain inductions at each of the four predetermined levels of intensity (P_1 to P_4). Each stimulus lasted 4s, followed by a recovery period of 8 to 12s. Additionally, 20 baseline measurements were taken at 32°C (NP), resulting in 100 stimulations per

participant, delivered randomly. The data was then processed to create segments of 5.5s starting 1s after reaching the target temperature for each stimulation. This processing yielded 8,700 samples, each lasting 5.5s, distributed evenly across the five pain intensity classes for each modality, covering all 87 subjects. The videos are recorded at a frame rate of 25 frames per second (FPS), while the biosignal (ECG, EMG, GSR) recordings are sampled at a rate of 512 Hz. It should be noted that data from *Part A* of the dataset were used, where the biosignals had been pre-filtered according to [41].

2) *AI4Pain Dataset*: The AI4Pain Grand Challenge 2024 dataset is a more recent addition to the field and is specifically designed for advanced pain recognition tasks using fNIRS and facial video data. Sixty-five ($n=65$) volunteers, including 23 females, took part in the experiment, with their ages ranging from 17 to 52 years. While the dataset also includes physiological signals such as photoplethysmography (PPG), electrodermal activity (EDA), and respiration (RESP), these additional data are not yet publicly available. The current version of the AI4Pain Challenge dataset is divided into three parts: training (41 Volunteers), validation (12 Volunteers), and testing (12 Volunteers). The fNIRS and video recording setup used in this dataset provides comprehensive data on both brain activity and facial movements. The fNIRS data was recorded using an *Artinis* device (Artinis Medical Systems, Gelderland, the Netherlands). This device measures changes in the concentrations of oxygenated haemoglobin (HBO₂) and deoxygenated haemoglobin (HHB) (in $\mu\text{mol/L}$). The fNIRS system includes 24 channels covering the prefrontal cortex, with optodes (10 sources and 8 detectors) separated by 30 mm. Near-infrared light is emitted by sources with wavelengths of 760 nm and 840 nm, with a sampling rate of 50 Hz. The second sensing technology is a video camera (*Logitech StreamCam*) that captures facial video data and movements at a sampling rate of 30 FPS. The *AI4Pain* dataset is stratified into three distinct categories of varying levels of pain intensity: *No Pain*, *Low Pain*, and *High Pain*. Specifically, the dataset consists of 65 instances (each lasting 60s) of *No Pain*, 780 instances of *Low Pain* (each lasting 10s), and 780 instances of *High Pain* (each lasting 10s). The *No Pain* category includes instances from the baseline period at the start of each experiment, providing fNIRS and facial video data for comparison with pain-induced responses. The *Low Pain* category comprises instances of mild pain based on the pain tolerance test, capturing subtle neurological and behavioural changes in the corresponding fNIRS and video data. Finally, the *High Pain* category consists of instances where subjects experienced significant pain, also based on the pain tolerance test, leading to notable physiological and behavioural responses in the fNIRS and video data. It should be noted that the raw fNIRS signals were obtained directly from the dataset providers, with no pre-processing or filtering applied.

IV. EXPERIMENTAL EVALUATION & RESULTS

In the context of the present study, several scenarios were designed, both unimodal and multimodal, and used to evaluate the effectiveness of the proposed foundation model. The objective is to leverage various behavioral and physiological




Fig. 4: A high-level overview of the presented framework. *PainFormer*, the foundation model, is capable of extracting high-quality embeddings from a wide range of different behavioral and physiological modalities. Evaluating RGB, thermal, and depth videos and various representations of ECG, EMG, GSR, and fNIRS, including waveforms and spectrograms, demonstrate the comprehensive information encapsulated within these embeddings. Utilizing the embeddings from the *PainFormer* enables the development of diverse unimodal and multimodal pipelines for the pain assessment task. Each pipeline can be customized according to the modalities used, dataset characteristics, and the requirements of the target application or clinical environment. Our evaluations involved developing and applying various pipelines in unimodal and multimodal settings, achieving state-of-the-art results across different modalities and data representations.

modalities to evaluate *PainFormer*'s ability to extract and deliver high-quality embeddings for pain assessment. The experimental setup incorporates a wide array of modalities, including RGB, synthetic thermal imaging, depth videos, and physiological measurements like ECG, EMG, GSR, and fNIRS, which feature both waveform and spectrogram representations. Furthermore, particular pipelines were designed to suit either individual modalities or a combination thereof, depending on their specific integration requirements. This customization is central to our proposal since different pipelines may need to be developed based on particular needs, available data, or specific application needs. We want to be able to provide rich feature representations regardless of the input modality and perform exceptionally well across every modality and scenario. In Fig. 4, a high-level overview of the proposed framework is presented. It should be noted that all images, whether video frames or visual representations of biosignals are standardized to a resolution of 224×224 pixels.

From the *BioVid* dataset, we utilized *Part A* for experiments in a binary setting, specifically NP vs. P₄. Validation was conducted using the leave-one-subject-out (LOSO) cross-validation method. For the *AI4Pain* dataset, multilevel classification was performed using the three available pain levels. The validation

protocol employed is the hold-out method, as provided by the challenge organizers. For both datasets, the classification metrics utilized include accuracy, recall (sensitivity), and F1 score. Additionally, it should be noted that all conducted experiments adhere to a deterministic approach, ensuring they are not influenced by random initializations each time they are performed. This methodology guarantees that any observed differences in performance are solely the result of specific optimization parameters, modalities, and other deliberate variations rather than randomness.

A. BioVid

A plethora of experiments were conducted using the *BioVid* dataset. For the behavioral modalities beyond the original RGB, synthetic thermal and depth videos were also created in order to introduce additional visual representations, as described in III-B. Also, regarding biosignals, four different representations for ECGs, EMGs, and GSRs were evaluated, as presented in Section III-C. Additionally, combinations of these representations were tested. These are described in what follows.

1) *Video*: Regarding the behavioral modalities from *BioVid*, the *PainFormer* extracts an embedding of dimensionality $d =$

160 for each input frame. The embeddings from the respective frames of a specific video are subsequently concatenated to form a unified embedding representation of the entire video:

$$\mathcal{V}_D = [d_1 \| d_2 \| \dots \| d_m], \quad D \in \mathbb{R}^{N_1}, \quad (11)$$

where m denotes the number of frames in a video, and N_1 represents the dimensionality of the unified embedding, calculated as $m \times d \rightarrow 138 \times 160 = 22,080$. The resulting embedding is fed into the *Embedding-Mixer* for the final pain assessment. Beginning with 200 epochs and using only augmentation techniques for the RGB videos, an accuracy of 71.83% was achieved, with recall reaching 74.52%. Similarly, the thermal and depth videos recorded accuracies of 69.83% and 69.00%, respectively. Increasing the training to 300 epochs with more intense augmentations and applying *Label Smoothing* for regularization improved RGB accuracy to 72.50%, though recall slightly decreased by 0.46%. In contrast, performance metrics for the thermal modality declined, suggesting a higher sensitivity to augmentations and regularizations in this synthetic modality. Meanwhile, the depth modality showed improved results, achieving 70.08% accuracy, 71.27% recall, and 69.63% F1 score, indicating a positive response to the adjusted training parameters. In the final experimental setup, training was extended to 600 epochs with lighter augmentations at a 0.7 probability, coupled with 0.1 for *Label Smoothing* and 0.5 for *DropOut*. This configuration yielded the highest results for the RGB videos, achieving 76.29% accuracy and 77.56% recall. Notably, the F1 score saw the most significant improvement, with an increase of over 5%, reaching 75.56%. Similar patterns, albeit with smaller gains, were observed for the thermal and depth videos in the final experimental configuration. Accuracy reached 71.55% for thermal and 71.67% for depth videos, with recall rates nearly identical at 72.83% and 72.84%, respectively. These results underline a consistent improvement across all visual modalities with the adjusted training parameters and increased training time. Table V summarizes all the aforementioned experiments and results, highlighting that the RGB modality outperforms the others while the thermal and depth modalities exhibit similar performance levels. Additionally, although thermal and depth videos increase performances, the improvements are modest, suggesting that they may have reached their maximum potential results.

2) *ECG*: The training configuration applied to the video data was replicated for the ECG signals. Furthermore, as previously mentioned, four visual representations were employed. Similar to a video frame, each representation corresponds to an image of 224×224 pixels, from which an embedding with a dimensionality of $d = 160$ is extracted before being supplied to the *Embedding-Mixer*. Starting again with 200 epochs, using minimal augmentation, and without applying regularization techniques, the *waveform* representation achieved an accuracy of 69.58%, with recall and F1 scores reaching 72.67% and 68.10%, respectively. The *spectrogram-angle* showed lower performance across all metrics with an accuracy of 65.58%, whereas the *spectrogram-phase* recorded better accuracy, outperforming the previous two by 0.5% and 4.5%, respectively. The *spectrogram-PSD* delivered the best results, achieving

TABLE V: Classification results utilizing the video modality, NP vs. P₄ task, reported on accuracy, recall and F1 score.

