1 Emotion recognition in humans and machine using posed and spontaneous facial expression

2 Damien Dupré¹, Eva G. Krumhuber², Dennis Küster³, & Gary McKeown⁴

3 ¹ Dublin City University

4 ² University College London

5 ³ University of Bremen

6 ⁴ Queen's University Belfast

7 Author Note

8 Correspondence concerning this article should be addressed to Damien Dupré, Dublin
9 City University, Dublin 9, Republic of Ireland. E-mail: damien.dupre@dcu.ie

Abstract

In the wake of rapid advances in automatic affect analysis, commercial software tools for facial affect recognition have attracted considerable attention in the past few years. While several options now exist to analyze dynamic video data, less is known about the relative performance of these classifiers, in particular when facial expressions are spontaneous rather than posed. In the present work, we tested eight out-of-the-box machine classifiers (Affectiva's Affdex, CrowdEmotion's FaceVideo, Emotient's Facet, Microsoft's Cognitive Services, MorphCast's EmotionalTracking, Neurodatalab's EmotionRecognition, VicarVision's FaceReader and VisageTechnologies' FaceAnalysis), and compared their emotion recognition performance to that of human observers. For this, a total of 938 videos were sampled from two large databases that conveyed the basic six emotions (happiness, sadness, anger, fear, surprise, and disgust) either in posed (BU-4DFE) or spontaneous (UT-Dallas) form. Results revealed a recognition advantage for human observers over machine classification. Among the eight classifiers, there was considerable variance in recognition accuracy ranging from 49 % to 62 %. Subsequent analyses per type of expression revealed that performance by the two best performing classifiers approximated those of human observers, suggesting high agreement for posed expressions. However, accuracy was consistently lower for spontaneous affective behavior. Overall, the recognition of happiness was most successful across classifiers and databases, whereas confusion rates suggested system-specific biases favoring classification of certain expressions over others. The findings indicate shortcomings of existing out-of-the-box classifiers for measuring emotions, and highlight the need for more spontaneous facial datasets that can act as a benchmark in the training and testing of automatic emotion recognition systems.

Keywords: Emotion, Facial Expression, Automatic Recognition, Machine Learning

Word count: 5985

Emotion recognition in humans and machine using posed and spontaneous facial expression

Introduction

Sensing what other people feel is an important element of social interaction. Only if we can perceive the affective state of an individual, will we be able to communicate in a way that corresponds to that experience. In the quest for finding a “window to the soul” that reveals a view onto another’s emotion, the significance of the face has been a focus of popular and scientific interest alike. Since the publication of Charles Darwin’s book *The Expression of the Emotions in Man and Animals* (Darwin, 1872), facial behaviour has been considered to play an integral role in signalling emotional experience. According to Darwin, facial movements became associated with emotions as biological remnants of actions that once served survival-related purposes (Parkinson, 2005). Whilst he did not postulate an intrinsic link between emotions and facial expressions, his work became foundational to the contemporary emotion-expression view of Basic Emotion Theory (BET). Originally proposed by Tomkins (Tomkins, 1962), BET assumes that there are a limited number of emotions (*e.g.*, happiness, sadness, anger, fear, surprise, and disgust) that are characterized by signature expressions (Ekman, 1972; Ekman & Friesen, 1969). The emotions with which these expressions are associated are claimed to be basic, primary, or fundamental in the sense that they form the core emotional repertoire (Ekman, 1972, 1992). Facial behaviour, accordingly, is seen as a “readout” (Buck, 1984) of these subjective feeling states, comprising specific configurations of facial muscle actions that are prototypical, innate, and universal. Although the strength of the support for BET has been challenged within the last decades (*e.g.*, Barrett & Wager, 2006; Fridlund, 1994; Russell & Fernández-Dols, 1997; for a review see Kappas, Krumhuber, & Küster, 2013), this theoretical perspective remains highly influential in the domain of affective computing.

Inspired by the vision of an emotionally intelligent machine, efforts have been targeted towards computer systems that can detect, classify, and interpret human affective states.

