

# Personalized Automatic Estimation of Self-reported Pain Intensity from Facial Expressions

Daniel Lopez Martinez<sup>1,2</sup>, Ognjen (Oggi) Rudovic<sup>2</sup> and Rosalind Picard<sup>2</sup>

<sup>1</sup> Harvard-MIT Division of Health Sciences and Technology, Massachusetts Institute of Technology

<sup>2</sup> MIT Media Lab, Massachusetts Institute of Technology, Cambridge, MA, USA

{dlmocdm, orudovic, picard}@media.mit.edu

## Abstract

*Pain is a personal, subjective experience that is commonly evaluated through visual analog scales (VAS). While this is often convenient and useful, automatic pain detection systems can reduce pain score acquisition efforts in large-scale studies by estimating it directly from the participants' facial expressions. In this paper, we propose a novel two-stage learning approach for VAS estimation: first, our algorithm employs Recurrent Neural Networks (RNNs) to automatically estimate Prkachin and Solomon Pain Intensity (PSPI) levels from face images. The estimated scores are then fed into the personalized Hidden Conditional Random Fields (HCRFs), used to estimate the VAS, provided by each person. Personalization of the model is performed using a newly introduced facial expressiveness score, unique for each person. To the best of our knowledge, this is the first approach to automatically estimate VAS from face images. We show the benefits of the proposed personalized over traditional non-personalized approach on a benchmark dataset for pain analysis from face images.*

## 1. Introduction

Pain is a distressing experience associated with actual or potential tissue damage with sensory, emotional, cognitive and social components [4]. Accurate, valid, and reliable measurement of pain is essential for the diagnosis and management of pain in the hospital and at home, and to ensure accurate evaluation of the relative effectiveness of different therapies. In order to develop automatic methods of objectively quantifying an individual's experience of pain, several physiologic variables have been measured, such as skin conductance and heart rate [43]. In general, however, these markers do not correlate strongly enough with pain to warrant their use as a surrogate measure of pain [50]. In the absence of a valid and reliable objective, physiologic marker

of pain, patient's self-report provides the most valid measure of this fundamentally subjective experience [28]. Various pain rating scales have been developed to capture patient's self-report of pain intensity [16, 14]. The Visual Analogue Scale (VAS), Numerical Rating Scale (NRS), Verbal Rating Scale (VRS), and Faces Pain Scale-Revised (FPS-R) are among the most commonly used scales. While evidence supports the reliability and validity of each of these measures across many populations [11], each measure has strengths and weaknesses that make them more appropriate for different applications [49]. In the research setting, for example, VAS is usually preferred since it is statistically the most robust as it can provide ratio level data [49, 32], allowing parametric tests to be used in statistical analysis. Furthermore, VAS is the most common pain intensity scale in clinical trials [42, 10, 31, 19, 2, 20]. When using the VAS, patients are asked to mark the pain that they are experiencing on a 10cm-long horizontal line labeled "no pain" on the far left and "worst pain ever" on the far right. Pain intensity is determined by the length of the line as measured from the left-hand side to the point marked [29, 16]. The following cut points on the pain VAS are recommended: no pain (0-0.4 cm), mild pain (0.5-4.4 cm), moderate pain (4.5-7.4 cm), and severe pain (7.5-10 cm) [16]. Unfortunately, VAS is highly impractical and inefficient to measure and lacks utility in key patient populations, as it requires patients to have intact fine motor skills that may be limited by illness or injury.

To circumvent the challenge of acquiring patients' self-reported pain scores, facial indicators of pain [48] have been used in automatic pain detection systems. However, all existing studies have focused on the estimation of the Prkachin and Solomon Pain Intensity (PSPI) score, a metric that measures pain as a linear combination of the intensities of facial action units (see Sec.2) [33], instead of VAS. While the facial response to pain certainly consists of a core set of facial actions, the relationship between facial cues and VAS

scores is not clear [33]. For example, it has been shown that men and women show different levels of pain expression for the same stimulus [41]. In fact, many studies have found low correlations between facial expression and self-reported pain [33]. Therefore, an approach to automatically estimate self-reported pain intensities from facial cues should also account for individual’s differences in facial expressiveness of pain.

This work addresses the challenges of estimating VAS scores from facial cues. To this end, we propose a hierarchical learning framework that exploits the modeling power of sequence classification models. Specifically, we first employ a bidirectional long short-term memory recurrent neural net (LSTM-RNN) [18] to estimate PSPI from facial landmarks extracted from face images. Then, PSPI is used as input to the hidden conditional random field (HCRF) [34, 45] to estimate VAS of target person. The key to our approach is the personalization of the target classifier via the newly introduced individual facial expressiveness score (I-FES). Specifically, I-FES quantifies the relative disagreement between an externally observed pain intensity (OPI) rating and the patient’s self-reported VAS. By using I-FES to augment the PSPI input to HCRF, we account for person-specific biases in the VAS rating of pain levels. To the best of our knowledge, this is the first approach to automatically estimate VAS from face images. We show on the UNBC-McMaster Shoulder Pain Expression Archive Database [26] that the proposed personalized approach for automatic VAS estimation outperforms largely traditional (unpersonalized) approaches.

## 2. Related Work

Although there is much research in the automated recognition of affect from facial expression (for surveys see: [52, 5]), until recently only a handful of works have focused on automated pain estimation. Due to rapid advances in computer vision and also the recent release of the UNBC-McMaster dataset [26], pain analysis from face images have seen significant advances. Specifically, this dataset provides videos with each frame coded in terms of Prkachin and Solomon Pain Intensity (PSPI) [33] score, defined on an ordinal scale 0-15. This is considered to be an objective pain score, in contrast to subjective pain ratings such as VAS. Despite the fact that VAS is still the most commonly accepted pain score in clinical settings, all existing automatic methods for pain estimation from pain images focused on prediction of PSPI scores. We outline below some of the recently published works.