Input	Epochs	Augmentation		Regularization		Metric		
		<i>Basic</i>	<i>Masking</i>	<i>LS</i>	<i>DropOut</i>	Acc	Rec	F1
RGB	200	0.5	0.5 10-20	0.0	0.0	71.83	74.52	70.29
	300	0.7	0.7 15-20	0.1	0.0	72.50	74.06	70.93
	600	0.5	0.5 15-20	0.1	0.5	76.29	77.56	75.56
Thermal	200	0.5	0.5 10-20	0.0	0.0	69.83	71.51	69.17
	300	0.7	0.7 15-20	0.1	0.0	68.83	69.77	68.41
	600	0.5	0.5 15-20	0.1	0.5	71.55	72.83	71.12
Depth	200	0.5	0.5 10-20	0.0	0.0	69.00	69.44	67.94
	300	0.7	0.7 15-20	0.1	0.0	70.08	71.27	69.63
	600	0.5	0.5 15-20	0.1	0.5	71.67	72.84	71.26

LS: *Label Smoothing* For Augmentation and Regularization, the number denotes the probability of application, while in *Masking*, the number in | | indicates the size of the mask applied.

TABLE VI: Classification results utilizing the ECG modality, NP vs. P₄ task, reported on accuracy, recall and F1 score.

Input	Epochs	Augmentation		Regularization		Metric		
		<i>Basic</i>	<i>Masking</i>	<i>LS</i>	<i>DropOut</i>	Acc	Rec	F1
Wave	200	0.5	0.5 10-20	0.0	0.0	69.58	72.67	68.10
	300	0.7	0.7 15-20	0.1	0.0	71.08	72.74	70.22
	600	0.5	0.5 15-20	0.1	0.5	73.36	74.75	72.52
Angle	200	0.5	0.5 10-20	0.0	0.0	65.58	66.68	64.89
	300	0.7	0.7 15-20	0.1	0.0	66.33	68.22	65.22
	600	0.5	0.5 15-20	0.1	0.5	68.25	71.24	66.99
Phase	200	0.5	0.5 10-20	0.0	0.0	70.08	71.54	69.40
	300	0.7	0.7 15-20	0.1	0.0	72.33	73.73	71.69
	600	0.5	0.5 15-20	0.1	0.5	72.70	74.19	72.14
PSD	200	0.5	0.5 10-20	0.0	0.0	71.08	73.13	70.19
	300	0.7	0.7 15-20	0.1	0.0	71.50	73.14	70.18
	600	0.5	0.5 15-20	0.1	0.5	75.49	77.15	74.90

71.08%, 73.13%, and 70.19% across the three metrics. The same trend continued in the 300-epoch configuration, which improved results across all representations and metrics. In the final experimental setup of 600 epochs, while increases were observed across the board, the *spectrogram-PSD* demonstrated the most significant improvements, nearly 4%, achieving 75.49% accuracy, 77.15% recall, and 74.90% F1 score. This suggests that for ECG signals, integrating amplitude and frequency information provided by the PSD representation is the most valuable and effective for this analysis. Table VI presents the results for the ECG modality.

3) *EMG*: Regarding the EMG signals, the initial configuration of 200 epochs yielded similar accuracy results across the *waveform*, *spectrogram-phase*, and *spectrogram-PSD* representations, with scores of 68.75%, 68.33%, and 69.25%. In contrast, the *spectrogram-angle* representation underperformed, recording an accuracy of 66.42%, a trend also observed in the ECG modality. Interestingly, in the subsequent training configuration with extended epochs and more intensive augmentation and regularization, the *spectrogram-*

TABLE VII: Classification results utilizing the EMG modality, NP vs. P₄ task, reported on accuracy, recall and F1 score.

Input	Epochs	Augmentation		Regularization		Metric		
		<i>Basic</i>	<i>Masking</i>	<i>LS</i>	<i>DropOut</i>	Acc	Rec	F1
Wave	200	0.5	0.5 10-20	0.0	0.0	68.75	70.55	67.93
	300	0.7	0.7 15-20	0.1	0.0	69.83	72.52	68.68
	600	0.5	0.5 15-20	0.1	0.5	72.07	73.64	71.48
Angle	200	0.5	0.5 10-20	0.0	0.0	66.42	68.57	65.26
	300	0.7	0.7 15-20	0.1	0.0	63.92	66.33	62.67
	600	0.5	0.5 15-20	0.1	0.5	65.32	68.15	63.77
Phase	200	0.5	0.5 10-20	0.0	0.0	68.33	69.75	67.68
	300	0.7	0.7 15-20	0.1	0.0	68.58	70.00	67.97
	600	0.5	0.5 15-20	0.1	0.5	69.37	71.17	68.66
PSD	200	0.5	0.5 10-20	0.0	0.0	69.25	70.38	68.84
	300	0.7	0.7 15-20	0.1	0.0	69.67	71.06	69.12
	600	0.5	0.5 15-20	0.1	0.5	72.10	72.82	71.82

angle representation exhibited a notable decline in performance across all metrics, unlike the other representations. Specifically, in the 300-epoch configuration, despite some improvement, the angle representation still lagged behind the initial results, achieving an accuracy of 65.32%, with recall and F1 scores of 68.15% and 63.17%, respectively. This indicates that the phase spectrum without phase unwrapping is not well-suited for pain assessment tasks using EMG signals. Conversely, the other visual representations consistently showed performance improvements in each configuration. Notably, the *spectrogram-PSD* reached the highest accuracy at 72.10% and an F1 score of 71.82%. The *waveform* representation achieved the highest recall at 73.64%. In Table VII, the results for the EMG modality are presented.

4) *GSR*: Regarding GSR, we observe an apparent difference in performance among the four representations. *Waveform*-based representations significantly outperform the others, with an initial accuracy of 87.75% in the 200-epoch configuration, surpassing other metrics by over 14%. During longer training sessions, there is a slight improvement across representations, suggesting that the GSR modality may have reached its performance plateau. The *spectrogram-phase* is the most informative among spectrograms, achieving final accuracy, recall, and F1 scores of 76.41%, 77.23%, and 76.47%, respectively. The *waveform* representation is the most effective, achieving the highest scores of 88.99% for accuracy, 89.55% for recall, and 88.88% for the F1 metric. This variation in performance can be attributed to the inherent nature of the GSR signal. As shown in Fig. 4, GSR typically appears as a smooth curve with gradual slopes, reflecting the slow and steady changes in skin conductivity due to variations in sweat gland activity triggered by stress or arousal. In contrast, EMG signals display sharp spikes and erratic fluctuations, indicative of the rapid electrical activity associated with skeletal muscle contractions. Similarly, ECG signals are characterized by distinct cyclical patterns such as the P and T waves and the QRS complex. This suggests that the absence of more complicated patterns in GSR is not well-suited for spectral and frequency domain analysis,

TABLE VIII: Classification results utilizing the GSR modality, NP vs. P₄ task, reported on accuracy, recall and F1 score.

Input	Epochs	Augmentation		Regularization		Metric		
		<i>Basic</i>	<i>Masking</i>	<i>LS</i>	<i>DropOut</i>	Acc	Rec	F1
Wave	200	0.5	0.5 10-20	0.0	0.0	87.75	88.68	87.56
	300	0.7	0.7 15-20	0.1	0.0	88.50	89.16	88.34
	600	0.5	0.5 15-20	0.1	0.5	88.99	89.55	88.88
Angle	200	0.5	0.5 10-20	0.0	0.0	73.67	75.00	73.26
	300	0.7	0.7 15-20	0.1	0.0	73.08	74.60	72.66
	600	0.5	0.5 15-20	0.1	0.5	73.24	75.02	72.83
Phase	200	0.5	0.5 10-20	0.0	0.0	75.17	76.13	74.79
	300	0.7	0.7 15-20	0.1	0.0	75.92	76.60	75.57
	600	0.5	0.5 15-20	0.1	0.5	76.41	77.23	76.47
PSD	200	0.5	0.5 10-20	0.0	0.0	72.83	73.91	72.34
	300	0.7	0.7 15-20	0.1	0.0	73.08	73.96	72.68
	600	0.5	0.5 15-20	0.1	0.5	73.96	74.81	73.50

which is better leveraged by spectrograms. However, *waveform* representations yield the best results for GSR and outperform all other modalities and visual representations, underscoring their effectiveness in capturing essential physiological information from GSR signals. Table VIII summarizes the results for the GSR modality.

5) *Fusion*: Multiple fusion approaches have been explored to determine whether integrating multiple representations or modalities leads to performance enhancements. In this study, a combination of inputs from the same sensor type, such as RGB and depth-estimation videos or ECG *waveforms* and ECG *spectrogram-PSD*, is categorized as an unimodal fusion approach. Conversely, combinations involving inputs from different sensor types, such as GSR and EMG, are classified as multimodal. There are three primary fusion methods, including two feature fusion techniques and one decision fusion technique. Feature fusion includes addition, where embeddings from inputs are combined before proceeding to the subsequent module, and concatenation, which concatenates them along the *y-axis*. Decision fusion involves processing each embedding through the final module, the *Embedding-Mixer*, which compiles predictions from each source and aggregates them to produce the final prediction. All corresponding experiments follow the 600-epoch training configuration as previously presented for the individual input representations. Table IX presents the results of the fusion approaches that were conducted.

In exploring video modality fusion, four combinations were assessed: RGB and thermal, RGB and depth, thermal and depth, and the aggregation of all three—RGB, thermal, and depth. For the RGB and thermal video combination, regardless of the fusion technique employed, performance was inferior to using RGB alone, with the highest accuracy of 75.66% attained via decision fusion. Similarly, the RGB and depth video combination yielded the best results with decision fusion, with a 75.53% accuracy rate, yet this still fell short of the performance achieved with RGB alone. Interestingly, the fusion of thermal and depth videos marked an improvement over using depth alone, which was more effective when used

independently. With decision fusion, this combination achieved a 73.02% accuracy rate, surpassing depth alone by 1.35%. Other metrics also observed improvements, with recall and F1 scores reaching 74.46% and 72.59%, respectively. Finally, the three video inputs of RGB thermal and depth were the only combination that led to improved results compared to the RGB in isolation. Even though some improvements were observed with the addition method, the best results came through the decision fusion, with accuracy, recall, and F1 scores of 76.55%, 77.91%, and 76.11%, respectively—marking improvements of 0.26%, 0.35%, and 0.55%. It is noted that in all experiments involving video representations, decision fusion consistently outperformed the addition method.