This involves the ability to recognize emotional signals that are emitted by the face (Picard & Klein, 2002; Poria, Cambria, Bajpai, & Hussain, 2017), including post-hoc from video recordings as well as in real-time from a camera stream (Gunes & Pantic, 2010). In the wake of rapid advances in computer vision and machine learning, competing computational approaches focus today on the analysis of facial expressions. Automatic facial affect recognition has significant advantages in terms of time and labour costs over human coding (Cohn & De la Torre, 2014) and has been envisioned to increasingly give rise to numerous applications in fields as diverse as security, medicine, education, and telecommunication (Picard, 1997). While the computational modelling of emotional expressions forms a narrow, although increasingly common, approach, the ultimate aim is to build human-computer interfaces that not only detect but also respond to emotional signals of the user (D’mello & Kory, 2015; Schröder et al., 2012). To this end, computer algorithms generally follow three steps in classifying emotions from human facial behaviour. First, they identify and track one or more faces in a video stream based on morphological features and their configuration. Second, they detect facial landmarks and evaluate their changes over time. Finally, they classify the configuration of landmarks according to specific labels, categories, or dimensions (Sariyanidi, Gunes, & Cavallero, 2015). It is within the context of the last step where BET has exerted a profound impact on how expressive behaviour is analysed. To date, most computer models have adopted a BET perspective by focusing on the basic six emotions (Calvo & D’Mello, 2010; Gunes & Hung, 2016). That is, they output a categorical emotion label from a limited set of candidate labels (*i.e.*, happiness, sadness, anger, fear, surprise, and disgust), derived from the assumption that emotional expressions correspond to prototypical patterns of facial activity (Ekman, 1992).

In the last three decades, substantial progress has been made in the area of automated facial expression analysis by recognizing BET’s six categories. Zeng, Pantic, Roisman and Huang (Zeng, Pantic, Roisman, & Huang, 2009), for example, reviewed 29 vision-based affect detection methods, pointing towards the proliferation of software programs and platforms

that are concerned with classifying distinct emotions. As demonstrated by the first Facial Expression Recognition and Analysis (FERA) challenge, emotion recognition by the top performing algorithm already reached in 2011 a rate of 84% (Valstar, Mehu, Jiang, Pantic, & Scherer, 2012). Together with recent news reports that forecast a bright future for emotionally intelligent machines (*e.g.*, Murgia, 2016; Perez, 2018), the impression arises that the automatic inference of basic emotions may soon be a solved problem (Martinez, Valstar, Jiang, & Pantic, 2017). The majority of these past efforts, however, have relied on in-house techniques for facial affect recognition. As such, they involve classification algorithms that have been developed and benchmarked in individual laboratories, often using proprietary databases of emotion-related images and videos; these historically have not been easily accessible for systematic interdisciplinary and cross-laboratory research. Given that automated methods of measuring emotions on the basis of facial expression patterns have now matured, 16 providers of commercially available classifiers were recently identified (Deshmukh & Jagtap, 2017; Dupré, Andelic, Morrison, & McKeown, 2018). Those are marketed for monitoring and evaluating humans' affective states in a range of domains (*e.g.*, for automotive, health, security or marketing purposes). In consequence, their performance can be assessed more freely and openly. Interestingly, the overall relative performance of these software tools is still largely unknown.

In a study by Lewinski, den Uyl and Butler (Lewinski, Uyl, & Butler, 2014), the commercial FaceReader software (VicarVision) was tested on static facial images of posed expressions, achieving a recognition rate of 89%. Using similar sets of static basic emotions stimuli, Stockli, Schulte-Mecklenbeck, Borer and Samson (Stöckli, Schulte-Mecklenbeck, Borer, & Samson, 2018) reported performance indices of 97% and 73% for the software packages Facet (Emotient) and Affdex (Affectiva), respectively. While Facet was found to outperform human judges in classifying emotions on these standardized sets of static emotional portrayals, its accuracy dropped to 63% for dynamic expressions that had been recorded in the laboratory. Such performance drop was also observed for Facet (Calvo,

107 Fernández-Martín, Recio, & Lundqvist, 2018) and for CERT (a precursor of Facet) in the
108 context of dynamic expressions that were subtle and non-prototypical in their facial
109 configuration (21%) rather than highly intense and stereotypical (89%) (Yitzhak et al., 2017).
110 In none of the above studies, however, emotion recognition was tested in spontaneous
111 affective displays.

