Faces of Pain: Automated measurement of spontaneous facial expressions of genuine and posed pain

Gwen Littlewort, Marian Stewart Bartlett, Kang Lee

Institute for Neural Computation
University of California, San Diego
San Diego, CA 92093-0523

Abstract

We present initial results from the application of a CERT system (Computer Expression Recognition Toolbox), developed in the Machine Perception Laboratory, to spontaneous facial expressions of pain. In this study, participants submerged their hand in a bath of water for three minutes. Each subject experienced three experimental conditions: baseline, real pain and posed pain. In the real pain condition, the water was 3 degrees Celsius, whereas in the baseline and posed pain conditions the water was 20 degrees Celsius. The video was coded for facial action units (FACS) by both human and computer. Our initial goal was to correctly determine which experimental condition is shown in a 60 second clip from a previously unseen subject. We chose a machine learning approach, previously used successfully to categorize basic emotional facial expressions in posed datasets as well as to code for individual facial actions (Littlewort, 2006). For this study, we trained individual Action Unit (AU) classifiers on 3000 single frames selected from three datasets: two posed expression sets, the Cohn-Kanade dataset of FACS-coded facial expressions (Kanade, Cohn, & Tian, 2000), the Ekman-Hager dataset of directed facial actions (Bartlett et al., 1999), and a spontaneous expression data set from Mark Frank (Bartlett et al., 2006). We used an automatic face and eye location system that employs boosting techniques (Fasel et al., 2005) to find, scale and align face images. We convolved Gabor filters with 96 by 96 pixel cropped face boxes and trained linear Support Vector Machines for each of 20 facial action units, in one versus all mode, irrespective of combinations with other actions. The output of the system was a real valued number indicating the distance to the separating hyperplane for each classifier. Applying this system to the pain video data produced a 20 channel output stream, consisting of one real value for each learned AU, for each frame of the video. This data was further analyzed to predict the difference between baseline and pained faces, the difference between expressions of real pain and fake pain, and the difference between faked pain with or without recent experience of real pain. In a preliminary analysis of 5 subjects, the system correctly identified the experimental condition (baseline, posed pain, real pain) for 93 percent of samples in a 3-way forced choice. This is considerably higher than the performance of naive human observers.