Regarding biosignals, experiments were conducted solely with ECG and EMG, focusing on the two most effective representations: *waveform* and *spectrogram-PSD*. Fusion experiments for the GSR were not performed, as the waveform representations significantly outperform other visualizations, rendering further fusion unnecessary. For the ECG, all fusion methods yielded inferior results compared to the *spectrogram-PSD*, with the exception of the addition method, which showed an improvement of 0.21% in recall. In EMG biosignals, enhanced performance was noted across all approaches. The most effective method was concatenation, which led to increases of 0.74%, 0.36%, and 0.64% in the three metrics, achieving respective scores of 72.84%, 74.00%, and 72.46%.

In our multimodal scenario, physiological and behavioral modalities were combined, utilizing GSR signals for the physiological aspect and RGB, synthetic thermal, and estimated depth videos for the behavioral aspect. The *waveform* representation of the GSR signal was used to extract the corresponding embedding vector, where the individual video feature representations for RGB, thermal, and depth, described in 11, were aggregated to create a unified vector of dimension 22,080. This vector was then fed into the *Video-Encoder* and encoded into a lower dimensional space of 40. The representations from the GSR and the videos were then concatenated to create a fused vector with a dimension of $160 + 40 = 200$. The entire process can be expressed as follows:

$$\mathcal{M}_h = \mathcal{G}_d \| \text{Enc}[(\mathcal{V}_D^{\text{RGB}} + \mathcal{V}_D^{\text{Thermal}} + \mathcal{V}_D^{\text{Depth}})], \quad h \in \mathbb{R}^{N_2}, \quad (12)$$

where $\mathcal{G}$ represents the GSR embedding, $\mathcal{M}$ is the final fused vector and N_2 equal to 200. Fig. 4 (bottom right) depicts the specific process. This combination achieved the highest results in the study, reaching 89.08%, 89.88%, and 88.87% for accuracy, recall, and F1, respectively. Notably, this combination marginally outperformed the GSR modality when used as an individual input in terms of accuracy and recall.

B. AI4Pain

In the *AI4Pain* dataset, experiments were conducted using unimodal and multimodal approaches. The original RGB videos were utilized for the behavioral modality, while the physiological modality involved waveforms from the fNIRS's HBO2 channels. We note that out of the 24 available HBO2 channels, 2 were excluded due to malfunctions. Table X presents the corresponding results.

TABLE IX: Classification results on fusion settings*, NP vs. P₄ task, reported on accuracy, recall and F1 score.

Modality	Input	Fusion	Metric		
			Acc	Rec	F1
Video	RGB, Thermal	Add	75.09	76.97	73.98
			-1.20	-0.59	-1.58
		DF	75.66	77.23	75.08
			-0.63	-0.33	-0.48
	RGB, Depth	Add	74.93	76.41	73.38
			-1.36	-1.15	-2.18
		DF	75.53	77.18	75.00
			-0.76	-0.38	-0.56
	Thermal, Depth	Add	71.44	73.15	70.73
			-0.23	+0.31	-0.50
		DF	73.02	74.46	72.59
			+1.35	+1.62	+1.33
RGB, Thermal, Depth	Add	76.26	77.70	75.78	
		-0.03	+0.14	+0.22	
	DF	76.55	77.91	76.11	
		+0.26	+0.35	+0.55	
ECG	Wave, PSD	Add	75.43	77.36	74.75
			-0.06	+0.21	-0.15
		Concat	74.74	76.77	74.00
			-0.75	-0.38	-0.90
EMG	Wave, PSD	Add	72.79	74.15	72.28
			+0.69	+0.51	+0.46
		Concat	72.84	74.00	72.46
			+0.74	+0.36	+0.64
Video, GSR	RGB, Thermal, Depth, Wave	Add & Concat	89.08	89.88	88.87
			+0.09	+0.33	-0.01

*: All experiments follow the augmentation and regularization settings for the 600 epoch configuration outlined in the unimodal experiments. + and - indicate an increase or decrease in performance, respectively, compared to the best unimodal input approach. DF: Decision Fusion Add: Addition Concat: Concatenation

1) *Video*: Similar to IV-A1, an embedding of $d = 160$ is extracted for every frame, but here, the extracted embeddings are aggregated into a fused vector:

$$\mathcal{V}_d = [d_1 + d_2 + \dots + d_m], \quad d \in \mathbb{R}^{N_3}, \quad (13)$$

where m denotes the number of frames in a video, and N_3 represents the dimensionality of the unified embedding, equal to 160. After feeding the embedding into the *Embedding-Mixer* and following the same 600-epoch training configuration as in previous experiments, the setup achieved an accuracy of 49.77%, with recall and F1 scores of 50.11% and 49.77% respectively. Increasing the *DropOut* rate to 0.3 improved the accuracy and F1 scores to 51.39% and 51.31%. Further elevating the *DropOut* to 0.8 enhanced the recall to 52.74%.

2) *fNIRS*: Similarly, embeddings were aggregated for the 22 HBO2 channels, resulting in a feature representation of $\mathcal{O}_d = 160$. The 600-epoch training setup initially yielded 43.06%, 42.80%, and 42.07% for the three metrics. By increasing the *DropOut* to 0.3, a peak performance of 44.44%, 45.55%, and 43.74% was achieved.

3) *Fusion*: For the fusion of video and fNIRS data, an aggregation approach was employed:

$$\mathcal{F}_d = \mathcal{V}_d + \mathcal{O}_d, \quad d \in \mathbb{R}^{N_3}, \quad (14)$$

TABLE X: Classification results on the validation set of *AI4Pain* dataset, multilevel classification task, reported on accuracy, recall and F1 score.

Input	Epochs	Augmentation		Regularization		Metric		
		<i>Basic</i>	<i>Masking</i>	<i>LS</i>	<i>DropOut</i>	Acc	Rec	F1
Video	600	0.5	0.5 15-20	0.1	0.5	49.77	50.11	49.77
	600	0.5	0.5 15-20	0.1	0.3	51.39	51.50	51.31
	600	0.5	0.5 15-20	0.1	0.8	48.38	52.74	46.69
fNIRS	600	0.5	0.5 15-20	0.1	0.5	43.06	42.80	42.07
	600	0.5	0.5 15-20	0.1	0.3	44.44	45.55	43.74
	600	0.4	0.4 15-20	0.1	0.1	43.06	44.18	42.44
Fusion	600	0.5	0.5 15-20	0.1	0.5	50.00	51.01	48.54
	600	0.1	0.1 15-20	0.1	0.8	50.23	50.25	50.24
	600	0.4	0.4 15-20	0.1	0.6	51.85	51.87	51.35

Fusion: the Addition method of the modalities applied

where $\mathcal{F}_d$ is the combined feature representation. Beginning with the same 600-epoch training setup, 50.00%, 51.01%, and 48.54% results were reached for the three metrics. Increasing the *DropOut* to 0.8 improved the accuracy and F1 score by 0.23% and 1.7%, respectively, with the recall decreased by 0.75%. The optimal *DropOut* setting of 0.6 achieved peak performances of 51.85%, 51.57%, and 51.35% for the accuracy, recall, and F1 scores.

V. COMPARISON WITH EXISTING METHODS

To evaluate *PainFormer*, we benchmark our approach with studies from the literature. Specifically, studies that employ the *BioVid* dataset (Part A), utilize all available subjects (87), conduct the same task, follow the leave-one-subject-out (LOSO) validation protocol, and report accuracy metrics. Similarly, for the *AI4Pain* dataset, our comparisons were with studies that adhered strictly to the evaluation guidelines outlined in the corresponding challenge.

In *BioVid*, the proposed approach utilizing RGB, thermal, and depth video inputs in video-based studies ranks among the highest performance. Achieving an accuracy of 76.55% surpasses all approaches that utilize hand-crafted features, such as those in references [41], [50], [87]. Additionally, it outperforms the majority of deep learning-based methods, such as [51], [88]–[90]. However, there are exceptions, specifically, the results reported in [49] at 77.10%, in [91] at 77.30%, and in [52] at 78.90%. Moreover, the authors in [53] achieved 77.50% using a 3D CNN approach. Combined with the pseudo heart rate extracted from the videos, they reached the highest reported results reported in the literature at 88.10%. Table XI presents the results.

Regarding the biosignals, it is noted that in ECG-based studies, *PainFormer* achieved the highest results in the literature, with an accuracy of 75.49% using the *spectrogram-PSD* representation. Compared to the subsequent leading studies [49], [92], our approach shows a significant improvement, outperforming them by more than 6% and 8%, respectively. In the relatively few EMG-based studies employing the addition of *waveform* and *spectrogram-PSD* representations, we yielded

a 72.84% accuracy, significantly outperforming the following closest study [41], which achieved 57.90%. In the GSR-based studies, our approach, utilizing solely *waveform* representation and achieving an 88.99% accuracy, leads to performance. It is also observed that studies using raw biosignals rather than extracting domain-specific features generally exhibit better results. The second [93] and third [94] ranked studies achieved 85.56% and 84.80% accuracy, respectively. The Table XII presents the corresponding results for the biosignals.