112   Given that there are fundamental differences between posed and spontaneous stimuli in
113 terms of their appearance and timing (for a review see Calvo & Nummenmaa, 2016),
114 subjecting only deliberately displayed expressions to machine analyses and benchmarking
115 may provide insufficiently robust validation results. In consequence, affective analyses based
116 on deliberate and often prototypical displays are likely to be substantially less reliable with
117 respect to spontaneous expressive behaviour. This issue is further exacerbated by the general
118 trend to train computer algorithms on posed expressions that are highly intense and
119 homogeneous (Pantic & Bartlett, 2007). The third step in automated facial expression
120 analysis typically involves a training set of human-labelled stimuli to make inferences about a
121 much larger population of faces and facial expressions in which they occur (Zeng et al., 2009).
122 Unless a computer system is validated on deliberate as well as spontaneous facial actions, its
123 use in the public and private sector may likely prove inadequate. As the affective computing
124 market is believed to grow considerably, with growth estimations reaching \$41 billion by 2022
125 (Knowledge-Sourcing-Intelligence-LLP, 2017), the time is ripe for a systematic multi-system
126 evaluation of commercial software tools using both types of emotional expressions.

127   The present research aims to fill this gap by testing 8 commercially available machine
128 classifiers and comparing their recognition performance to human observers. To this end,
129 facial stimuli were sampled from two large datasets that depict emotions either in a posed or
130 spontaneous form. All of the examined expressions are dynamic to reflect the realistic nature
131 of human facial behaviour (Krumhuber, Kappas, & Manstead, 2013; Krumhuber & Skora,
132 2016). Following common approaches in the field of affective computing, we focused on the

recognition of the basic six emotions identified by BET. To assess the emotional content of expressions, participants selected the emotion label that best fits with a stimulus (forced choice). We predicted the classification accuracy of posed stimuli to exceed that of spontaneous ones, with generally reduced performance of the machine classifiers compared to human observers in the context of spontaneously occurring expressions. Given the predominance of posed data sets for the training of classifiers, confusion patterns found for machine classification should also be more similar to the human observers when analysing deliberate affective displays. Finally, in line with previous findings (*e.g.*, Palermo & Coltheart, 2004) we hypothesized that happiness achieves the highest overall recognition rate, with mixed results for all other emotional expressions.

Materials and Methods

For the present research, two well-known dynamic facial expression databases were chosen: BU-4DFE (Yin, Chen, Sun, Worm, & Reale, 2008) and UT-Dallas (O'Toole et al., 2005). Both are annotated in terms of emotion categories, and contain either posed or spontaneous facial expressions. To evaluate the accuracy of emotion recognition, we compared the performance achieved by human judges with those of 8 commercially available machine classifiers. To this end, we first conducted a judgment study with naive human observers. Second, we assessed the performance of the machine classifiers on the same databases, and employed standard metrics for all human-machine and classifier-based comparisons.

Stimulus Material

A total of 938 video stimuli were retrieved from both facial expression datasets (Krumhuber, Skora, Küster, & Fou, 2017).

The *BU-4DFE* database contains 468 videos of posed expressions recorded from 78 individuals. They represent male and female subjects of mostly White Caucasian ethnicity,

with an age range of 18-45 years. Each subject was instructed to facially portray the six basic emotions: anger, disgust, fear, happiness, sadness, and surprise. Expression sequences lasted on average 4s ($M = 4.05$, $SD = 0.43$), and started and ended with a neutral face.

