

Towards Pain Monitoring: Facial Expression, Head Pose, a new Database, an Automatic System and Remaining Challenges

Philipp Werner¹
Philipp.Werner@ovgu.de

Ayoub Al-Hamadi¹
Ayoub.Al-Hamadi@ovgu.de

Robert Niese¹
Robert.Niese@ovgu.de

Steffen Walter²
Steffen.Walter@uni-ulm.de

Sascha Gruss²
Sascha.Gruss@uni-ulm.de

Harald C. Traue²
Harald.Traue@uni-ulm.de

¹ Institute for Information Technology and Communications,
Otto-von-Guericke University,
Magdeburg, Germany

² Department for Psychosomatic Medicine and Psychotherapy,
University of Ulm,
Ulm, Germany

Abstract

Pain is what the patient says it is. But what about these who cannot utter? Automatic pain monitoring opens up prospects for better treatment, but accurate assessment of pain is challenging due to the subjective nature of pain. To facilitate advances, we contribute a new dataset, the BioVid Heat Pain Database which contains videos and physiological data of 90 persons subjected to well-defined pain stimuli of 4 intensities. We propose a fully automatic recognition system utilizing facial expression, head pose information and their dynamics. The approach is evaluated with the task of pain detection on the new dataset, also outlining open challenges for pain monitoring in general. Additionally, we analyze the relevance of head pose information for pain recognition and compare person-specific and general classification models.

1 Introduction

Accurate assessment of pain is the key to successful pain management [20]. Adequate treatment of pain is important as it does not only bring relief from mental suffering, but may also speedup patient's recovery [8, 13]. Pain is a subjective experience and thus self-report is considered the gold standard for pain measurement. However, self-report is not always reliable and valid, e.g. for demented patients [26]. Furthermore, it cannot be applied at all for unconscious or newborn patients. Observational and physiological measures [21] can help in these cases. They may also facilitate to overcome the weaknesses of simple rating

scales [23] which are common practice in clinics. For adequate pain management, the assessment must be repeated regularly, especially if the patient cannot call for help by himself. For providing a continuous assessment, we work towards an automatic system for measurement and monitoring of pain, which can alert hospital staff timely and provide additional information for the patient’s medical record. In this paper, we focus on facial expressions and head movement to detect pain of different intensities.

1.1 Related Work

Zhihong *et al.* [25] survey on the extensive previous research in facial expressions recognition and Murphy-Chutorian and Trivedi [24] give an overview on previous head pose estimation approaches. In the more specific field of pain recognition from facial expressions, the list of contributions is less extensive. Lucey *et al.* [22] use Active Appearance Models to track and align the face based on manually labeled key-frames. They extract shape and appearance features and apply a Support Vector Machine (SVM) to classify at frame level whether there is a facial expression of pain, i.e. whether any of the pain related action units previously found by Prkachin [29] is present. At sequence level they classify three levels of intensity by fusing frame level results via majority voting. As they try to mimic an expert observer, the ground truth intensity labels were assigned by considerably trained observers. They also analyze the correlation between observed intensity and change in the head’s orientation, because they observed head movements coinciding with facial expression of pain. All experiments are conducted using the UNBC-McMaster Shoulder Pain Expression Archive Database, which is publicly available now. On the same database, Chen *et al.* [3] recently compared different learning algorithms for training *person-specific* classifiers. Based on the landmarks provided by Lucey *et al.* and Local Binary Pattern features, their Inductive Transfer Learning outperforms Lucey *et al.* [22] on the frame level. Also on the UNBC database, Hammal and Cohn [6] used Log-Normal filter based features and SVMs to classify four pain expression intensities (no pain and three pain levels).

Niese *et al.* [16] distinguishes pain from five other expression classes (neutral and four basic emotions) based on facial distances and angles taken as input of an SVM. Werner *et al.* [27] extend their work by measuring the pain expression intensities in a continuous scale and integrating gradient based features for measuring facial wrinkles. Hammal and Kunz [12] also utilize distances and measure nasal wrinkles, which are automatically extracted from frontal face. Based on the Transferable Belief Model which is designed for handling noisy and partial input, they provide a prediction for each frame, but incorporate dynamics using a temporal evolution model to refine the prediction until the end of the sequence. They propose the inclusion of context variables to bias the classifier towards the most relevant expression. Littlewort *et al.* [18] use Gabor filters and SVMs to classify real versus posed pain at sequence level. Their automatic system outperforms untrained observers in this task.