In multimodal approaches, our method utilizing video inputs and GSR achieved an 89.08% accuracy, marking the highest result reported in the literature (refer to Table XIII). Additionally, with one exception [49], all documented studies incorporated the GSR signal as one of the inputs. This is consistent with findings that GSR is the most effective modality for pain assessment. For instance, the second-highest-performing study [95], which used a combination of GSR and ECG signals, achieved 87.06%. At the same time, the authors in [58], which included videos, ECG, EMG, and GSR, reached an accuracy of 86.00%.

Finally, concerning the *AI4Pain* dataset, *PainFormer* achieved 53.67% accuracy using the RGB video modality, outperforming [96] with 49.00% but falling short of [97], which reached 55.00% using a transformer-based masked autoencoder. Utilizing only the fNIRS, an accuracy of 52.60% was achieved, which is approximately 1% lower than the 53.66% reported in [98]. In a multimodal approach combining videos with waveform representations, an accuracy of 55.69% was achieved, surpassing [66] nearly 1%, [67] by more than 9%, [99] by over 4%, and [100] (using only EDA) by 0.52%, establishing the highest performance reported on this dataset to date. Table XIV presents the corresponding results.

VI. INTERPRETATION

Enhancing the interpretability of models is essential for their acceptance and effective integration into clinical settings. In this study, the *PainFormer* has been used to generate attention maps, as shown in Fig. 5. The weights from the “Stage 4” self-attention heads have been applied by interpolating them onto the input images, allowing us to visualize the model’s attention areas.

In Fig. 5(a), (1st row), we showcase examples from the RGB, thermal, and depth modalities, and in Fig. 5(a), (2nd row), we present the corresponding attention maps. Observations indicate that the model primarily focuses on the glabella region (the area between the eyebrows) in the RGB frame, a zone for manifesting facial expressions. Additional focus is seen on the mental protuberance area (the chin), which is also associated with expressions of pain. For the thermal frame, the model concentrates on areas around the eyes and the left and right parts of the mouth. Interestingly, these areas correspond to brighter colors in the thermal imagery, indicating higher temperatures rather than direct facial expressions. This suggests that temperature variations influence the model’s attention in the thermal frame rather than facial movements and expressions. In the depth frame, the model targets areas showcasing variations in depth, particularly across the horizontal eye region. There

TABLE XI: Comparison of video-based studies utilizing *BioVid* (Part-A), NP vs. P₄ task and LOSO cross-validation.

Study	Method		Acc%
	Features	ML	
[88]	raw	SLSTM	61.70
[89]	raw	2D CNN, biLSTM	69.25
[87]	optical flow	RF	70.20
[90]	raw	2D CNN	71.00
[34]	raw [⊙]	Vision-MLP	71.03
[101] [†]	raw	2D CNN	71.30
[41]	facial landmarks, 3D distances	RF	71.60
[102]	facial 3D distances	Deep RF	72.10
[102]	facial action descriptors	Deep RF	72.40
[59]	facial landmarks, 3D distances	RF	72.70
[50]	fiducial points	GNN	73.20
[51]	raw	Transformer	73.28
[103] [†]	raw	2D CNN, GRU	73.90
[49]	raw	Transformer	77.10
[91]	facial landmarks	STAGCN	77.30
[53]	raw	3D CNN	77.50
[52]	raw, rPPG [♦]	3D CNN	78.90
[53]	raw, heart rate [★]	3D CNN	88.10
Our	raw [✱]	Transformer	76.55

†: reimplemented for pain intensity estimation on *BioVid* by [53] ⊙: RGB, synthetic thermal videos ♦: remote photo plethysmography (estimated from videos) ★: pseudo heart rate gain (estimated from videos) ✱: RGB-thermal-depth (DF) ML: Machine Learning SLSTM: Sparse Long Short-Term Memory biLSTM: Bidirectional LSTM RF: Random Forest MLP: Multi-Layer Perceptron GNN: Graph Neural Networks GRU: Gated Recurrent Unit STAGCN: Spatio-temporal Attention Graph Convolution Network

is also slight attention to the frame's lower left and right edges, illustrating depth differences in body parts other than the face, indicating a nuanced understanding of the model's representation of depth.

The ECG attention maps shown in Fig. 5(b), (top left), emphasize primarily a distinct R peak in the trace's center. Notable attention is also directed towards the T waves, especially in the first part of the signal and the T waves that follow the central R peak. Interestingly, in the EMG attention maps of Fig. 5(b), (top right), the *PainFormer* primarily focuses on the initial and middle sections of the signals. Despite the presence of a muscle contraction burst later in the sequence, the model exhibits less attention to this portion. The *Silent-EMG* dataset [78], on which the *PainFormer* was pre-trained, could relate to this observation. This pre-training background might influence the model's attention and responsiveness to specific sections of the EMG signals. For the GSR signal in Fig. 5(b), (bottom left), mild attention is observed at the onset of the response, marking the beginning of the conductance increase. The most intense attention is near the peak amplitude, where the conductance reaches its maximum level. For the fNIRS signal in Fig. 5(b) (bottom right), the attention map predominantly highlights regions where attention colors align with peaks and rapid changes in HbO₂ levels. Notable attention is concentrated in the left, middle, and right sections of the map, where distinct peaks and dips in the signal are observed. This indicates that the *PainFormer* consistently focuses on

TABLE XII: Comparison of biosignal-based studies utilizing *BioVid* (Part-A), NP vs. P₄ task and LOSO cross-validation.

Study	Modality	Method		Acc%
		Features	ML	
[54]	ECG	raw	1D CNN	57.04
[55]	ECG	domain-specific [*]	LR	57.40
[104]	ECG	domain-specific [*]	LR	57.69
[105]	ECG	domain-specific [*]	SVM	58.39
[106]	ECG	domain-specific [*]	SVM	58.62
[94]	ECG	raw	1D CNN, biLSTM	61.20
[41]	ECG	domain-specific [*]	RF	62.00
[59]	ECG	domain-specific [*]	SVM	62.40
[107]	ECG	raw	2D CNN, biLSTM	63.20
[53]	ECG	heart rate [★]	3D CNN	65.00
[49]	ECG	heart rate	Transformer	67.04
[92]	ECG	domain-specific [*]	FCN	69.40
Our	ECG	raw [✱]	Transformer	75.49
[105]	EMG	domain-specific [*]	LSTM	56.83
[41]	EMG	domain-specific [*]	RF	57.90
[55]	EMG	domain-specific [*]	LR	58.59
[54]	EMG	raw	2D CNN	58.65
Our	EMG	raw [✱]	Transformer	72.84
[41]	GSR	domain-specific [*]	RF	73.80
[104]	GSR	domain-specific [*]	LR	74.21
[59]	GSR	domain-specific [*]	RF	74.40
[105]	GSR	domain-specific [*]	LSTM	76.86
[108]	GSR	domain-specific [*]	RF	80.40
[56]	GSR	domain-specific [*]	RF	81.90
[55]	GSR	domain-specific [*]	LR	82.36
[109]	GSR	domain-specific [*]	SVM	83.30
[107]	GSR	raw	1D CNN, biLSTM	83.60
[110]	GSR	domain-specific [*]	MLP	84.22
[54]	GSR	raw	1D CNN	84.57
[94]	GSR	raw	1D CNN, biLSTM	84.80
[93]	GSR	raw	1D CNN Transformer	85.56
[111]	GSR	raw	1D CNN Transformer	86.21
Our	GSR	raw [✱]	Transformer	88.99

✱: numerous features ★: pseudo heart rate gain (estimated from videos) ♦: PSD ✱: waveform-PSD (Concat) ✱: waveform SVM: Support Vector Machines LR: Logistic Regression

significant fluctuations in the HbO₂ signal, which are likely correlated with pain conditions. Additionally, regions with lower or moderate attention correspond to parts of the time series with stable or minor variations in HbO₂, reflecting lower levels of brain activation typically associated with mild or no pain responses.

VII. COMPUTATIONAL RESOURCES AND RUNTIME

The development of *Painformer* and all presented experiments was conducted on a workstation running Ubuntu 22.04 LTS, with Python 3.9.18, Torch 1.11.0, and CUDA 12.9. The hardware configuration featured an NVIDIA RTX 3090 GPU paired with an AMD Ryzen 9 5950X 16-core CPU. Pre-training the foundation model using the proposed multi-task learning

TABLE XIII: Comparison of multimodal-based studies utilizing *BioVid* (Part-A), NP vs. P₄ task and LOSO cross-validation.

Study	Modality	Method		Acc%
		Features	ML	
[104]	ECG, GSR	domain-specific*	SVM	72.20
[41]	ECG, EMG, GSR	domain-specific*	RF	74.10
[105]	ECG, EMG, GSR	domain-specific*	LSTM	77.21
[41]	Video ¹ , ECG ² , EMG ² , GSR ²	facial landmarks ¹	RF	77.80
		3D distances ¹		
[59]	Video ¹ , ECG ² , GSR ²	domain-specific ² *	RF	78.90
		facial landmarks ¹		
[49]	Video ¹ , ECG ²	3D distances ¹	Transformer	82.74
		domain-specific ² *		
[56]	Video ¹ , ECG ² , EMG ² , GSR ²	raw ¹ , heart rate ²	RF	83.10
		geometric ¹		
[55]	ECG, EMG, GSR	appearance ¹	RF	83.20
		domain-specific ²		
[112]	ECG, EMG, GSR	domain-specific	LR	83.30
[113]	ECG, EMG, GSR	domain-specific	biLSTM	83.30
[114]	ECG, EMG, GSR	raw	DDCAE	83.99
[54]	ECG, EMG, GSR	raw	DDCAE, NN	84.25
		raw	2D CNN	
[115]	GSR, ECG	domain-specific*	NN	84.58
[107]	Video, GSR	raw	2D CNN	84.80
			biLSTM	
[94]	ECG, GSR	raw	ID CNN	84.80
			biLSTM	
[57]	ECG, EMG, GSR	domain-specific	RF	85.70
[116]	ECG, EMG, GSR	domain-specific	RF	85.80
[58]	Video ¹ , ECG ² , EMG ² , GSR ²	facial descriptors ¹	RF	86.00
		domain-specific ²		
[95]	GSR, ECG	domain-specific*	NN	87.06
Our	Video ⁺ , GSR*	raw	Transformer	89.08

+: RGB-thermal-depth *: waveform *: numerous features DDCAE: Deep Denoising Convolutional Autoencoders NN: Neural Network

TABLE XIV: Comparison of studies on the testing set of *AI4Pain* dataset.