The *UT-Dallas* database consists of 470 videos of spontaneous expressions recorded from 148 individuals. They represent male and female subjects of mostly White Caucasian ethnicity, with an age range of 18-25 years. Each subject watched a 10-minute video that included scenes from different movies and television programs intended to elicit distinct emotions. Selected emotive instances were extracted by the database authors, with expressive behaviour corresponding to the six basic emotions: anger, disgust, fear, happiness, sadness, and surprise. Spontaneous clips lasted on average 6s ($M = 6.11$, $SD = 0.68$), and started/ended with a neutral or expressive face.

Besides conceptual differences in elicitation condition and underlying theoretical approach, stimuli from the two databases are similar in the sense that they depict frontal head shots at close distance with comparable expressive intensity envelopes, a static camera view, and adequate illumination. All videos are rendered in colour and captured with a frame-rate of 25 frames per second. While BU-4DFE contains particularly high-resolution video data (1094x1392; UT-Dallas: 720x480), both provide adequate resolution for facial analysis that meets the expected requirements for machine classification (Krumhuber et al., 2017) (see Figure 1).

Human Observers

Fourteen participants (10 females, $M_{age} = 24.0$, $SD = 6.62$), recruited from the academic community in Germany, Turkey, and the UK, volunteered to participate for free or a monetary reward. The study was approved by the departmental ethics committee at University College London, UK. Informed consent was obtained prior to participation. Participants were told that short videos of facial expressions would be presented. Their task

Figure 1. Snapshots of automatic face tracking (green box) and facial landmark detection (yellow feature points) for database stimuli expressing happiness (left: BU-4DFE) and disgust (right: UT-Dallas). Stimulus subjects agreed to the publication of their photographs.

was to indicate the label which best described the displayed expression. They were instructed to watch all 938 videos attentively and with sufficient rest periods. Videos were shown in an individually randomized order and with scrambled file names to avoid guessing of the correct labels.

In line with common categorization paradigms, emotion recognition was assessed through a forced-choice task. This required participants to make a selection among the following emotion labels: *anger, disgust, fear, happiness, sadness, surprise, no/other emotion*. We opted for this response format to allow for direct comparability with the machine recognition data using pre-specified emotion labels.

In addition to the standard classification task, participants were asked to evaluate each video in terms of the perceived genuineness of the expressed emotion, using a 7-point Likert scale (1 -*very posed*, 7 -*very genuine*). An expression was defined as genuine if the person is truly feeling the emotion, in contrast to a posed expression which is simply put onto the face while not much may be felt. Results showed that participants judged posed expressions as significantly less genuine than spontaneous ones (BU-4DFE: $M = 3.42$, $SD = 1.79$;

UT-Dallas: $M = 4.60$, $SD = 1.81$; $t(13, 022.99) = -37.40$, $p < .001$, $d = 0.66$), thereby validating the two different emotion elicitation approaches for database construction.

Machine Classification

The 938 video stimuli (468 BU-4DFE, 470 UT-Dallas) were submitted to automatic facial expression analysis by the following 8 computer-based systems: Affectiva’s Affdex, CrowdEmotion’s FaceVideo, Emotient’s Facet, Microsoft’s Cognitive Services, MorphCast’s EmotionalTracking, Neurodata Lab’s EmotionRecognition, VicarVison’s FaceReader and VisageTechnologies’ FaceAnalysis. All of them offer a prototypical basic emotion approach by classifying facial expressions in terms of the basic six emotions (anger, disgust, fear, happiness, sadness, and surprise) or synonymy.

Affdex (SDK v3.4.1) was developed by Affectiva which is a spin-off company resulting from the research activities of the MIT media lab created in 2009 (Picard, 2011). At present, it is distributed by Affectiva (API and SDK) as well as iMotions (SDK integrated in a software). Affdex’s algorithm uses Histogram of Oriented Gradient (HOG) features and Support Vector Machine (SVM) classifiers for facial expression recognition (McDuff et al., 2016).

FaceVideo (API v1.0) was developed by the company CrowdEmotion founded in 2013. Its algorithm used Convolutional Neural Networks (CNN) and can recognize 6 emotions plus neutral expression.

Facet (SDK v6.3) was originally developed by Emotient and distributed by iMotions. Initially a spin-off company by the University of California San Diego (Bartlett et al., 2005), Emotient was bought by Apple Inc. in 2017 and is available for further use by existing customers until 2020.