Face shape features have been used in artificial neural networks to classify images of persons’ faces in a typical mood versus a pain inducing task [30]. Likewise, [25, 3] used Active Appearance Models (AAM) - based features combined with Support Vector Machine (SVM) classifiers

to classify pain versus no pain images. On the other hand, instead of treating pain as binary (pain - no pain), [15] attempted estimation of pain intensity on a 4 level scale, using one-versus-all SVM classifiers. [22] performed estimation of the full pain range (0-15), using a framework of Relevance Vector Regression models (RVR). Recently, [38] have directly exploited the temporal axis information by using long short-term memory (LSTM) neural networks. In their approach, raw images are fed to a convolutional neural network that extracts frame features. These features are subsequently used by the LSTM to predict a PSPI score for each frame. Therefore, this deep learning approach is able to leverage the temporal information associated with facial expressions. [8] also used a deep learning approach to perform automatic continuous PSPI prediction using a combination of hand-crafted and deep-learned features.

Due to subtle changes in facial expressions and inter-subject variability, per frame PSPI estimation is very challenging. Hence, some of the works attempted pain intensity estimation per sequence. For instance, [40] proposed a multi-instance learning framework to identify the most expressive (in terms of pain) segments within image sequences. Likewise, [37] proposed a HCRF framework for semi-supervised learning of the peak of the pain in image sequences, using weakly labeled pain images. However, these per sequence labels are derived using heuristics, and do not relate to established pain ratings, such as VAS scores.

Aside from the UNBC-McMaster dataset [26], there have been a variety of pain recognition studies based on other datasets (e.g., see [46, 47]). Note also that pain detection has been attempted from other modalities including upper body movements combined with face [13, 21], physiological signals such as skin conductance and heart rate [43], and brain hemodynamic responses using NIRS [1, 51] and fMRI [27, 44]. Nevertheless, none of these works attempted automatic estimation of VAS, and in a personalized manner.

## 3. Personalized RNN-HCRFs: The Model

**Notation.** We consider a sequence classification problem, where we are given  $N_i$  image sequences of person  $i = 1, \dots, L$ , and  $L$  is the number of target persons. The sequences of each person are annotated in terms of VAS as  $\mathcal{V} = \{V_1, \dots, V_L\}$ , where  $V_i = \{v_i^1, \dots, v_i^{N_i}\}$  and individual per-sequence VAS scores are  $v_i \in \{0, \dots, 10\}$ . Likewise, for observed pain intensity (OPI) scores we have:  $\mathcal{O} = \{O_1, \dots, O_L\}$ ,  $O_i = \{o_i^1, \dots, o_i^{N_i}\}$  and  $o_i \in \{0, \dots, 5\}$ . Furthermore, each sequence is represented by a set of input features (in our case, locations of facial landmarks – see Fig.3C) and the corresponding PSPI scores, given as pairs  $(X, S)$ , respectively. Here,  $X = \{x_1, \dots, x_T\}$  and  $S = \{s_1, \dots, s_T\}$ , and duration  $T$  can vary for each sequence. The pairs  $\{x_t, s_t\}_{t=1}^T$  are assumed to be i.i.d. sam-

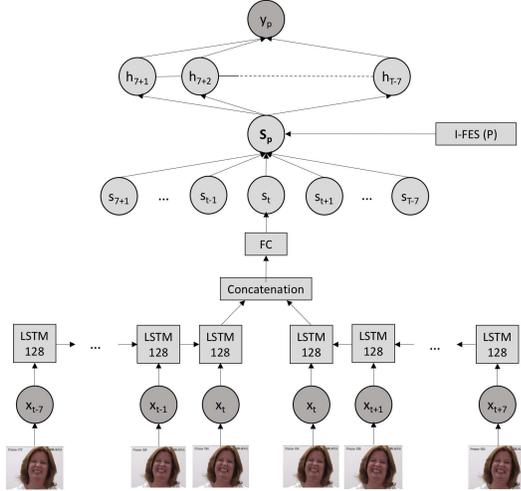


Figure 1: Graphical representation of the proposed pRNN-HCRF for estimation of VAS from facial landmarks ( $x$ ) at each frame in a sequence. The bidirectional LSTM-RNN and HCRF layers are connected via the estimated PSPI ( $s$ ) and individual facial expressiveness ratio (I-FES). The latter is used to personalize the model parameters.

ples from an underlying but unknown distribution, and  $x_t \in \mathcal{R}^D$ , where  $D$  is the input dimension, and  $s_t \in \{0, \dots, 15\}$ . We exploit all this information (as described below) to define and learn our model.

We exploit the inherent hierarchical nature of the target problem to define our model for personalized VAS score prediction from image sequences. To this end, we first leverage the modeling power of Long Short-Term Memory Recurrent Neural Networks (LSTM-RNNs)[18] to estimate (in an unpersonalized manner) the PSPI scores from input features  $X$  (Sec.3.1). These are then endowed with the newly introduced Individual Facial Expressiveness Score (I-FES) denoted by  $p_i$  (Sec.3.2), and used as input to the top layer of our model, based on Hidden Conditional Random Fields (HCRF) [34, 45], that performs sequence classification in terms of target VAS scores  $v_i$ . To account for the "personalized dynamics", we use the notion of latent states of HCRF, denoted as  $H = \{h_1, \dots, h_T\}$  (Sec.3.3). The graph model of the proposed personalized RNN-HCRF (pRNN-HCRF) is depicted in Fig.1.

### 3.1. RNNs for PSPI Estimation

The first step in our approach is to obtain an objective estimate of the pain intensity, as encoded by the (manually coded) PSPI [26]. To this end, various approaches based on static classifiers/regressors (e.g., SVMs, SVRs) have been proposed [22]. To capture temporal information within an image sequence of a person with varying pain intensity levels, dynamic classification methods based on

Hidden Markov Models (HMMs) or Conditional Random Fields, and their extensions are typically used [35]. However, these methods usually assume first order Markov dependency, thus failing to capture long-term dependencies in image sequences.