Study	Modality						ML	Acc%
	Video	fNIRS	EDA	BVP	Resp	SpO ₂		
[67] [†]	✓	✓	–	–	–	–	Transformer	46.67
[96] [†]	✓	–	–	–	–	–	2D CNN	49.00
[99] [†]	✓	✓	–	–	–	–	CNN-Transformer	51.33
[98] [†]	–	✓	–	–	–	–	ENS	53.66
[97] [†]	✓	–	–	–	–	–	Transformer	55.00
[117] [‡]	–	–	–	–	✓	–	Transformer	42.24
[66] [‡]	–	–	✓	✓	✓	✓	MoE	54.89
[100] [‡]	–	–	✓	–	–	–	Transformer	55.17
Our [†]	–	✓	–	–	–	–		52.60
	✓	–	–	–	–	–	Transformer	53.67
	✓	✓	–	–	–	–		55.69

BVP: Blood Volume Pulse Resp: Respiration Signal SpO₂: Peripheral Oxygen Saturation ENS: Ensemble Classifier MoE: Mixture of Experts †: AI4PAIN-First Multimodal Sensing Grand Challenge ‡: AI4PAIN-Second Multimodal Sensing Grand Challenge




Fig. 5: Attention maps from the *PainFormer*: (a)(1st row) frames from RGB, thermal, and depth video modalities; (a)(2nd row) corresponding attention maps; (b)(1st row) attention maps for ECG and EMG; (b)(2nd row) attention maps for EDA and fNIRS modalities.

approach took approximately 6 months, while pain-related tasks required 2-7 days, depending on the number of modalities and the selected fusion strategy.

VIII. DISCUSSION

In this study, we introduced *PainFormer*, a vision foundation model designed for pain assessment tasks across various input modalities. *PainFormer* leverages a vision-transformer architecture, pre-trained on 14 tasks/datasets encompassing 10.9 million samples using a multi-task learning approach. The foundation model is supported by supplementary models: *Embedding-Mixer*, a transformer-based model that processes extracted embeddings for pain assessment, and *Video-Encoder*, which encodes video embeddings into a lower-dimensional space. This specific setup facilitates the effective integration of behavioral and physiological modalities. Furthermore, our approach was evaluated using two pain-related datasets, *BioVid* and *AI4Pain*. We explored a wide array of representations, including RGB, synthetic thermal, depth videos, and waveforms and spectrograms for biosignals such as ECG, EMG, GSR, and fNIRS. Utilizing this diverse set of modalities, we developed multiple pipelines in both unimodal and multimodal settings to assess the quality and effectiveness of the embeddings created by *PainFormer*.

The experiments revealed that the RGB video modality outperformed other video modalities, such as thermal and depth, achieving 76.29% accuracy in the *BioVid* dataset. Notably, thermal and depth also yielded robust results, with accuracies of 71.55% and 71.67%, respectively. When combining these video

representations, we noted subtle yet important improvements in performance compared to using RGB alone. Interestingly, the fusion of thermal and depth modalities enhanced accuracy by 1.35%, reaching 73.02%, which is closely aligned with RGB's performance, suggesting a possible equivalence. It is crucial to emphasize the significance of this finding, as synthetic thermal and depth estimation videos serve as an intermediary between the highly informative and effective RGB modality, which poses privacy concerns, and facial descriptions such as facial action units, which, while offering high privacy due to the absence of visible faces, tend to yield mediocre performance [41], [102]. Regarding the physiological modalities, the ECG demonstrated solid performance across all tested representations, with the *spectrogram-PSD* achieving the highest accuracy at 75.49%. However, combining various ECG representations did not yield any improvement. The notably challenging EMG modality performed exceptionally well, achieving over 72% accuracy with *waveform* and *spectrogram-PSD* representations. Interestingly, combining these representations resulted in improvements exceeding 0.5%. GSR signals, recognized as the most effective modality for pain assessment, achieved the highest results compared to other modalities, reaching an 88.9% accuracy with *waveform* representations. Crucially, our approach to integrating behavioral and physiological modalities with the *BioVid* dataset resulted in substantial performance improvements. By combining GSR embeddings with video embeddings from RGB, thermal, and depth modalities, we achieved an 89.08% accuracy. The experiments conducted with the *AI4Pain* dataset showed strong performances across the available modalities. Specifically, using RGB videos, an accuracy of 53.67% was achieved, higher than the 52.60% recorded with fNIRS, while the combination of modalities yielded an increase to 55.69%, demonstrating that data fusion, with proper optimization, can indeed enhance performance.

Creating attention maps to interpret how *PainFormer* processes inputs to generate embeddings yielded intriguing results. The foundation model consistently focuses on regions of interest. For instance, RGB videos emphasize distinct facial expressions, while thermal videos target areas displaying more vivid colors that correspond to higher temperatures. Regarding the depth videos, the analysis shows that the model focuses primarily on the eye area, which is prominent in terms of depth variations. Similarly, it highlights other body areas that exhibit notable depth differences compared to the head. The situation is analogous to the biosignals, where *PainFormer* can detect subtle yet significant variations across all inputs, including ECG, EMG, GSR, and fNIRS. This capability suggests that the model can capture a comprehensive representation of pain-related data, beneficial for research and clinical applications. However, it is important to note that further investigation is necessary. The pre-training process might bias the model's focus towards specific areas of interest, potentially leading to conflicts and inconsistencies with actual pain events.

Finally, the results are intriguing when comparing the proposed approach to existing methods in the literature. Across datasets and various modalities, our performances are consistently state-of-the-art. Our method ranks among the best reported using the *BioVid* dataset, particularly for video-based

approaches. However, it falls short of methods integrating facial features with cardiac-related information extracted directly from the videos [52], [53]. Interestingly, our method excels with biosignals, including ECG, EMG, and GSR, achieving the highest recorded performances. *PainFormer* also consistently delivers high results in unimodal and multimodal settings, underscoring its robustness regardless of the input and the pipelines. When using the *AI4Pain* dataset, our multimodal approach achieved the highest results reported to date. However, it should be noted that this dataset is relatively new, so further studies and extensive research are needed.

IX. CONCLUSION

This study presented *PainFormer*, a vision transformer-based foundation model for pain recognition tasks. Pre-trained on 10.9 million samples across various modalities and datasets, including facial recognition, emotion recognition, and biosignal datasets. This makes it the first general-purpose and foundation model in automatic pain assessment documented in the literature. Extensive evaluations across various input modalities and representations, including RGB, synthetic thermal, depth videos, and biosignals like ECG, EMG, GSR, and fNIRS, have demonstrated that the proposed model consistently extracts high-quality feature representations, regardless of the input. A direct comparison with 75 methodologies from the literature using two datasets exhibited that the model delivers state-of-the-art performance across all modalities in unimodal and multimodal configurations. Moreover, by generating attention maps, we offered insights into the foundation model's functioning, highlighting the specific focus areas within the inputs. Finally, we recommend that future research explore multi-modality approaches, which have proven to be the most effective for assessing pain phenomena in real-world settings. Additionally, synthetic data shows considerable promise, and further exploration is needed. Developing methods for interpretation is also essential, especially for the potential integration of these frameworks into clinical practice.

REFERENCES

- [1] V. Santiago, "Painful truth: The need to re-center chronic pain on the functional role of pain," *Journal of Pain Research*, vol. 15, pp. 497–512, 2022.
- [2] S. P. Cohen, L. Vase, and W. M. Hooten, "Chronic pain: an update on burden, best practices, and new advances," *The Lancet*, vol. 397, no. 10289, pp. 2082–2097, 2021.
- [3] A. C. d. C. Williams and K. D. Craig, "Updating the definition of pain," *Pain*, vol. 157, no. 11, pp. 2420–2423, nov 2016.
- [4] R.-D. Treede, W. Rief, A. Barke, Q. Aziz, M. I. Bennett, R. Benoliel, M. Cohen, S. Evers, N. B. Finnerup, M. B. First *et al.*, "A classification of chronic pain for icd-11," *Pain*, vol. 156, no. 6, pp. 1003–1007, 2015.
- [5] R. Sinatra, "Causes and consequences of inadequate management of acute pain," *Pain Medicine*, vol. 11, no. 12, pp. 1859–1871, 12 2010.
- [6] L. De Ruddere and R. Tait, *Facing Others in Pain: Why Context Matters*. Cham: Springer International Publishing, 2018, pp. 241–269.
- [7] A. Abdulla, N. Adams, M. Bone, A. M. Elliott, J. Gaffin, D. Jones, R. Knaggs, D. Martin, L. Sampson, P. Schofield, and British Geriatric Society, "Guidance on the management of pain in older people," *Age and ageing*, vol. 42 Suppl 1, p. i1–57, March 2013.
- [8] G. Collaborators *et al.*, "Global, regional, and national incidence, prevalence, and years lived with disability for 354 diseases and injuries for 195 countries and territories, 1990–2017: A systematic analysis for the global burden of disease study 2017," *The Lancet*, vol. 392, no. 10159, pp. 1789–1858, nov 2018.