1.2 Contributions

This work contributes the following. For a visual overview, see Fig. 1.

- We introduce the BioVid Heat Pain Database (see Sec. 2), a newly collected experimental pain research dataset. In contrast to existing databases known to the authors, it does not only provide video information, but also physiological data. Thus, it promotes advancements in both, observational and physiological pain measurement and

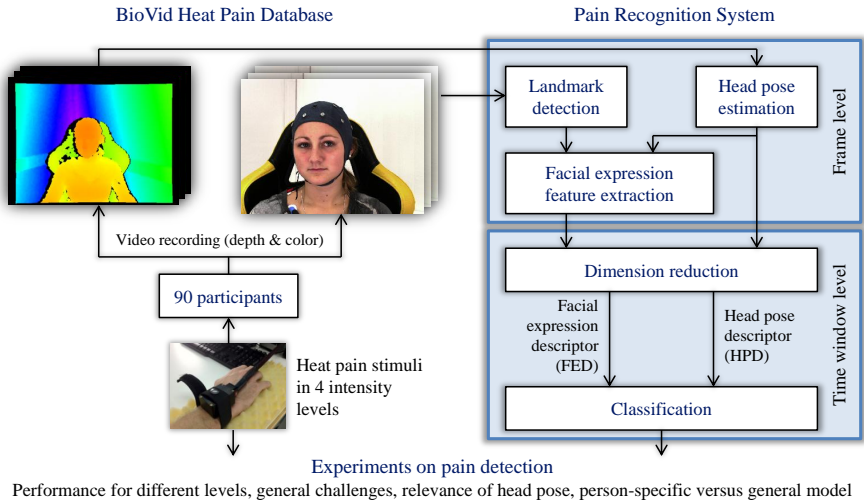


Figure 1: Overview on the new database, the recognition concept and the experiments.

facilitates the fusion of both in one automatic recognition system. The database will be made available for non-commercial research purposes.

- A new approach to recognize pain from facial expression and head pose information is proposed (see Sec. 3). It utilizes dynamics through a time window descriptor which is calculated from frame level features. To the best knowledge of the authors, our pain recognition system is the first which takes up the challenges arising with fully automatic feature extraction when permitting non-frontal head poses. It is also the first approach using head pose information for classification.
- We conduct experiments applying our recognition system to the BioVid Database. Whereas the others evaluate their systems based on observational measures (e.g. facial actions) or do not even specify the source of their ground truth labels, we compete with the gold standard of pain assessment, i.e. self-report. Since our focus is on recognizing pain, not only facial expressions of pain, we try to classify the stimulus. This does not only reveal strengths and weaknesses of the recognition system, but also challenges caused by the nature of pain. In Sec. 4.1 we investigate which level of pain can be detected by our system, also outlining some open challenges for pain monitoring in general. Next, we analyze the role of head pose for pain recognition (see Sec. 4.2), an issue recently raised by Lucey *et al.* [14]. While focusing on person-specific models, we also compare with a general model. In the conclusion (Sec. 5), we summarize our current work and directions for further research.

2 BioVid Heat Pain Database

The BioVid Heat Pain Database was collected in a study with 90 participants, who were recruited from three age groups (18-35, 36-50 and 51-65 years) with each of the groups consisting of 15 men and 15 women. The pain was induced experimentally by a thermode (Medoc

PATHWAY, <http://www.medoc-web.com>) at the right arm (see Fig. 1). The experiments were recorded with video cameras and physiological sensors. We utilized three synchronized AVT Pike F145C cameras (1388 x 1038 colored pixels, 25 Hz), one directly in front of the study participant and two at the side. The side cameras captured a frontal face in case the participant turned the head 45° left or right, respectively. The participants were explicitly allowed to move their head freely. We also acquired depth information using a Microsoft Kinect camera (640 x 480 pixels, ca. 30 Hz). In terms of the participant's physiology, we recorded the skin conductance level (SCL), the electrocardiogram (ECG), the electromyogram (EMG) of three pain related muscles, and the electroencephalogram (EEG).