Cognitive Services: Face (API v1.0) was developed by the company Microsoft on its

Azure platform and first released in 2015. It provides a suite of artificial intelligence tools for face, speech and text analysis.

EmotionalTracking (SDK v1.0) was developed by the company MorphCast founded in 2013. EmotionalTracking SDK is a javascript engine less than 1MB and works directly on mobile browsers (*i.e.*, without remote server and API processing).

EmotionRecognition (API v1.0) was developed by the company Neurodata Lab founded in 2016. Neurodata Lab provides a suite of tool for emotion recognition/annotation experiments such as face recognition, speaker diarization, body pose estimation, heart rate and respiration rate tracking. Neurodata Lab's EmotionRecognition is available both in API and SDK.

FaceReader (software v7.0) was developed by VicarVison and is now distributed by Noldus (Lewinski et al., 2014). Initially presented in 2005 (Den Uyl & Van Kuilenburg, 2005), the software uses Active Appearance Models (AAM) for face modelling and Convolutional Neural Networks (CNN) for facial expression classification (Gudi, Tasli, Den Uyl, & Maroulis, 2015). All default settings were used for the video processing.

FaceAnalysis (SDK v1.0) was developed by the company Visage Technologies founded in 2002. Visage Technologies provides solutions for facial expression recognition as well as for ID verification using face recognition.

For all computer-based systems, performance indicators as reported in the present research are based on the respective software version indicated above. Results may be subject to change with the release of newer versions. Because the type of output is slightly different for each system, classification scores were rescaled to the odds ratios of recognition probability ranging from 0 to 1.

Data Analysis

The data analysis focuses on a comparison in recognition performance between humans and machine. To ensure that performance measures applied to the machine classifications are analogous to those of human judgments, a common accuracy index was calculated (Sokolova & Lapalme, 2009). For the human data, this index reflects the number of correctly classified videos within an emotion category, divided by the total number of videos per emotion category aggregated across all human judges. For each video, the process to determine the recognised label follows the equation (1)

$$video_i.label = \max \left(\frac{n_{video_i.label_j}}{N_{video_i}} \right)$$

where i is a judged video, j is a category of emotion recognized, n is the number of human annotators choosing the label j and N is the total number of human annotators for the video i .

In the context of the machine data, this index reflects the number of videos within an emotion category for which a given machine classifier correctly indicated the highest recognition confidence score, divided by the total number of videos per emotion category. The recognition confidence score (see Dente, Küster, Skora, & Krumhuber, 2017) corresponds to the sum of the odds ratios for a specific emotion (*e.g.*, happiness) aggregated per video-frame relative to the sum of the odds ratios for all other emotions (*e.g.*, anger, disgust, fear, sadness, surprise) (Dente et al., 2017). For each video, the process to determine the recognised label follows the equation (2)

$$video_i.label = \max \left(\frac{\sum_{t=1}^T x_{t.video_i.label_j} / (1 - x_{t.video_i.label_j})}{\sum_{t=1}^T \sum_{j=1}^J x_{t.video_i.label_j} / (1 - x_{t.video_i.label_j})} \right)$$

where i is a processed video, j is a category of emotion recognized, t corresponds to the

timestamp of the processed video.

By selecting the aggregated maximum confidence score as the indicator for emotion recognition, it is possible that more than one emotion label applies to the same video if they share identical overall confidence scores but this occurred very rarely (Supplementary Figure S1).