- [9] U. B. of Disease Collaborators, "The State of US Health, 1990-2010: Burden of Diseases, Injuries, and Risk Factors," *JAMA*, vol. 310, no. 6, pp. 591–606, 08 2013.
- [10] D. J. Gaskin and P. Richard, "The economic costs of pain in the united states," *The Journal of Pain*, vol. 13, no. 8, pp. 715–724, 2012.
- [11] H. Breivik, E. Eisenberg, and T. O'Brien, "The individual and societal burden of chronic pain in europe: the case for strategic prioritisation and action to improve knowledge and availability of appropriate care," *BMC public health*, vol. 13, pp. 1–14, 2013.
- [12] D. A. Economics, (2019) The cost of pain in australia. [Online]. Available: <https://www.deloitte.com/content/dam/assets-zone1/au/en/docs/services/economics/deloitte-au-economics-cost-pain-australia-040419.pdf>
- [13] P. Dinakar and A. M. Stillman, "Pathogenesis of pain," *Seminars in Pediatric Neurology*, vol. 23, no. 3, pp. 201–208, aug 2016.
- [14] P. Seth, R. A. Rudd, R. K. Noonan, and T. M. Haegerich, "Quantifying the epidemic of prescription opioid overdose deaths," *American Journal of Public Health*, vol. 108, no. 4, pp. 500–502, 2018.
- [15] R. Benyamin, A. M. Trescot, S. Datta, R. M. Buenaventura, R. Adlaka, N. Sehgal, S. E. Glaser, and R. Vallejo, "Opioid complications and side effects," *Pain Physician*, vol. 2s;11, no. 3;2s, pp. S105–S120, 03 2008.
- [16] S. Gkikas and M. Tsiknakis, "Automatic assessment of pain based on deep learning methods: A systematic review," *Computer Methods and Programs in Biomedicine*, vol. 231, p. 107365, 2023.
- [17] L. A. Joel, "The fifth vital sign: pain," *AJN The American Journal of Nursing*, vol. 99, no. 2, p. 9, 1999.
- [18] A. Badura, A. Masłowska, A. Myśliwiec, and E. Pietka, "Multimodal signal analysis for pain recognition in physiotherapy using wavelet scattering transform," *Sensors*, vol. 21, no. 4, 2021.
- [19] K. Vader, G. P. Bostick, L. C. Carlesso, J. Hunter, G. Mesaroli, K. Perreault, Y. Tousignant-Laflamme, S. Tupper, D. M. Walton, T. H. Wideman, and J. Miller, "The revised iasp definition of pain and accompanying notes: Considerations for the physiotherapy profession," *Physiotherapy Canada*, vol. 73, no. 2, pp. 103–106, 2021.
- [20] L. N. Straatman, M. J. Lukacs, J. Y. Lee, M. Ghodrati, E. A. Lalone, and D. M. Walton, "Are people good prognosticators of their own pain? an exploration of the relationship between sex-specific pain beliefs and clinical pain evaluation," *Musculoskeletal Science and Practice*, vol. 62, p. 102667, December 2022.
- [21] R. Fernandez Rojas, N. Brown, G. Waddington, and R. Goecke, "A systematic review of neurophysiological sensing for the assessment of acute pain," *NPJ Digital Medicine*, vol. 6, no. 1, p. 76, 2023.
- [22] J. Farmani, G. Bargshady, S. Gkikas, M. Tsiknakis, and R. Fernandez Rojas, "A crossmod-transformer deep learning framework for multi-modal pain detection through eda and ecg fusion," *Scientific Reports*, vol. 15, no. 1, p. 29467, 2025. [Online]. Available: <https://doi.org/10.1038/s41598-025-14238-y>
- [23] R. Fernandez Rojas, M. Liao, J. Romero, X. Huang, and K.-L. Ou, "Cortical network response to acupuncture and the effect of the hegu point: An fnirs study," *Sensors*, vol. 19, no. 2, 2019.
- [24] G. Bargshady, S. Aziz, S. Gkikas, M. Tsiknakis, R. Goecke, and R. Fernandez Rojas, "Pain assessment using multi-kernel-fcn-lstm and haemoglobin difference in fnirs," *ACM Trans. Comput. Healthcare*, 2025.
- [25] S. Gkikas, "A pain assessment framework based on multimodal data and deep machine learning methods," 2025, arXiv preprint arXiv:2505.05396. [Online]. Available: <https://arxiv.org/abs/2505.05396>
- [26] P. Werner, D. Lopez-Martinez, S. Walter, A. Al-Hamadi, S. Gruss, and R. Picard, "Automatic recognition methods supporting pain assessment: A survey," *IEEE Transactions on Affective Computing*, 2019.
- [27] G. D. De Sario, C. R. Haider, K. C. Maita, R. A. Torres-Guzman, O. S. Emam, F. R. Avila, J. P. Garcia, S. Borna, C. J. McLeod, C. J. Bruce, R. E. Carter, and A. J. Forte, "Using ai to detect pain through facial expressions: A review," *Bioengineering*, vol. 10, no. 5, 2023.
- [28] M. M. Al Qudah, A. S. Mohamed, and S. L. Lutfi, "Affective state recognition using thermal-based imaging: A survey," *Computer Systems Science & Engineering*, vol. 37, no. 1, 2021.
- [29] S. Ioannou, V. Gallese, and A. Merla, "Thermal infrared imaging in psychophysiology: Potentialities and limits," *Psychophysiology*, vol. 51, no. 10, pp. 951–963, 2014.
- [30] S. Jarlier, D. Grandjean, S. Delplanque, K. N'Diaye, I. Cayeux, M. I. Velazco, D. Sander, P. Vuilleumier, and K. R. Scherer, "Thermal analysis of facial muscles contractions," *IEEE Transactions on Affective Computing*, vol. 2, no. 1, pp. 2–9, 2011.
- [31] A. Merla, L. Di Donato, P. Rossini, and G. Romani, "Emotion detection through functional infrared imaging: preliminary results," *Biomedizinische Technik*, vol. 48, no. 2, pp. 284–286, 2004.
- [32] V. K. Erel and H. S. Özkan, "Thermal camera as a pain monitor," *Journal of Pain Research*, vol. 10, pp. 2827–2832, 2017.
- [33] M. A. Haque, R. B. Bautista, F. Noroozi, K. Kulkarni, C. B. Laursen, R. Irani, M. Bellantonio, S. Escalera, G. Anbarjafari, K. Nasrollahi, O. K. Andersen, E. G. Spaich, and T. B. Moeslund, "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)*, 2018, pp. 250–257.
- [34] S. Gkikas and M. Tsiknakis, "Synthetic thermal and rgb videos for automatic pain assessment utilizing a vision-mlp architecture," in *2024 12th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, 2024, pp. 4–12.
- [35] A. Savran, R. Gur, and R. Verma, "Automatic detection of emotion valence on faces using consumer depth cameras," in *2013 IEEE International Conference on Computer Vision Workshops*, 2013, pp. 75–82.
- [36] L. Ballihi, A. Lablack, B. B. Amor, I. M. Bilasco, and M. Daoudi, "Positive/negative emotion detection from rgb-d upper body images," in *International Workshop on Face and Facial Expression Recognition from Real World Videos*. Springer, 2014, pp. 109–120.
- [37] M. Imamura, A. Tashiro, S. Kumano, and K. Otsuka, "Synergistic functional spectrum analysis: A framework for exploring the multifunctional interplay among multimodal nonverbal behaviours in conversations," *IEEE Transactions on Affective Computing*, vol. 16, no. 2, pp. 1056–1073, 2025.
- [38] M. Szwoch and P. Pieniazek, "Facial emotion recognition using depth data," in *2015 8th International Conference on Human System Interaction (HSI)*, 2015, pp. 271–277.
- [39] J. Li, Z. Dong, S. Lu, S.-J. Wang, W.-J. Yan, Y. Ma, Y. Liu, C. Huang, and X. Fu, "Cas(me)3: A third generation facial spontaneous micro-expression database with depth information and high ecological validity," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 3, pp. 2782–2800, 2023.
- [40] G. Kalliatakis, A. Stergiou, and N. Vidakis, "Conceiving human interaction by visualising depth data of head pose changes and emotion recognition via facial expressions," *Computers*, vol. 6, no. 3, 2017.
- [41] P. Werner, A. Al-Hamadi, R. Niese, S. Walter, S. Gruss, and H. C. Traue, "Automatic pain recognition from video and biomedical signals," Institute of Electrical and Electronics Engineers Inc., 2014, pp. 4582–4587.
- [42] R. Bommasani, D. A. Hudson, E. Adeli, R. Altman, S. Arora, S. von Arx, M. S. Bernstein, J. Bohg, A. Bosselut, E. Brunskill *et al.*, "On the opportunities and risks of foundation models," *arXiv preprint arXiv:2108.07258*, 2021.
- [43] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, "Segment anything," in *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023, pp. 3992–4003.
- [44] J. Ma, Y. He, F. Li, L. Han, C. You, and B. Wang, "Segment anything in medical images," *Nature Communications*, vol. 15, no. 1, p. 654, 2024.
- [45] J. Wu, W. Ji, Y. Liu, H. Fu, M. Xu, Y. Xu, and Y. Jin, "Medical sam adapter: Adapting segment anything model for medical image segmentation," *arXiv preprint arXiv:2304.12620*, 2023.
- [46] S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J. T. Springenberg *et al.*, "A generalist agent," *arXiv preprint arXiv:2205.06175*, 2022.
- [47] M. Awais, M. Naseer, S. Khan, R. M. Anwer, H. Cholakkal, M. Shah, M.-H. Yang, and F. S. Khan, "Foundational models defining a new era in vision: A survey and outlook," *arXiv preprint arXiv:2307.13721*, 2023.
- [48] P. Rodriguez, G. Cucurull, J. González, 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*, vol. 52, no. 5, pp. 3314–3324, 2022.
- [49] S. Gkikas, N. S. Tachos, S. Andreadis, V. C. Pezoulas, D. Zaridis, G. Gkois, A. Matonaki, T. G. Stavropoulos, and D. I. Fotiadis, "Multi-modal automatic assessment of acute pain through facial videos and heart rate signals utilizing transformer-based architectures," *Frontiers in Pain Research*, vol. 5, 2024.
- [50] S. Patania, G. Boccignone, S. Buršić, A. D'Amelio, and R. Lanzarotti, "Deep graph neural network for video-based facial pain expression assessment," ser. SAC '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 585–591.
- [51] S. Gkikas and M. Tsiknakis, "A full transformer-based framework for automatic pain estimation using videos," in *2023 45th Annual*