Before the data recording was started, we determined the participant's individual pain threshold and pain tolerance, i.e. the temperatures for which the participant's sensing changes from heat to pain respectively for which the pain becomes unacceptable. We used these thresholds as the temperatures for lowest and highest pain levels and added two additional intermediate levels, whereas the temperatures were equally distributed in the range (see Fig. 2a). In the first part of the main experiment, each of the four pain levels was stimulated 20 times in randomized order (see Fig. 2b). For each stimulus, the maximum temperature was held for 4 s. The pauses between the stimuli were randomized between 8-12 s. In the second part, the participant posed pain and basic emotions. Afterwards, images and video clips were shown to elicit spontaneous emotions. This part is intended to facilitate analyses on the distinction of genuine pain from emotions and posed pain. Subsequently, the pain stimulation of the first part was repeated with EMG electrodes attached to the facial muscles zygomaticus major and corrugator supercilii. In the first part, these had been omitted to avoid the partial occlusion of the face. After this third part, the recording was stopped and the participant filled in a set of questionnaires about his personality and psychosomatic health.

In this work, we consider the first part of the main experiment (pain stimuli without facial EMG) viewed from the center cameras (one Pike and Kinect). Analyses of the remaining data will be presented in follow-up publications. The database will be made available to other researchers for non-commercial purposes at the beginning of next year.




Figure 2: Pain stimuli in the BioVid Heat Pain Database. (a) Pain levels and according person-specific temperatures. (b) Exemplary temperature curve with alternating stimuli and pauses. The time windows used for the experiments are defined at the top (see Sec. 3.4).




Figure 3: Feature extraction at frame level. (a) Automatic face and facial landmark detection with Haar-like feature detector cascade. (b) Mouth and brow point (yellow) detection. Mouth ROI (red) and contour (green). Eye axis (blue), line segment for brow search (green). (c) Measured 3D point cloud (from Kinect), model fitting residuals (cyan) and nose-tip coordinate system illustrating the determined pose. (d) Facial distance features. (e) Regions for mean gradient magnitude features (blue: nasal wrinkles, yellow: nasolabial folds, green: eye closing) based on anchors and facial axes (white).

3 Pain Recognition System

3.1 Facial Feature Point Detection

Our facial expression analysis is based on a set of landmarks, which are extracted automatically (see Fig. 3a-b). For each image of the video stream, we first find the face employing the Haar-like feature detector cascade by Lienhart *et al.* [10]. Inside the facial region we apply the eye detector cascades trained by Castrillón *et al.* [11] and the mouth corner cascades by Panning *et al.* [12]. To identify false detections, we compare the candidates to an estimate given by a generic face model which is placed inside the face’s bounding box. Based on the obtained eye center and mouth corner points we find the remaining landmarks as described by Niese *et al.* [13]. The upper and lower lip points are determined by a color based lip segmentation approach. It is based on the histogram of the normalized green channel for the mouth’s region of interest. The segmentation contour is also used to redefine the mouth corner points, as the results are more accurate than the detector in most cases. Each of the eyebrow points is selected from a line segment perpendicular to the eye axis by finding the maximum peak of the vertical gradient. The eye axis is also used to compensate for in-plane head rotation by rotating the detector input image of the next frame.

3.2 Head Pose Estimation

To estimate the head pose, we utilize the depth information recorded by the Kinect camera. For a volume of interest the depth map is converted into a 3D point cloud by a pinhole camera model (see Fig. 3c). Afterwards, a generic face model is registered with the measured point cloud using a variant of the Iterative Closest Point (ICP) algorithm as presented by Niese *et al.* [13]. It provides a 6D head pose vector including the 3D position and 3D orientation (pitch, yaw and roll). The ICP is initialized with the previous frame’s pose vector. If the residual registration error is greater than a threshold, the pose estimation is declared to be

failed. This case triggers a reset of the pose vector relative to the center of gravity in the upper part of the point cloud.