We use the long short-term memory recurrent neural nets (LSTM-RNNs) [18], as they have recently shown great success in sequence learning tasks such as speech recognition [12] and sleep/wake classification [6]. In this work, we employ a bidirectional LSTM-RNNs architecture [6] (shown in Fig.1) to estimate the PSPI values from input features  $X$  (i.e., a sequence of facial landmarks). While the traditional RNNs are unable to learn temporal dependencies longer than a few time steps due to the vanishing gradient problem [17], LSTM-RNNs overcome this by introducing recurrently connected memory blocks instead of traditional neural network nodes. In an LSTM-RNN, each block contains one or more recurrently connected memory cells, along with three multiplicative gate units: the input, output, and forget gates [18]. Their role is to give access to long range context information to each block (in our case, the relationships between facial expressions across a time-window). This, in turn, allows the network to store and retrieve information over longer time windows.

To train our LSTM-RNN<sup>1</sup> for PSPI estimation, we consider the standard regression setting: each PSPI score at time  $t$  in an image sequence is treated as a continuous value  $s_t$  (PSPI  $\in \{0, \dots, 15\}$  scaled to  $s_t = (0 - 1)$ ). The inputs are the normalized locations of 66 facial landmarks ( $x_t$ ), provided by the database creators and depicted in Fig.3. Our model then tries to estimate  $s_t$  for each frame  $t$ . To do so, the LSTM units in our RNN consider windows of 15 frames:  $\{x_{t-7}, \dots, x_t, \dots, x_{t+7}\}$  (see Fig.1). This has been shown to be sufficient to capture local (dynamic) changes in facial expressions due to the pain onset/offset phases [38]. To select the number of units in the LSTM blocks, we performed a validation on the subset of training data. The number of units was set to 128. The output of the two LSTM-RNNs is concatenated and fed to a fully-connected layer with rectified linear units, which produces a single score for each 15-frame window. The neural network was trained using root mean square propagation with mean squared error loss. The whole algorithm was implemented using deep learning frameworks Theano 0.8.2 and Keras 1.2.1.

### 3.2. Individual Facial Expressiveness Score (I-FES)

Most work on automatic pain estimation from faces has focused on estimating the PSPI, a metric that measures pain as a linear combination of the intensities of facial action

<sup>1</sup>Further in the text, we refer to this model as RNN.

units (AUs):

$$\text{PSPI} = \text{AU4} + \max(\text{AU6}, \text{AU7}) + \max(\text{AU9}, \text{AU10}) + \text{AU43}, \quad (1)$$

where the included AU's (see Fig.3B) are defined on a 0-5 ordinal scale, and correspond to brow lowering (AU4), orbital tightening (AU6&7), eye closure (AU43), nose wrinkling, and lip raise (AU9&10) [7]. These are derived by manual coding of each face image using the Facial Action Coding System (FACS) [9].

Estimating VAS poses a big challenge since facial expressions and self-reports are typically poorly correlated due to person-specific biases in facial expressions of pain [33]. To account for these individual differences (e.g., see Fig.2), we introduce the Individual Facial Expressiveness Score (I-FES), which captures the ratio of OPI ( $o$ ), obtained by independent observers, to VAS ( $v$ ). Specifically, for person  $i$ , we define I-FES ( $p_i$ ), as:

$$p_i = \begin{cases} \frac{1}{\alpha} \sum_{k=1}^{\alpha} \frac{o_i^k + 1}{v_i^k + 1}, & \text{iff } \alpha > 0 \\ 1, & \text{otherwise} \end{cases} \quad (2)$$

where  $\alpha$  is the number of sequences used to calculate I-FES for target person. When  $\alpha = 0$ , I-FES is set to 1 (assuming the perfect agreement), whereas  $\alpha = 1, \dots, N_{L_i}$  indicates the number of sequences of person  $i$ . To avoid division by 0 we add +1 in Eq.2. Note the intuition behind the proposed I-FES: for a given self-reported VAS, the observed pain intensity (OPI), as noted by an external observer, is expected to vary significantly from person to person depending on their level of facial expressiveness. Using I-FES, we quantify these relative differences in the ratings, with the OPI rater acting as the 'reference point' for all target persons.

### 3.3. HCRF for Personalized VAS Estimation

HCRFs are a class of models for dynamic multi-class classification of sequential data [34, 45]. They are a generalization of Conditional Random Fields (CRFs) [24] proposed for modeling/decoding of the state-sequence (e.g., PSPI scores for each image frame [35]). In HCRFs, the first-order log-linear CRFs are usually employed as building blocks. This is attained by introducing a top node in the graph structure of CRFs (see Fig.1), representing the sequence label (in our case, the target VAS score), and the temporal states are treated as hidden variables in the model, and denoted by  $H = \{h_1, \dots, h_T\}$ . Formally, using the notation introduced above, HCRFs combine the score functions of  $K$  CRFs, one for each class  $v \in \{1, \dots, K\}$ , via the following score function:

$$F(v, S_p, H; \Omega) = \sum_{k=1}^K I(k = v) \cdot f(S_p, H; \theta_k). \quad (3)$$

We define the personalized HCRF features as  $S_p = [X p]$ , where  $X$  is the sequence of input features and  $p$  the I-FES

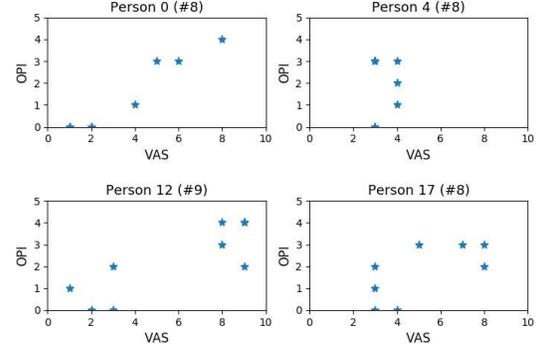


Figure 2: The VAS vs OPI scores for four different subjects (along with the number of available sequences). Note the different correlation patterns exhibited for each person: linear (top-left), no-correlation (top-right), and bi-modal (bottom row). All these pose challenges in predicting the individual VAS. For this reason, we account for these individual biases via the proposed I-FES (Eq.2) in our pRNN-HCRF.