Results

Chance was set to a stringent level of 25% (rather than 17% or 1/6) which is considered to be a more conservative test (Frank & Stennett, 2001). Two-sided Student's t-tests reveal that the mean target emotion recognition was significantly higher than chance for human observers ($M = 72.79$, 95% CI [72.03, 73.55], $t(13,061) = 122.73$, $p < .001$, $d = 1.07$), Affectiva's Affdex ($M = 50.70$, 95% CI [47.48, 53.91], $t(932) = 15.69$, $p < .001$, $d = 0.51$), CrowdEmotion's FaceVideo ($M = 48.55$, 95% CI [45.34, 51.77], $t(932) = 14.39$, $p < .001$, $d = 0.47$), Emotient's Facet ($M = 62.17$, 95% CI [59.05, 65.28], $t(932) = 23.40$, $p < .001$, $d = 0.77$), Microsoft's Cognitive Services ($M = 52.84$, 95% CI [49.63, 56.05], $t(932) = 17.03$, $p < .001$, $d = 0.56$), MorphCast's EmotionalTracking ($M = 49.62$, 95% CI [46.41, 52.84], $t(932) = 15.04$, $p < .001$, $d = 0.49$), Neurodata Lab's EmotionRecognition ($M = 57.02$, 95% CI [53.84, 60.20], $t(932) = 19.75$, $p < .001$, $d = 0.65$), VicarVision's FaceReader ($M = 57.56$, 95% CI [54.38, 60.73], $t(932) = 20.11$, $p < .001$, $d = 0.66$), VisageTechnologies' FaceAnalysis ($M = 55.31$, 95% CI [52.11, 58.50], $t(932) = 18.61$, $p < .001$, $d = 0.61$). Hence, accuracy scores by both humans and all machine classifiers exceeded the chance level.

A classifier (9) x expression type (posed *vs.* spontaneous) x emotion (6) ANOVA was conducted to identify differences in performance. We found a significant main effect of classifier on recognition accuracy ($F(8,20418) = 130.47$, $MSE = 1,704.32$, $p < .001$, $\eta^2 = .039$). Pairwise comparisons with Tukey HSD correction (Figure 2, see also Supplementary Table S1 in the Supplementary Materials) showed that human observers

generally outperformed the machine classifiers ($ps < .001$). Facet and FaceReader performed comparably and significantly better than the rest ($ps < .001$).

Figure 2. Mean recognition accuracy of human observers and machine classifiers. Errors bars represent 95% Confidence Interval. Bars not sharing a common letter denote statistically significant differences between classifiers ($p < .05$).

In line with previous research, target emotion recognition was on average superior for posed ($M = 68.31$, $SD = 46.53$) compared to spontaneous expressions ($M = 63.76$, $SD = 48.07$), $F(1, 20418) = 62.20$, $MSE = 1,704.32$, $p < .001$, $\hat{\eta}^2 = .002$. A significant interaction between classifier and expression type ($F(8, 20418) = 12.28$, $MSE = 1,704.32$, $p < .001$, $\hat{\eta}^2 = .004$) further demonstrated that performance depended on the specific classifier considered (Figure 3, see also Supplementary Table S2). While there was a relatively large variance in accuracy across classifiers in the context of posed expressions, results showed a more homogeneous pattern for spontaneous expressions. For both types of expressions, human observers achieved the highest recognition rates, followed by VicarVision's FaceReader and/or Emotient's Facet.

When comparing the accuracy levels of each classifier between posed and spontaneous expressions, human judges performed significantly better on posed displays ($\Delta M = 2.29$, 95% CI [0.71, 3.88], $t(20, 508) = 2.83$, $p = .005$, $d = 0.05$). The same applied to Emotient's Facet ($\Delta M = 11.01$, 95% CI [5.07, 16.96], $t(20, 508) = 3.63$, $p < .001$, $d = 0.23$), Microsoft's Cognitive Services ($\Delta M = 7.82$, 95% CI [1.87, 13.76], $t(20, 508) = 2.58$, $p = .010$, $d = 0.16$), MorphCast's EmotionalTracking ($\Delta M = 9.54$, 95% CI [3.60, 15.48], $t(20, 508) = 3.15$, $p = .002$, $d = 0.19$), Neurodata Lab's EmotionRecognition ($\Delta M = 6.74$, 95% CI [0.79, 12.68], $t(20, 508) = 2.22$, $p = .026$, $d = 0.14$), VicarVision's FaceReader ($\Delta M = 21.53$, 95% CI [15.58, 27.47], $t(20, 508) = 7.10$, $p < .001$, $d = 0.45$) and VisageTechnologies' FaceAnalysis ($\Delta M = 14.89$, 95% CI [8.94, 20.83], $t(20, 508) = 4.91$, $p < .001$, $d = 0.30$). Interestingly, Affectiva's Affdex ($\Delta M = -8.47$, 95% CI [-14.41, -2.53], $t(20, 508) = -2.79$, $p = .005$, $d = 0.17$) showed a performance advantage on spontaneous compared to posed expressions and CrowdEmotion's FaceVideo ($\Delta M = 4.83$, 95% CI [-1.12, 10.77], $t(20, 508) = 1.59$, $p = .111$, $d = 0.10$) did not show a difference between spontaneous and posed expressions.