3.3 Frame Level Facial Expression Features

For each image frame we extract a set of distance and gradient features. They are selected to capture several pain related facial actions which have been identified by several previous studies, e.g. by Prkachin [19]. These actions include lowering of the brows, tightening of the lid, closing of the eyes, raising of the cheeks and the upper lip, wrinkling of the nose, and stretching and opening of the mouth. To uncouple facial expression from head pose, distances are calculated in 3D, as proposed by Niese *et al.* [20]. Using a pinhole camera model, the detected landmarks (see Sec. 3.1) are projected onto the surface of the generic face model placed according to the current head pose (see Sec. 3.2). From the obtained 3D points (depicted in Fig. 3d), we calculate the distances between brows and eyes ($\overline{B_l E_l}$, $\overline{B_r E_r}$), eyes and mouth ($\overline{E_l M_l}$, $\overline{E_r M_r}$), brows and mouth ($\overline{B_l M_l}$, $\overline{B_r M_r}$), as well as the width and height of the mouth ($\overline{M_l M_r}$, $\overline{M_l M_b}$). Next to these distances, some facial changes are measured from the texture. Based on the landmarks we define rectangular regions of interest (see Fig. 3e) and calculate the mean gradient magnitude for each of these regions. This way, we measure nasal wrinkles and the nasolabial furrows as done by Werner *et al.* [21]. In addition, the idea is applied to measure the closing of the eyes, which involves a decrease in the mean gradient magnitude of the eye region. All regions are anchored by the eye center and mouth corner points which are also utilized to define eye axis ($\overline{E_l E_r}$) and the vertical face axis (line between the center of $\overline{E_l E_r}$ and $\overline{M_l M_r}$). Based on the anchor points and the axes, the regions are placed according to assumptions derived from empirical investigations of our data.

3.4 Time Window Descriptors

Whereas a single image may contain enough information to estimate the intensity of a facial expression of pain, it misses the dynamics which we consider to contain valuable information about the underlying feeling of pain. Thus, we classify time windows rather than single images. The previously described processing steps provide (1) a 6D vector of pose parameters per depth frame and (2) a 13D vector of facial expression features per color frame. For our experiments we use time windows of 6 s length, clipping 6 discrete-time signals of pose parameters and 13 discrete-time signals of facial expression features. To reduce the number of dimensions, we calculate a descriptor of each signal. We first apply a 3 frame median filter for noise reduction and estimate the first and second temporal derivation of the signal. Then, we calculate 7 statistic measures (mean, median, range, standard and median absolute deviation, interquartile and interdecile range) for each, the smoothed signal and its first and second derivation, resulting in a 21D descriptor per signal. The 6 descriptors of the head pose signals are combined into the head pose descriptor (HPD), the 13 descriptors of the facial expression feature signals into the facial expression descriptor (FED).

3.5 Classification

We classify the time windows with a Support Vector Machine (SVM) implemented by Chang and Lin [22] using a radial basis function kernel. Before fed into the SVM, the previously defined descriptors are standardized, i.e. each variable is converted to a z-score based on the




Figure 4: No pain (BLN) versus each pain intensity (PA n). Distributions of person-specific classification results subdivided by age groups. Each box plot visualizes the minimum and maximum (left and right whisker), the 25% and 75% quartile (left and right border of the box), the median (middle of the box) and the mean (black dot) of the distribution. The chance and perfect performance levels are shown with dotted lines.

corresponding mean and standard deviation determined from the training set. Our sample set consists of 5 classes: no pain (baseline, BLN), stimulus at pain threshold (PA1), stimulus at pain tolerance (PA4) and at the two intermediate levels (PA2 and PA3). Depending on the classification task (see Sec. 4), the samples of a subset are used. In the training phase, we first apply a grid search with 5-fold stratified cross validation on the training set to select optimal values for the SVM parameters γ , C and C_{BLN} . The latter is a separate penalty parameter for the baseline class, which is necessary to handle the unbalanced sample counts. In the second step, the SVM is trained with the whole training set and the determined parameters.