defined in Sec.3.2. Furthermore,  $\Omega = \{\theta_k\}_{k=1}^K$  denotes the model parameters, and  $f(\cdot, \cdot; \theta_k)$  is the  $k$ -th CRF score function. We assume the linear-chain graph structure  $G = (U, E)$ , encoded by the sum of the unary (U) and edge (E) potentials:

$$\sum_{r \in U} \Psi^{(U)}(S_p, h_r) + \sum_{e=(r,z) \in E} \Psi^{(E)}(S_p, h_r, h_z). \quad (4)$$

The unary potentials are defined as the linear classifier:

$$\Psi_r^{(U)}(S_p, h_r) = \sum_{c=1}^C I(h_r = c) \cdot u_c^\top S_p, \quad (5)$$

and the edge potentials  $\Psi^{(E)}(S_p, h_r, h_s)$  are defined as:

$$\Psi^{(E)}(S_p, h_r, h_z) = [I(h_r = c \wedge h_z = l)]_{C \times C} \otimes m_k(r, z), \quad (6)$$

where  $I(\cdot)$  is the indicator function that returns 1 (0) if the argument is true (false), and  $\otimes$  is Kronecker product. The role of the edge potentials is to assure the temporal consistency of the hidden states within a sequence. The model parameters are stored as  $\theta_k = \{u_k, m_k\}$ .

Using the score function defined above, the joint conditional distribution of the class and state-sequence is defined as:

$$P(v, H|S_p) = \frac{\exp(F(v, S_p, H; \Omega))}{Z(S_p)}. \quad (7)$$

The state-sequence  $H = (h_1, \dots, h_T)$  is unknown, and it is integrated out by directly modeling the class conditional distribution:

$$P(v|S_p) = \sum_H P(v, H|S_p) = \frac{\sum_H \exp(F(v, S_p, H; \Omega))}{Z(S_p)}. \quad (8)$$

Evaluation of the class-conditional  $P(v|S_p)$  depends on the partition function  $Z(S_p) = \sum_k Z_k(S_p) = \sum_k \sum_H \exp(F(k, S_p, H; \Omega))$ , and the class-latent joint posteriors  $P(k, h_r, h_z|x) = P(h_r, h_z|S_p, k) \cdot P(k|S_p)$ . Both can be computed from independent consideration of  $K$  individual CRFs.

The parameter optimization in the HCRF is carried out by maximizing the (regularized) negative log-likelihood of the class conditional distribution in Eq.8. Formally, the regularized objective function is given by:

$$RLL(\Omega) = - \sum_{i=1}^N \log P(\mathbf{v}_i | \mathbf{S}_p^i; \Omega) + \lambda \|\Omega\|^2, \quad (9)$$

where  $N$  is the number of training sequences,  $P(\mathbf{v}_i | \mathbf{S}_p^i; \Omega)$  is defined by Eq.7, and  $\lambda$  controls the regularization penalty. To estimate the model parameters, the standard quasi-Newton (such as Limited-memory BFGS) gradient descent algorithm is used.

### 3.4. Learning and Inference

Given facial landmarks ( $X_{tr}$ ) extracted from image sequences of  $L_{tr}$  training persons (multiple sequences per person), and the corresponding PSPI scores for each frame (stored in  $S_{tr}$ ), we first focus on the learning of the RNNs, as described in Sec.3.1. This step is critical as it transforms the high-dimensional input features (i.e., facial landmarks) into a low-dimensional (1-D) representation of the ‘painful faces’. One can think of this as an efficient supervised dimensionality reduction approach, aimed at capturing the variance in the input that is directly related to the observed pain levels observed in face images. After learning the RNN model parameters ( $Q^{opt}$ ), we compute the I-FES, as denoted by  $p$ , for each person in the training data. Note that I-FES is unique to each person and is estimated from the ratio of VAS ( $v$ ) and OPI ( $o$ ) using all sequences of target persons in training set (Eq.2). Since we use all sequences of the training persons to estimate I-FES for each person (stored in  $\mathcal{P}$ ) we obtain a robust measure of the discrepancy between the subjective pain rated by the person experiencing it and an external observer’s (e.g., a caregiver in the clinical setting) score.

The actual personalization of our approach occurs in the second stage of the learning. Namely, the estimated PSPI scores ( $\tilde{S}$ ), obtained using the learned RNN model, are augmented with the I-FES ( $\mathcal{P}$ ), to obtain 2-D feature vectors for each image frame in a sequence of the target person. Such personalized feature vectors are then used as input to the HCRF model for the estimation of target VAS scores ( $v$ ) (Sec.3.3). It is important to mention that while the estimated PSPI can vary from frame to frame within

---

### Algorithm 1 Personalized RNN-HCRFs

---

**Learning:** Input  $\mathcal{D}_{tr} = \{X_i^{tr}, \mathcal{V}_i^{tr}, \mathcal{O}_i^{tr}, S_i^{tr}\}_{i=1}^{L_{tr}}$

**Step 1:** Optimize RNNs ( $Q^{opt}$ ) given  $\{S_{tr}, X_{tr}\}$

**Step 2:** Compute I-FES

**for**  $i = 1 : L_{tr}, p_i \leftarrow \{\mathcal{V}_i^{tr}, \mathcal{O}_i^{tr}\}$  **end**,  $\mathcal{P} = \{p_i\}_{i=1}^{L_{tr}}$