- International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2023, pp. 1–6.
- [52] R. Yang, Z. Guan, Z. Yu, X. Feng, J. Peng, and G. Zhao, “Non-contact pain recognition from video sequences with remote physiological measurements prediction,” in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Z.-H. Zhou, Ed. International Joint Conferences on Artificial Intelligence Organization, 8 2021, pp. 1231–1237.
- [53] D. Huang, X. Feng, H. Zhang, Z. Yu, J. Peng, G. Zhao, and Z. Xia, “Spatio-temporal pain estimation network with measuring pseudo heart rate gain,” *IEEE Transactions on Multimedia*, vol. 24, pp. 3300–3313, 2022.
- [54] P. Thiam, P. Bellmann, H. A. Kestler, and F. Schwenker, “Exploring deep physiological models for nociceptive pain recognition,” *Sensors*, vol. 19, p. 4503, 10 2019.
- [55] M. S. Patil and H. D. Patil, “Logistic regression based model for pain intensity level detection from biomedical signal,” *International Research Journal of Multidisciplinary Scope*, 2024.
- [56] M. Kächele, P. Thiam, M. Amirian, P. Werner, S. Walter, F. Schwenker, and G. Palm, “Multimodal data fusion for person-independent, continuous estimation of pain intensity,” in *Engineering Applications of Neural Networks*, L. Iliadis and C. Jayne, Eds. Cham: Springer International Publishing, 2015, pp. 275–285.
- [57] M. Kächele, P. Thiam, M. Amirian, F. Schwenker, and G. Palm, “Methods for person-centered continuous pain intensity assessment from bio-physiological channels,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 10, no. 5, pp. 854–864, 2016.
- [58] R. Zhi and J. Yu, “Multi-modal fusion based automatic pain assessment,” in *2019 IEEE 8th Joint International Information Technology and Artificial Intelligence Conference (ITAIC)*, 2019, pp. 1378–1382.
- [59] M. Kächele, P. Werner, A. Al-Hamadi, G. Palm, S. Walter, and F. Schwenker, “Bio-visual fusion for person-independent recognition of pain intensity,” in *Multiple Classifier Systems*. Springer International Publishing, 2015, pp. 220–230.
- [60] K. Han, A. Xiao, E. Wu, J. Guo, C. XU, and Y. Wang, “Transformer in transformer,” in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 15 908–15 919.
- [61] Y. Tang, K. Han, J. Guo, C. Xu, Y. Li, C. Xu, and Y. Wang, “An image patch is a wave: Phase-aware vision mlp,” in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 10 925–10 934.
- [62] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, “Swin transformer: Hierarchical vision transformer using shifted windows,” in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 9992–10 002.
- [63] J. Guibas, M. Mardani, Z. Li, A. Tao, A. Anandkumar, and B. Catanzaro, “Efficient token mixing for transformers via adaptive fourier neural operators,” in *International Conference on Learning Representations*, 2021.
- [64] A. Jaegle, F. Gimeno, A. Brock, O. Vinyals, A. Zisserman, and J. Carreira, “Perceiver: General perception with iterative attention,” in *International conference on machine learning*. PMLR, 2021, pp. 4651–4664.
- [65] L. Yang, B. Kang, Z. Huang, X. Xu, J. Feng, and H. Zhao, “Depth anything: Unleashing the power of large-scale unlabeled data,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 10 371–10 381.
- [66] S. Gkikas, I. Kyprakis, and M. Tsiknakis, “Tiny-biomoe: a lightweight embedding model for biosignal analysis,” 2025.
- [67] S. Gkikas and M. Tsiknakis, “Twins-painvit: Towards a modality-agnostic vision transformer framework for multimodal automatic pain assessment using facial videos and fnirs,” in *2024 12th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, 2024, pp. 13–21.
- [68] J. Li, Y. Zhao, Y. Liu, H. Zhang, P. Shan, Y. Wu, W. Wang, and Y. Wang, “Sparse emotion dictionary and cwt spectrogram fusion with multi-head self-attention for depression recognition in parkinson’s disease patients,” *IEEE Transactions on Affective Computing*, vol. 16, no. 2, pp. 1159–1176, 2025.
- [69] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, “Vggface2: A dataset for recognising faces across pose and age,” in *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*. IEEE, 2018, pp. 67–74.
- [70] G. Bae, M. de La Gorce, T. Baltrušaitis, C. Hewitt, D. Chen, J. Valentin, R. Cipolla, and J. Shen, “Digiface-1m: 1 million digital face images for face recognition,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, January 2023, pp. 3526–3535.
- [71] A. Mollahosseini, B. Hasani, and M. H. Mahoor, “Affectnet: A database for facial expression, valence, and arousal computing in the wild,” *IEEE Transactions on Affective Computing*, vol. 10, no. 1, pp. 18–31, 2019.
- [72] S. Li, W. Deng, and J. Du, “Reliable crowdsourcing and deep locality-preserving learning for expression recognition in the wild,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [73] M. Abdrakhmanova, A. Kuzdeuov, S. Jarju, Y. Khassanov, M. Lewis, and H. A. Varol, “Speakingfaces: A large-scale multimodal dataset of voice commands with visual and thermal video streams,” *Sensors*, vol. 21, no. 10, 2021.
- [74] F. Boutros, M. Huber, P. Siebke, T. Rieber, and N. Damer, “Sface: Privacy-friendly and accurate face recognition using synthetic data,” in *2022 IEEE International Joint Conference on Biometrics (IJCB)*, 2022, pp. 1–11.
- [75] D. Yi, Z. Lei, S. Liao, and S. Z. Li, “Learning face representation from scratch,” *arXiv preprint arXiv:1411.7923*, 2014.
- [76] S. Du, Y. Tao, and A. M. Martinez, “Compound facial expressions of emotion,” *Proceedings of the National Academy of Sciences*, vol. 111, no. 15, pp. E1454–E1462, 2014.
- [77] J. M. Ford, V. A. Palzes, B. J. Roach, and D. H. Mathalon, “Did I Do That? Abnormal Predictive Processes in Schizophrenia When Button Pressing to Deliver a Tone,” *Schizophrenia Bulletin*, vol. 40, no. 4, pp. 804–812, 07 2013.
- [78] D. Gaddy and D. Klein, “Digital voicing of silent speech,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics, 2020, pp. 5521–5530.
- [79] M. Kachuee, S. Fazeli, and M. Sarrafzadeh, “Ecg heartbeat classification: A deep transferable representation,” in *2018 IEEE International Conference on Healthcare Informatics (ICHI)*, 2018, pp. 443–444.
- [80] S. G. Müller and F. Hutter, “Trivialaugment: Tuning-free yet state-of-the-art data augmentation,” in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 754–762.
- [81] D. Hendrycks, N. Mu, E. D. Cubuk, B. Zoph, J. Gilmer, and B. Lakshminarayanan, “Augmix: A simple data processing method to improve robustness and uncertainty,” *arXiv preprint arXiv:1912.02781*, 2019.
- [82] G. Larsson, M. Maire, and G. Shakhnarovich, “Fractalnet: Ultra-deep neural networks without residuals,” *arXiv preprint arXiv:1605.07648*, 2016.
- [83] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the inception architecture for computer vision,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [84] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov, “Improving neural networks by preventing co-adaptation of feature detectors,” *arXiv preprint arXiv:1207.0580*, 2012.
- [85] S. Walter, S. Gruss, H. Ehleiter, J. Tan, H. C. Traue, S. Crawcour, P. Werner, A. Al-Hamadi, A. O. Andrade, and G. M. D. Silva, “The biovid heat pain database: Data for the advancement and systematic validation of an automated pain recognition,” 2013, pp. 128–131.
- [86] R. Fernandez Rojas, N. Hirachan, C. Joseph, B. Seymour, and R. Goecke, “The ai4pain grand challenge 2024: Advancing pain assessment with multimodal fnirs and facial video analysis,” in *2024 12th International Conference on Affective Computing and Intelligent Interaction*. IEEE, 2024.
- [87] P. Werner, A. Al-Hamadi, and S. 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)*, 2017, pp. 176–180.
- [88] R. Zhi and M. Wan, “Dynamic facial expression feature learning based on sparse rnn,” in *Proceedings of 2019 IEEE 8th Joint International Information Technology and Artificial Intelligence Conference, ITAIC 2019*. Institute of Electrical and Electronics Engineers Inc., may 2019, pp. 1373–1377.
- [89] P. Thiam, H. A. Kestler, and F. Schwenker, “Two-stream attention network for pain recognition from video sequences,” *Sensors (Switzerland)*, vol. 20, no. 3, p. 839, feb 2020.
- [90] M. Tavakolian, M. Bordallo Lopez, and L. Liu, “Self-supervised pain intensity estimation from facial videos via statistical spatiotemporal distillation,” *Pattern Recognition Letters*, vol. 140, pp. 26–33, 2020.
- [91] F. Alhamdoosh, P. Pala, and S. Berretti, “Pain level estimation from videos by analyzing the dynamics of facial landmarks with a spatio-temporal graph neural network,” 2025.