4 Experimental Results

In our experiments we focus on person-specific classification, because it allows to reveal some challenges for automatic facial expression based pain recognition that has not been addressed by other researchers so far. Additionally, we think person-specific models are the best way to go, as they facilitate better recognition performance than generalized models. Getting person-specific training data may be considered a problem for real-world applications, but Chen *et al.* [9] recently proposed a promising approach to handle it.

All experiments are conducted with samples taken from the BioVid Heat Pain Database (see Sec. 2). For each study participant, there are 20 samples per pain level (PA1 up to PA4) and 80 no-pain samples (baseline, BLN). The corresponding time windows are selected as shown in Fig. 2b. The classification performance is measured with the macro-averaged F_1 measure [24]. The F_1 measure is the harmonic mean of recall and precision. Thus, it involves both, false alarms and misses. Macro-averaging weights each class equally, which is desirable in our context, since the BLN class has more samples without being more important than the others. The performance of each person-specific model is evaluated through 10-fold stratified cross validation. The general model is evaluated using leave-one-subject-out cross validation, also yielding one performance value per subject.

4.1 Pain versus no pain

The first important aspect of pain monitoring is the detection of pain, e.g. for alerting hospital staff. But how to define pain? Some other researchers working on facial expression based pain recognition seem to equate pain and facial expressions of pain. Whereas there are strong

correlations, feeling pain does not necessarily cause a facial expression. Thus, equating both disregards pain without facial expressions. Among others, personality, social context and previous pain experiences influence the expression of pain [6]. Viewing videos of our dataset, we observed large differences in expressiveness. Whereas some persons already show considerable facial expressions at low pain intensities, other seem not to react at all, even when stimulated with high temperature at their pain tolerance level. To investigate this, we trained our pain recognition system to distinguish no-pain (BLN) from each of the different pain intensities (PAN). The cross validation results of the various person-specific models are summarized in Fig. 4 using box plots. As there are significant differences between the age groups, we present their results separately in the columns. The rows list the evaluated class pairs, i.e. no-pain classified versus each of the pain intensities. The chance level, i.e. 50% recall for both classes is at 0.45 for the considered classification task, as the precision is factored in as well. The first aspect to notice is, that the mean performance is increasing with the pain intensity for all age groups. This confirms that a higher pain intensity results in more observable reaction and thus is easier to classify. Whereas the performance is above chance level for the vast majority of persons, it is also not clearly above it for the majority at the low pain levels PA1 and PA2. It may be possible to improve the results with a more sophisticated system, but we think it is not possible to detect these levels of pain with high reliability across all subjects only with observational measures, as many do not react visibly to low levels of pain. As a possible solution for pain monitoring we anticipate the fusion with physiological information to improve results. Nevertheless, as you can see from the right whisker, there are some expressive persons with high classification performance even for the low pain levels. This confirms the observed diversity in expressiveness. Level PA4 and PA3 can be reliably distinguished from BLN for the vast majority of persons, whereas the results are best for the young and worst for the old age group. Whereas a study by Kunz *et al.* [9] did not find any impact of age on the facial expression of pain, we observed that with older age there are more people with very low to no facial response to pain, even for very high temperature pain tolerance levels. In contrast to the study of Kunz *et al.*, our participants did not undergo any health examination. Since "older age is associated with an increase in pathological conditions that can affect pain processing and pain experience" [9], our study does not exclude these conditions to be more representative of the prospective clinical use of pain monitoring, where we cannot assume the patients to be healthy. The diversity of individuals is one of the major challenges for pain recognition. Thus, we will analyze age, expressiveness and other group effects further in future work, also utilizing the facial action coding system and the physiological signals. A secondary factor which negatively affects the classification performance of the middle and old age group, is that the detection of the facial feature point (see Sec. 3.1) is less robust and accurate, mainly because of wrinkles and beards. Although better, the robustness and accuracy are an issue for the young age group as well, because every error introduces noise which can hide subtle, but relevant reactions to pain. Therefore, robust and accurate detection of facial landmarks across the diverse appearance of faces is a further challenge to be solved, especially when considering out-of-plane rotations.