**Step 3:** Optimize HCRFs ( $\Omega$ )

a) estimate PSPI  $\tilde{S} \leftarrow \text{RNN}(X_{tr}; Q^{opt})$  – **Step 1**,  $S_p = \{\tilde{S}, \mathcal{P}\}$

b)  $\min_{\Omega^{opt}} - \sum_{i=1}^{L_{tr}} \sum_{j=1}^{N_i} \log P(\mathbf{v}_i^j | S_{p_i}^j; \Omega) + \lambda \|\Omega\|^2$

Output: RNNs( $\cdot; Q^{opt}$ ), HCRFs( $\cdot; \Omega^{opt}$ )

---

**Inference:** Input  $\mathcal{D}_{te} = \{X_*, p_*\}$

**Step 1:** Estimate PSPI  $S^* \leftarrow \text{RNNs}(X_*; Q^{opt})$

**Step 2:** Estimate VAS  $v^* \leftarrow \text{HCRFs}(S_p^* = \{S^*, p_*\}; \Omega^{opt})$

Output:  $S^*, v^*$

---

a sequence, the corresponding I-FES ( $p$ ) remains constant across all sequences of a target person. In this way, it encodes the person-specific biases coming from two different rating scales (VAS&OPI). Thus, by augmenting the person-agnostic PSPI scores ( $S$ ), we implicitly modulate the parameters of the HCRF model. Consequently, the model adjusts the classification boundaries of the HCRF so that it fits best the VAS scores for each target person. On the other hand, since  $p$  is computed using OPI as well, OPI acts as a reference point in estimating the person specific-biases in the model. We choose 11 hidden states in HCRF, the same as the number of VAS levels. The summary of our learning procedure is given in Alg.1. One important aspect of our model is that it separates the learning of the person-agnostic PSPI and person-specific VAS pain intensity scores. This is for the following reasons: while the first level (RNNs) focus on extracting efficient low-dimensional pain descriptor (thus, removing noise), we achieve more accurate and robust learning of target VAS scores than if the input facial landmarks are used directly in our HCRF, as confirmed by our experiments. More importantly, the contribution of the I-FES in the feature vector could easily be downplayed due to the large number of other (possibly noisy) input features (i.e., facial landmarks). We also avoid heavy parametrization of our HCRF, which can easily lead to overfitting. Lastly, by adopting the proposed two-stage approach, we allow our HCRF to focus its learning and dynamics on target person (‘personalized dynamics’), rather than on filtering of noisy data of multiple persons (as done at the RNN level).

To make inference of the VAS score ( $v^*$ ) for a (new) target person, we feed into our pRNN-HCRF the facial landmarks  $x^*$  and an estimate of the I-FES ( $p^*$ ) for that person. The output are the estimated PSPI ( $s^*$ ) and VAS ( $v^*$ ). Note that while this inference is straightforward, it requires, however, an estimate of the I-FES for the target person. In



Table 2: The performance of different methods tested for VAS (0-10) estimation. The mean and standard deviation are computed over 5 random selection of target sequences used to compute I-FES for test persons.

Algorithm	Performance	No personalization	Personalization approach	
		$\alpha = 0$	$\alpha = 1$	$\alpha = 2$
HCRF(PSPI)	MAE	2.65	2.18 (0.05)	2.12 (0.23)
	ICC	0.35	0.49 (0.03)	0.50 (0.13)
RNN-HCRF	MAE	3.67	<b>2.47 (0.18)</b>	<b>2.46 (0.23)</b>
	ICC	0.04	<b>0.36 (0.08)</b>	<b>0.34 (0.04)</b>
SVR-HCRF	MAE	5.44	2.98 (0.28)	2.83 (0.30)
	ICC	0.04	0.22 (0.06)	0.20 (0.17)
HCRF( $X$ )	MAE	<b>2.8</b>	2.88 (0.20)	2.91 (0.39)
	ICC	<b>0.19</b>	0.18 (0.07)	0.22 (0.14)

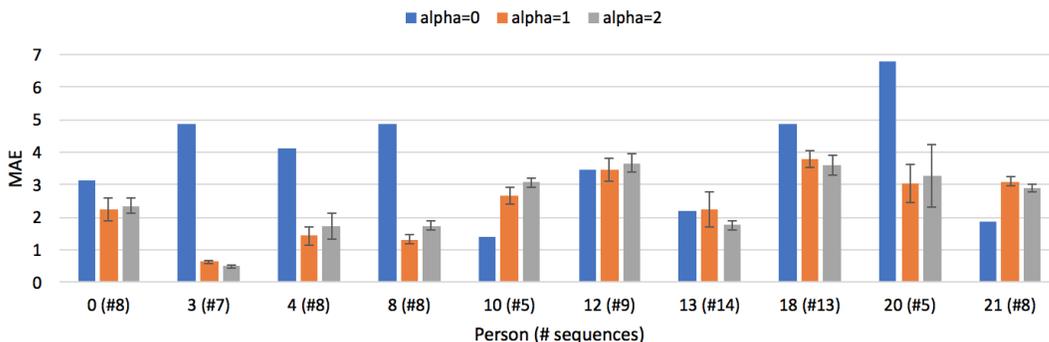


Figure 5: Mean MAE and standard deviation for the VAS estimation per test person when using pRNN-HCRF, and with varying  $\alpha = 0, 1, 2$ , used to compute I-FES. We also report the number of available sequences per person. The error bars are computed over 5-random selections of target sequences used to compute I-FES for target test person.

**Features.** As input features, we used the locations of 66 facial landmarks (see Fig.3), obtained by an Active Appearance Model and provided by the database creators [26]. To reduce redundant data in the input, we applied PCA, resulting in 40D feature vectors (preserving 95% of variance). Since majority of the frames are neutral (PSPI=0), for learning the PSPI estimator (the first stage), we balanced the training data so that the number of neutral frames occur with the same frequency as those with PSPI=1. We did this by removing neutral frames around the active pain segments. However, the whole sequences were used for learning VAS, and evaluation on test data. We performed the experiments in a subject independent setting, by randomly splitting the subjects in the training (15) and test (10) partitions. As evaluation measure, we used the Intra-Class Correlation ICC(3,1) [39], which is commonly used in behavioral sciences to measure agreement between annotators (in our case, the predicted and true PSPI/VAS levels). We also report the Mean Absolute Error (MAE), commonly used for ordinal prediction tasks [23, 36].