- [92] S. Gkikas, C. Chatzaki, and M. Tsiknakis, "Multi-task neural networks for pain intensity estimation using electrocardiogram and demographic factors," in *Information and Communication Technologies for Ageing Well and e-Health*. Springer Nature Switzerland, 2023, pp. 324–337.
- [93] Z. Lu, B. Ozek, and S. Kamarthi, "Transformer encoder with multiscale deep learning for pain classification using physiological signals," *Frontiers in Physiology*, vol. 14, 2023.
- [94] K. N. Phan, N. K. Iyortsuun, S. Pant, H.-J. Yang, and S.-H. Kim, "Pain recognition with physiological signals using multi-level context information," *IEEE Access*, vol. 11, pp. 20 114–20 127, 2023.
- [95] M. Jiang, Y. Li, J. He, Y. Yang, H. Xie, and X. Chen, "Physiological time-series fusion with hybrid attention for adaptive recognition of pain," *IEEE Journal of Biomedical and Health Informatics*, pp. 1–9, 2024.
- [96] P. Prajod, D. Schiller, D. W. Don, and E. André, "Faces of experimental pain: Transferability of deep-learned heat pain features to electrical pain*," in *2024 12th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*, 2024, pp. 31–38.
- [97] M.-D. Nguyen, H.-J. Yang, S.-H. Kim, J.-E. Shin, and S.-W. Kim, "Transformer with leveraged masked autoencoder for video-based pain assessment," 2024.
- [98] M. U. Khan, S. Aziz, L. Murtagh, G. Chetty, R. Goecke, and R. Fernandez Rojas, "Empirically transformed energy patterns: A novel approach for capturing fnirs signal dynamics in pain assessment," *Computers in Biology and Medicine*, vol. 192, p. 110300, 2025.
- [99] J. Vianto, A. Divakaran, H. Yang, S. Yeom, S. Kim, S. Kim, and J. Shin, "Multimodal model for automated pain assessment: Leveraging video and fnirs," *Applied Sciences*, vol. 15, no. 9, 2025.
- [100] S. Gkikas, I. Kyprakis, and M. Tsiknakis, "Multi-representation diagrams for pain recognition: Integrating various electrodermal activity signals into a single image," 2025.
- [101] D. Huang, Z. Xia, L. Li, K. Wang, and X. Feng, "Pain-awareness multistream convolutional neural network for pain estimation," *Journal of Electronic Imaging*, vol. 28, no. 04, p. 1, 2019.
- [102] P. Werner, A. Al-Hamadi, K. Limbrecht-Ecklundt, S. Walter, S. Gruss, and H. C. Traue, "Automatic pain assessment with facial activity descriptors," *IEEE Transactions on Affective Computing*, vol. 8, no. 3, pp. 286–299, 2016.
- [103] D. Huang, Z. Xia, J. Mwesigye, and X. Feng, "Pain-attentive network: a deep spatio-temporal attention model for pain estimation," *Multimedia Tools and Applications*, vol. 79, no. 37–38, pp. 28 329–28 354, 2020.
- [104] D. Lopez-Martinez and R. Picard, "Continuous pain intensity estimation from autonomic signals with recurrent neural networks," vol. 2018, 2018, pp. 5624–5627.
- [105] E. Pavlidou and M. Tsiknakis, "Multimodal pain assessment based on physiological biosignals: The impact of demographic factors on perception and sensitivity," in *Proceedings of the 11th International Conference on Information and Communication Technologies for Ageing Well and e-Health - ICT4AWE, INSTICC*. SciTePress, 2025, pp. 320–329.
- [106] S. Gkikas., C. Chatzaki., E. Pavlidou., F. Verigou., K. Kalkanis., and M. Tsiknakis., "Automatic pain intensity estimation based on electrocardiogram and demographic factors." SciTePress, 2022, pp. 155–162.
- [107] M. S. Patil and H. D. Patil, "Ensemble neural networks for multimodal acute pain intensity evaluation using video and physiological signals," *Journal of Computational Analysis and Applications (JoCAAA)*, vol. 33, no. 05, p. 779–791, Sep. 2024.
- [108] X. Ji, T. Zhao, W. Li, and A. Zomaya, "Automatic pain assessment with ultra-short electrodermal activity signal," in *Proceedings of the 38th ACM/SIGAPP Symposium on Applied Computing*, ser. SAC '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 618–625.
- [109] F. Pouroumran, S. Radhakrishnan, and S. Kamarthi, "Exploration of physiological sensors, features, and machine learning models for pain intensity estimation," *PLOS ONE*, vol. 16, no. 7, pp. 1–17, 07 2021.
- [110] P. Gouverneur, F. Li, W. M. Adamczyk, T. M. Szikszay, K. Luedtke, and M. Grzegorzec, "Comparison of feature extraction methods for physiological signals for heat-based pain recognition," *Sensors*, vol. 21, no. 14, 2021.
- [111] J. Li, J. Luo, Y. Wang, Y. Jiang, X. Chen, and Y. Quan, "Automatic pain assessment based on physiological signals: Application of multi-scale networks and cross-attention cross-attention," in *Proceedings of the 2024 13th International Conference on Bioinformatics and Biomedical Science*, ser. ICBBS '24. New York, NY, USA: Association for Computing Machinery, 2025, p. 113–122.
- [112] R. Wang, K. Xu, H. Feng, and W. Chen, "Hybrid rnn-ann based deep physiological network for pain recognition," in *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, vol. 2020-July. Institute of Electrical and Electronics Engineers Inc., jul 2020, pp. 5584–5587.
- [113] P. Thiam, H. A. Kestler, and F. Schwenker, "Multimodal deep denoising convolutional autoencoders for pain intensity classification based on physiological signals," in *ICPRAM 2020 - Proceedings of the 9th International Conference on Pattern Recognition Applications and Methods*. SCITEPRESS - Science and Technology Publications, 2020, pp. 289–296.
- [114] P. Thiam, H. Hihn, D. A. Braun, H. A. Kestler, and F. Schwenker, "Multi-modal pain intensity assessment based on physiological signals: A deep learning perspective," *Frontiers in Physiology*, vol. 12, 2021.
- [115] M. Jiang, R. Rosio, S. Salanterä, A. M. Rahmani, P. Liljeberg, D. S. da Silva, V. H. C. de Albuquerque, and W. Wu, "Personalized and adaptive neural networks for pain detection from multi-modal physiological features," *Expert Systems with Applications*, vol. 235, p. 121082, 2024.
- [116] P. Bellmann and F. Schwenker, "Automated pain assessment: Is it useful to combine person-specific data samples?" in *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2020, pp. 1588–1593.
- [117] S. Gkikas, I. Kyprakis, and M. Tsiknakis, "Efficient pain recognition via respiration signals: A single cross-attention transformer multi-window fusion pipeline," 2025.

Stefanos Gkikas received his M.Sc. in Computer Vision from the University of Burgundy, France, in 2020. In 2025, he earned his Ph.D. from the School of Electrical and Computer Engineering at the Hellenic Mediterranean University. His doctoral research focused on pain recognition, with the dissertation titled "A Pain Assessment Framework Based on Multimodal Data and Deep Machine Learning Methods". Throughout his Ph.D., he specialized in the analysis of videos and biosignals. He is now a postdoctoral researcher at the Computational Biomedicine Laboratory of FORTH/ICS. His research interests include affective computing, health monitoring systems, and computational physiology.

Raul Fernandez Rojas received his Ph.D. from the University of Canberra, Australia, in 2018, specialising in biomarker identification through computational methods and cognitive engineering. From 2018 to 2019, he served as a Postdoctoral Research Associate at the University of New South Wales, within the Australian Defence Force Academy. He is now an Associate Professor in the Faculty of Science and Technology at the University of Canberra. His research interests encompass signal processing, embedded systems design, affective computing, and multimodal sensing.

Manolis Tsiknakis received his M.Sc (1983) and Ph.D. (1989) degrees from the University of Bradford, UK. He performed postdoctoral training at the University of Bradford (1990-1991) and the Institute of Computer Science at FORTH (1992). Since 2012, he has been a Professor of Biomedical Informatics and eHealth at the Department of Electrical and Computer Engineering at the Hellenic Mediterranean University and a visiting Professor at the Computational Biomedicine Laboratory of FORTH/ICS. His main areas of expertise include approaches for semantic health data integration and interoperability of health information systems, affective computing and its application in developing smart eHealth solutions, and service platforms for pervasive eHealth and mHealth services.