4.2 Relevance of Head Pose

Now we address the question, whether head pose information can be utilized for pain recognition. First, the quality of our head pose estimation (see Sec. 3.2) was tested with the BIWI Kinect Head Pose Database [4]. We focused on the poses, which are relevant for our esti-




Figure 5: Histograms of the head’s pitch, yaw and roll estimation in utilized time windows.




Figure 6: Results for no pain versus pain tolerance classified from head pose only (HPD), facial expression only (FED) and using both. (a) Person-specific model. (b) General model.

mation task, i.e. we restricted the range of considered rotations (pitch and roll in between $\pm 40^\circ$). The mean \pm standard deviation of the absolute error of our estimation are $3.4 \pm 8.1^\circ$ for pitch, $5.8 \pm 12.4^\circ$ for yaw, $3.9 \pm 8.4^\circ$ for roll and $18.5 \pm 24.0mm$ for the position. The estimation failed in 7.6% of the frames. Compared to the results of Fanelli *et al.* [4], the angular error of our approach is lower whereas the rate of estimation failures (no valid result) is higher. As the database is more challenging than our data, i.e. the percentage of large rotations is much higher, we conclude that our approach is sufficient for the following analysis.

Fig. 5 depicts the histograms of the head orientation parameters estimated across the utilized time windows of the BioVid Heat Pain Database. Pose estimation failed in 1.7% of the frames. Compared to the UNBC McMaster Database as described by Lucey *et al.* [14], the variance is higher for pitch (77.4) and yaw (93.8), but slightly lower for roll (28.5). Whereas Lucey *et al.* report significant differences for yaw and roll when comparing pain and no-pain variances, our data’s variance differs most for pitch (PA4: 87.3, BLN: 77.5). We explain the differences by the differing pain stimuli. In the data of Lucey *et al.* the pain was located at shoulder. As this is near the muscles responsible for yaw and roll rotations, these may coincide with pain more often. In contrast, our stimulus is located at the lower arm. Thus, our stimulus location has less influence on the head movement. Lucey *et al.* did not find any direct relation between pain and head movement, which they initially expected. Nevertheless, we also observe movements of the head co-occurring with pain in our dataset. But often the movements are subtle and often they are indeed dominated by pitch rotations and raising of the head. To investigate whether head movement can be utilized for pain recognition, we trained and tested our system with different SVM input on the classification task BLN versus PA4. Fig. 6 presents the cross validation results for both, person-specific and general models. It also includes the always-predict-PA4 performance level (dashed line),

as it differs significantly from chance level. When only using head pose information (HPD), for the vast majority of persons the performance is significantly above chance for person-specific models. In contrast, for the general model the mean performance is at chance, suggesting that head movement during pain is an individual trait. However, when comparing results for facial expression only (FED) and both, FED and HPD, it can be seen that the performance can benefit from including head pose information for both, person-specific and general classifiers.

5 Conclusion

In this work, we introduced a new dataset for pain research, the BioVid Heat Pain Database, which includes videos and physiological data of 90 participants subjected to pain stimuli of four intensities. We also proposed a pain recognition system utilizing facial expression, head pose information and their dynamics. With this system, experiments in pain detection were conducted, trying to predict the stimulus rather than an observational measure. We also face the challenge of fully automatic feature extraction including severe out-of-plane head poses. It was shown that head pose information is valuable for pain recognition. Our system revealed high performance for high pain intensities across the majority of individuals when using person-specific models. However, there are persons which show only little or no facial expression, also during high stimulation. This, and the diversity of individuals needs further investigations, which we will start with a facial action coding system based analysis. For low pain intensities, the recognition is difficult as well. First, because often there is no facial reaction. Second, because some changes are too subtle to be measured. Further research is necessary in accurate and robust detection of facial landmarks, because less errors in face alignment and facial measurement are essential for reaching a performance level sufficient for the prospective clinical use. Next to the observational modality, we will analyze the recorded physiological data, utilize it for recognizing pain and combine it with this work to a multi-modal pain recognition system. We will also work towards taking the step from pain detection to measurement of pain intensity.