**Models.** For comparisons with our RNN-HCRF model,

we also include the results for the baseline model based on Support Vector Regression (SVR). We use SVR as the baseline method for PSPI estimation, as done in [22], and its personalized version, SVR-HCRF, using the proposed approach. Furthermore, to show the benefits of our hierarchical learning, we report the results obtained when the VAS estimation is attempted directly from input of the facial landmarks HCRF( $X$ ) thus, skipping the PSPI estimation. We also show the results for estimation of VAS directly from the ground-truth PSPI. Note that we do not compare here with state-of-the-art models for PSPI estimation, since in this work we focus on the evaluation of VAS estimation, which has not been reported before in the literature. However, these approaches can be explored for replacing the first step of our method as future work.

**Results.** From Fig.4, we also observe that PSPI estimated by RNN-LSTMs achieves accurate prediction of low intensity levels, while performing poorly on higher levels, likely due to the imbalance in the training data. However, Table 1 shows the improved PSPI estimation performance of the LSTM-RNNs over SVR [22] and non-recurrent NN in both measures. This is attributed to modeling of temporal

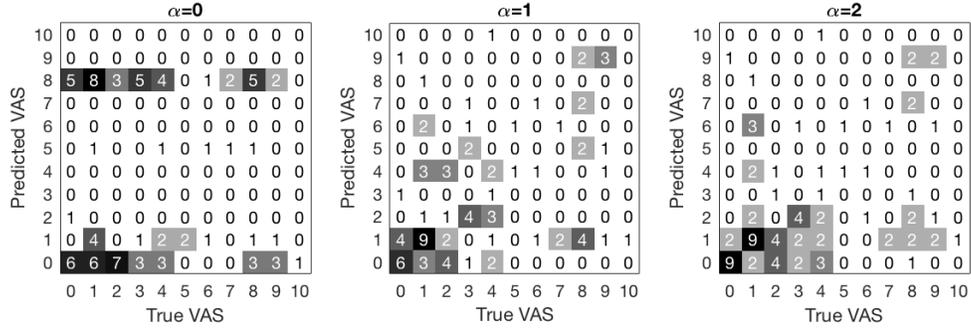


Figure 6: Confusion matrices for the predicted VAS scores using the proposed pRNN-HCRF model.

information as well as different layers (more powerful feature extractors) by the former. These results are also similar to those obtained by [26], and can further be improved using some of the state-of-the-art methods for pain intensity estimation (e.g., [38, 8]). Although the estimation of PSPI is not perfect when the base RNN/NN/SVR are used, they still can provide a good proxy of PSPI needed for estimation of VAS, as demonstrated below.

Table 2 shows the VAS estimation results. The methods considered, from top down, are: (1) HCRF(PSPI) using manually coded PSPI values as input and HCRFs, (2) the proposed two-stage RNN-HCRF model (with automatically estimated PSPI), (3) the same as (2) but using SVR for the PSPI prediction (SVR+HCRF), and (4) HCRF( $X$ ) using the facial landmarks (PCA pre-processed) as input. All four methods were evaluated with different personalization approaches, with  $\alpha$  from 0 (no personalization), to 1 and 2 sequences of target persons (personalized), and with five random repetitions. Note that without the personalization, directly estimating VAS scores with HCRF( $X$ ) outperforms the other models by a large margin. However, once the model personalization is performed, the proposed RNN-HCRF highly improves performance, outperforming both unpersonalized models ( $\alpha = 0$ ) as well as SVR+HCRF and HCRF( $X$ ), while the latter two achieve similar performance. Finally, note that using only a single sequence of test persons to compute their I-FES, we attain a high boost in the models’ performance.

To get more insights into the model’s performance, in Fig. 5, we show MAE<sup>2</sup> for the VAS estimation per test person. Again, we vary the number of sequences to compute the I-FES for test subjects as:  $\alpha = 0, 1, 2$ . Note that the proposed approach outperforms its unpersonalized counterpart ( $\alpha = 0$ ) on all test persons apart from 5&21, and performs similarly on 13. A possible reason for this is that more sequences of these persons are needed to estimate their I-FES. On the other hand, note the high improvement on persons,

<sup>2</sup>We do not report ICC here as it cannot be reliably computed per person because of small # of sequences.

e.g., 3,4,8 and 20. Finally, Fig. 6 shows confusion matrices for the estimated VAS scores using the proposed pRNN-HCRF model. Observe the region in the upper left corner for the unpersonalized model ( $\alpha = 0$ ), where significant confusion occurs. These are corrected for by the proposed personalized approach.

## 5. Conclusions

Existing work on automatic pain estimation focuses on (arguably) objective pain measures such as PSPI, derived directly from facial action units. However, as we showed in this paper, for estimation of self-reported VAS, this traditional one-size-fits-all approach fails to account for individual differences. Yet, these proved to be critical for improving the estimation of VAS for each person. In this paper, we introduced and evaluated a new approach that can provide automatic personalized estimation of self-reported VAS pain levels. We showed, on the benchmark pain dataset of face images, that structuring the model to provide personalization results in improved estimation of VAS. The focus of our future research will be on investigating the robustness of the proposed expressiveness measure and on alternative measures that can be used to capture the individual biases inherently present in self-reports of pain. We also plan to extend our inference algorithm to be able to learn jointly the PSPI and VAS estimation. Finally, we plan to perform a more extensive evaluation of the proposed approach and investigate further the relationship between different pain scores (VAS, OPI, PSPI), and their automatic estimation from facial expressions.

## Acknowledgements

The work of O. Rudovic is funded by European Union H2020, Marie Curie Action - Individual Fellowship no. 701236 (EngageMe).

## References

- [1] C. M. Aasted, M. A. Yucel, S. C. Steele, K. Peng, D. A. Boas, L. Becerra, and D. Borsook. Frontal lobe hemodynamic responses to painful stimulation: A potential brain marker of nociception. *PLoS ONE*, 11(11):1–12, 2016. **2**
- [2] B. Aicher, H. Peil, B. Peil, and H.-C. Diener. Pain measurement: Visual Analogue Scale (VAS) and Verbal Rating Scale (VRS) in clinical trials with OTC analgesics in headache. *Cephalalgia : an international journal of headache*, 32(3):185–97, 2012. **1**
- [3] A. B. Ashraf, S. Lucey, J. F. Cohn, T. Chen, Z. Ambadar, K. M. Prkachin, and P. E. Solomon. The painful face Pain expression recognition using active appearance models. *Image Vis. Comput.*, 27(12):1788–1796, nov 2009. **2**
- [4] A. C. de C Williams and K. D. Craig. Updating the definition of pain. *Pain*, Publish Ah(Box 1):1–14, 2016. **1**
- [5] R. A. Calvo and S. D’Mello. Affect detection: An interdisciplinary review of models method and their applications. *IEEE Trans. Affect. Comput.*, 1(1):18–37, 2010. **2**
- [6] W. Chen, A. Sano, D. Lopez Martinez, S. Taylor, A. W. Mchill, A. J. K. Phillips, L. Barger, E. B. Klerman, and R. W. Picard. Multimodal Ambulatory Sleep Detection. *IEEE International Conference on Biomedical and Health Informatics (BHI)*, pages 465–468, 2017. **3**
- [7] J. F. Cohn and F. De la Torre. Automated face analysis for affective. *The Oxford handbook of affective computing*, page 131, 2014. **4**
- [8] J. Egede, M. Valstar, and B. Martinez. Fusing Deep Learned and Hand-Crafted Features of Appearance, Shape, and Dynamics for Automatic Pain Estimation. 2017. **2, 8**
- [9] P. Ekman, W. V. Friesen, and J. C. Hager. *Facial action coding system*. A Human Face, Salt Lake City, UT, 2002. **4**
- [10] J. T. Farrar, J. T. Farrar, R. K. Portenoy, R. K. Portenoy, J. a. Berlin, J. a. Berlin, J. L. Kinman, J. L. Kinman, B. L. Strom, and B. L. Strom. Defining the clinically important difference in pain outcome measures. *Pain*, 88:287–294, 2000. **1**
- [11] M. A. Ferreira-Valente, J. L. Pais-Ribeiro, and M. P. Jensen. Validity of four pain intensity rating scales. *Pain*, 152(10):2399–2404, 2011. **1**
- [12] A. Graves and N. Jaitly. Towards End-to-End Speech Recognition with Recurrent Neural Networks. *Proc. 31st Int. Conf. Mach. Learn.*, 32(1):17641772, 2014. **3**
- [13] H. Gunes and M. Piccardi. Automatic temporal segment detection and affect recognition from face and body display. *IEEE Trans. Sys., Man Cybern. B*, 39(1):64–84, 2009. **2**
- [14] M. Haefeli and A. Elfering. Pain assessment. *European Spine Journal*, 15(1):S17–S24, 2006. **1**
- [15] Z. Hammal and J. F. Cohn. Automatic detection of pain intensity. In *Proc. 14th Int. Conf. Multimodal Interact.*, pages 47–52, Santa Monica, CA, USA, oct 2012. ACM. **2, 6**
- [16] G. A. Hawker, S. Mian, T. Kendzerska, and M. French. Measures of adult pain: Visual Analogue Scale for Pain (VAS Pain), Numeric Rating Scale for Pain (NRS Pain), McGill Pain Questionnaire (MPQ), Short-Form McGill Pain Questionnaire (SF-MPQ), Chronic Pain Grade Scale (CPGS), Short Form-36 Bodily Pain Scale (SF. *Arthritis Care and Research*, 63(SUPPL. 11):240–252, 2011. **1**
- [17] S. Hochreiter, Y. Bengio, P. Frasconi, and J. Schmidhuber. Gradient Flow in Recurrent Nets: the Difficulty of Learning Long-Term Dependencies. 2001. **3**
- [18] S. Hochreiter and J. Schmidhuber. Long Short-Term Memory. *Neural computation*, 9(8):1735–80, 1997. **2, 3**
- [19] M. P. Jensen, C. Chen, and A. M. Brugger. Interpretation of visual analog scale ratings and change scores: A reanalysis of two clinical trials of postoperative pain. *Journal of Pain*, 4(7):407–414, 2003. **1**
- [20] M. P. Jensen, S. A. Martin, and R. Cheung. The meaning of pain relief in a clinical trial. *Journal of Pain*, 6(6):400–406, 2005. **1**
- [21] J. Joshi, R. Goecke, G. Parker, and M. Breakspear. Can body expressions contribute to automatic depression analysis? In *2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, pages 1–7. IEEE, 4 2013. **2**
- [22] S. Kaltwang, O. Rudovic, and M. Pantic. Continuous Pain Intensity Estimation from Facial Expressions. In *Adv. Vis. Comput.*, volume 7432 of *Lecture Notes in Computer Science*, pages 368–377, Heidelberg, jul 2012. Springer. **2, 3, 7**
- [23] M. Kim and V. Pavlovic. Structured output ordinal regression for dynamic facial emotion intensity prediction. *ECCV*, pages 649–662, 2010. **7**
- [24] J. Lafferty, A. McCallum, and F. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. *ICML*, pages 282–289, 2001. **4**
- [25] P. Lucey, J. F. Cohn, I. Matthews, S. Lucey, S. Sridharan, J. Howlett, and K. M. Prkachin. Automatically detecting pain in video through facial action units. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 41(3):664–674, 2011. **2**
- [26] P. Lucey, J. F. Cohn, K. M. Prkachin, P. E. Solomon, and I. Matthews. Painful data: The UNBC-McMaster shoulder pain expression archive database. *2011 IEEE International Conference on Automatic Face and Gesture Recognition and Workshops, FG 2011*, pages 57–64, 2011. **2, 3, 6, 7, 8**
- [27] S. C. Mackey. Central neuroimaging of pain. *Journal of Pain*, 14(4):328–331, 2013. **2**
- [28] J. Melzack, Ronald; Katz. Pain Measurement in Adult Patients. In S. McMahon, M. Koltzenburg, I. Tracey, and D. C. Turk, editors, *Wall & Melzack’s Textbook of Pain*, chapter 21. Elsevier, 6th edition, 2013. **1**
- [29] H. Mohan, J. Ryan, B. Whelan, and A. Wakai. The end of the line? The Visual Analogue Scale and Verbal Numerical Rating Scale as pain assessment tools in the emergency department. *Emergency medicine journal : EMJ*, 27(5):372–375, 2010. **1**
- [30] M. Monwar and S. Rezaei. Pain recognition using artificial neural network. In *Proc. IEEE Int. Symp. Signal Process. Inf. Technol.*, pages 28–33. IEEE, 2006. **2**
- [31] A. Moore, O. Moore, H. McQuay, and D. Gavaghan. Deriving dichotomous outcome measures from continuous data in randomised controlled trials of analgesics: Use of pain intensity and visual analogue scales. *Pain*, 69(3):311–315, 1997. **1**