Acknowledgment

This work was funded by the German Research Foundation (DFG), project Schmerzerkennung. FKZ/GZ: Al 638/3-1 AOBJ 585843.

References

- [1] M. Castrillón, O. Déniz, C. Guerra, and M. Hernández. ENCARA2: real-time detection of multiple faces at different resolutions in video streams. *Journal of Visual Communication and Image Representation*, 18(2):130–140, April 2007. ISSN 1047-3203. doi: 10.1016/j.jvcir.2006.11.004. URL <http://www.sciencedirect.com/science/article/pii/S1047320306000782>.
- [2] Chih-Chung Chang and Chih-Jen Lin. LIBSVM – a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):27:1–27:27, 2011. URL <http://www.csie.ntu.edu.tw/~cjlin/libsvm/>.

- [3] Jixu Chen, Xiaoming Liu, Peter Tu, and Amy Aragonés. Person-specific expression recognition with transfer learning. In *Image Processing (ICIP), 2012 19th IEEE International Conference on*, pages 2621–2624, Orlando, 2012. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6467436.
- [4] Gabriele Fanelli, Thibaut Weise, Juergen Gall, and Luc Van Gool. Real time head pose estimation from consumer depth cameras. In *Pattern Recognition, Lecture Notes in Computer Science Volume 6835*, page 101–110. Springer, 2011. URL http://link.springer.com/chapter/10.1007/978-3-642-23123-0_11.
- [5] Thomas Hadjistavropoulos and Kenneth D. Craig. *Pain: psychological perspectives*. Lawrence Erlbaum Associates, Incorporated, 2004. ISBN 978-0-8058-4299-9.
- [6] Zakia Hammal and Jeffrey F. Cohn. Automatic detection of pain intensity. In *Proceedings of the 14th ACM International Conference on Multimodal Interaction*, pages 47–52, New York, 2012. ACM. ISBN 978-1-4503-1467-1. doi: 10.1145/2388676.2388688. URL <http://doi.acm.org/10.1145/2388676.2388688>.
- [7] Zakia Hammal and Miriam Kunz. Pain monitoring: A dynamic and context-sensitive system. *Pattern Recognition*, 45(4):1265–1280, April 2012. ISSN 0031-3203. doi: 10.1016/j.patcog.2011.09.014. URL <http://www.sciencedirect.com/science/article/pii/S0031320311003931>.
- [8] Henrik Kehlet. Acute pain control and accelerated postoperative surgical recovery. *Surgical Clinics of North America*, 79(2):431–443, April 1999. ISSN 0039-6109. doi: 10.1016/S0039-6109(05)70390-X. URL <http://www.sciencedirect.com/science/article/pii/S003961090570390X>.
- [9] Miriam Kunz, Veit Mylius, Karsten Schepelmann, and Stefan Lautenbacher. Impact of age on the facial expression of pain. *Journal of Psychosomatic Research*, 64(3):311–318, March 2008. ISSN 0022-3999. doi: 16/j.jpsychores.2007.09.010. URL <http://www.sciencedirect.com/science/article/pii/S0022399907004412>.
- [10] Rainer Lienhart, Er Kuranov, and Vadim Pisarevsky. Empirical analysis of detection cascades of boosted classifiers for rapid object detection. In *DAGM 25th Pattern Recognition Symposium*, pages 297–304, 2003.
- [11] Gwen C. Littlewort, Marian Stewart Bartlett, and Kang Lee. Automatic coding of facial expressions displayed during posed and genuine pain. *Image and Vision Computing*, 27(12):1797–1803, November 2009. ISSN 0262-8856. doi: 16/j.imavis.2008.12.010. URL <http://www.sciencedirect.com/science/article/pii/S0262885609000055>.
- [12] Patrick Lucey, Jeffrey F. Cohn, Kenneth M. Prkachin, Patricia E. Solomon, Sien Chew, and Iain Matthews. Painful monitoring: Automatic pain monitoring using the UNBC-McMaster shoulder pain expression archive database. *Image and Vision Computing*, 30(3):197–205, March 2012. ISSN 0262-8856. doi: 10.1016/j.imavis.2011.12.003. URL <http://www.sciencedirect.com/science/article/pii/S0262885611001363>.