- [32] D. D. Price, P. A. McGrath, A. Rafii, and B. Buckingham. The validation of visual analogue scales as ratio scale measures for chronic and experimental pain. *Pain*, 17(1):45–56, 1983. 1
- [33] K. M. Prkachin and P. E. Solomon. The structure, reliability and validity of pain expression: Evidence from patients with shoulder pain. *Pain*, 139(2):267–274, 2008. 1, 2, 4
- [34] A. Quattoni, M. Collins, and T. Darrell. Conditional random fields for object recognition. *NIPS*, pages 1097–1104, 2004. 2, 3, 4
- [35] O. Rudovic, V. Pavlovic, and M. Pantic. Automatic pain intensity estimation with heteroscedastic conditional ordinal random fields. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8034 LNCS(PART 2):234–243, 2013. 3, 4
- [36] O. Rudovic, V. Pavlovic, and M. Pantic. Context-sensitive dynamic ordinal regression for intensity estimation of facial action units. *TPAMI*, 37(5):944–958, 2014. 7
- [37] A. Ruiz, O. Rudovic, X. Binefa, and M. Pantic. Multi-Instance Dynamic Ordinal Random Fields for Weakly-Supervised Pain Intensity Estimation. *Asian Conference on Computer Vision*, 9905(1):35–35, 2016. 2
- [38] A. Ruiz, O. Rudovic, X. Binefa, and M. Pantic. Deep Pain: Exploiting Long Short-Term Memory Networks for Facial Expression Classification. In *IEEE transactions on cybernetics*, pages 171–186. 2017. 2, 3, 8
- [39] P. E. Shrouf and J. L. Fleiss. Intraclass correlations: Uses in assessing rater reliability. *Psychological Bulletin*, 86(2):420–428, 1979. 7
- [40] K. Sikka, A. Dhall, and M. Bartlett. Weakly supervised pain localization using multiple instance learning. In *2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*, pages 1–8. IEEE, 4 2013. 2
- [41] M. J. L. Sullivan, D. A. Tripp, and D. Santor. Gender differences in pain and pain behavior: The role of catastrophizing. *Cognitive Therapy and Research*, 24(1):121–134, 2000. 2
- [42] K. H. Todd. Clinical versus statistical significance in the assessment of pain relief. *Annals of Emergency Medicine*, 27(4):439–441, 1996. 1
- [43] R. Treister, M. Kliger, G. Zuckerman, I. G. Aryeh, and E. Eisenberg. Differentiating between heat pain intensities: The combined effect of multiple autonomic parameters. *Pain*, 153(9):1807–1814, 2012. 1, 2
- [44] T. D. Wager, L. Y. Atlas, M. a. Lindquist, M. Roy, C.-W. Woo, and E. Kross. An fMRI-based neurologic signature of physical pain. *The New England journal of medicine*, 368(15):1388–97, 2013. 2
- [45] S. Wang, A. Quattoni, L.-P. Morency, D. Demirdjian, and T. D. Hidden conditional random fields for gesture recognition. *CVPR*, pages 1097–1104, 2006. 2, 3, 4
- [46] P. Werner, A. Al-Hamadi, and R. Niese. Pain recognition and intensity rating based on Comparative Learning. In *Proc. IEEE Int. Conf. Image Process.*, pages 2313–2316. IEEE, 2012. 2
- [47] D. J. Wilkie. Facial expressions of pain in lung cancer. *Analgesia*, 1(2):91–99, 1995. 2
- [48] A. C. D. C. Williams. Facial expression of pain: an evolutionary account. *The Behavioral and brain sciences*, 25(4):439–455, 2002. 1
- [49] A. Williamson and B. Hoggart. Pain: A review of three commonly used pain rating scales. *Journal of Clinical Nursing*, 14(7):798–804, 2005. 1
- [50] J. Younger, R. McCue, and S. Mackey. Pain outcomes: A brief review of instruments and techniques. *Current Pain and Headache Reports*, 13(1):39–43, 2009. 1
- [51] M. A. Yücel, C. M. Aasted, M. P. Petkov, D. Borsook, D. A. Boas, and L. Becerra. Specificity of hemodynamic brain responses to painful stimuli: a functional near-infrared spectroscopy study. *Scientific reports*, 5:9469, 2015. 2
- [52] Z. Zeng, M. Pantic, G. I. Roisman, and T. S. Huang. A survey of affect recognition methods: audio, visual, and spontaneous expressions. *IEEE Trans. Patt. Anal. Mach. Intel.*, 31(1):39–58, 2009. 2