- [13] Henry McQuay, Andrew Moore, and Douglas Justins. Treating acute pain in hospital. *BMJ: British Medical Journal*, 314(7093):1531, 1997. URL <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2126780/>.
- [14] E. Murphy-Chutorian and M. M Trivedi. Head pose estimation in computer vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(4): 607–626, April 2009. ISSN 0162-8828. doi: 10.1109/TPAMI.2008.106.
- [15] Robert Niese, Ayoub Al-Hamadi, and Bernd Michaelis. A stereo and color-based method for face pose estimation and facial feature extraction. In *Pattern Recognition, 18th International Conference on*, volume 1, pages 299–302, 2006. URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1698892.
- [16] Robert Niese, Ayoub Al-Hamadi, Axel Panning, Dominik Brammen, Uwe Ebmeyer, and Bernd Michaelis. Towards pain recognition in post-operative phases using 3D-based features from video and support vector machines. *International Journal of Digital Content Technology and its Applications*, 3(4):21–33, 2009. ISSN 1975-9339. doi: 10.4156/jdcta.vol3.issue4.2.
- [17] Robert Niese, Ayoub Al-Hamadi, Axel Panning, and Bernd Michaelis. Emotion recognition based on 2D-3D facial feature extraction from color image sequences. *Journal of Multimedia*, 5, October 2010. ISSN 1796-2048. doi: 10.4304/jmm.5.5.488-500. URL <http://ojs.academypublisher.com/index.php/jmm/article/view/0505488500>.
- [18] A. Panning, A. K. Al-Hamadi, R. Niese, and B. Michaelis. Facial expression recognition based on haar-like feature detection. *Pattern Recognition and Image Analysis*, 18 (3):447–452, September 2008. ISSN 1054-6618. doi: 10.1134/S1054661808030139. URL <http://www.springerlink.com/content/t532j123213pk446/>.
- [19] Kenneth M. Prkachin. The consistency of facial expressions of pain: a comparison across modalities. *Pain*, 51(3):297–306, December 1992. ISSN 0304-3959. doi: 16/0304-3959(92)90213-U. URL <http://www.sciencedirect.com/science/article/pii/030439599290213U>.
- [20] Michael Serpell, editor. *Handbook of Pain Management*. Springer, 2008. ISBN 978-1-908517-12-8.
- [21] Jenny Strong, Anita M. Unruh, Antony Wright, and G. David Baxter. *Pain: a textbook for therapists*. Churchill Livingstone Edinburgh, Scotland, 2002.
- [22] Philipp Werner, Ayoub Al-Hamadi, and Robert Niese. Pain recognition and intensity rating based on comparative learning. In *Image Processing (ICIP), 2012 19th IEEE International Conference on*, pages 2313–2316, Orlando, 2012.
- [23] Amanda Williams, Huw T. O. Davies, and Yasmin Chadury. Simple pain rating scales hide complex idiosyncratic meanings. *Pain*, 85(3):457–463, April 2000. ISSN 0304-3959. doi: 10.1016/S0304-3959(99)00299-7. URL <http://www.sciencedirect.com/science/article/pii/S0304395999002997>.

- [24] Yiming Yang. An evaluation of statistical approaches to text categorization. *Information retrieval*, 1(1-2):69–90, 1999. URL <http://link.springer.com/article/10.1023/A%3A1009982220290>.
- [25] Zhihong Zeng, M. Pantic, G. I. Roisman, and T. S. Huang. A survey of affect recognition methods: Audio, visual, and spontaneous expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(1):39–58, January 2009. ISSN 0162-8828. doi: 10.1109/TPAMI.2008.52.
- [26] Sandra M. G. Zwakhalen, Jan P. H. Hamers, Huda H. Abu-Saad, and Martijn P. F. Berger. Pain in elderly people with severe dementia: A systematic review of behavioural pain assessment tools. *BMC Geriatrics*, 6(1):3, January 2006. ISSN 1471-2318. doi: 10.1186/1471-2318-6-3. URL <http://www.biomedcentral.com/1471-2318/6/3/abstract>.