The ANOVA further showed a significant main effect of emotion (6) on recognition accuracy ($F(5, 20418) = 734.58$, $MSE = 1, 704.32$, $p < .001$, $\hat{\eta}^2 = .136$). Overall, happiness was the best recognized emotion ($M = 86.76$, $SD = 33.89$), followed by sadness ($M = 72.81$, $SD = 44.50$) and surprise ($M = 63.76$, $SD = 48.07$), then anger ($M = 54.99$, $SD = 49.76$) and disgust ($M = 52.03$, $SD = 49.96$), and lastly fear ($M = 39.22$, $SD = 48.84$) (all pairwise comparisons at $ps < .01$, see Supplementary Table S3).

The main effect was qualified by a significant interaction between emotion and classifier ($F(40, 20418) = 24.54$, $MSE = 1, 704.32$, $p < .001$, $\hat{\eta}^2 = .036$). Pairwise comparisons with Tukey HSD correction (Figure 4, see also Supplementary Table S4) showed consistently high performance indices for all classifiers in recognizing happiness, with Microsoft's Cognitive Services achieving the best recognition score. A larger discrepancy between humans and machine was observed in the context of the other emotions. Here, accuracy rates were

Figure 3. Human and machine recognition accuracy separately for posed and spontaneous expressions. Errors bars represent 95% Confidence Interval. Bars not sharing a common letter denote statistically significant differences between classifiers ($p < .05$).

generally high or moderate for human observers, and often superior to the machine ratings.

Across all six basic emotions, we observed variable patterns of correct classifications of the target emotion, as well as confusion errors with non-target emotion categories. As shown in Supplementary Figure S1, confusion rates for human observers in each cell were generally below the 25% chance level, except for spontaneous anger which remained to some extent unidentified as “other” emotion. Whereas anger and disgust as well as fear and surprise were often attributed interchangeably, such confusion errors with other emotions did not apply to happiness. In comparison to humans, the prevalence of confusion with non-target emotions by the 8 computer-based systems was considerably higher, reaching rates of up to 100% for individual cells.

In order to quantify the similarity of confusion errors between the classifiers, each

Figure 4. Human and machine recognition accuracy for the six basic emotions. Errors bars represent 95% Confidence Interval. Bars not sharing a common letter denote statistically significant differences between classifiers ($p < .05$).

matrix was transformed into a single vector and correlations were calculated (Kuhn, Wydell, Lavan, McGettigan, & Garrido, 2017). From inspection of Table 1, there was considerable overlap between humans and the machine classifiers for posed expressions, suggesting that recognition patterns were positively related. For spontaneous expressions, humans and machine displayed much weaker correlations. Also, overall consistency between the computer-based systems tended to be high only for Emotient's Facet and Microsoft's Cognitive Services and between Neurodata Lab's EmotionRecognition, VicarVision's FaceReader and VisageTechnologies's FaceAnalysis. In contrast, correlations was low for Affectiva's Affdex. Together, the results point towards similar patterns of recognition performance for humans and machine in the context of posed expressions, but more classifier-specific patterns for spontaneous expressions.

In addition to assessing the relative performance of the classifiers (Davis & Goadrich, 2006), we computed metrics to obtain a comprehensive picture of generic performance levels. For this, average unweighted True Positive Rates (TPR, Sensitivity), Positive Predicted Values (PPV, Precision), True Negative Rates (TNR, Specificity) and F1 scores were calculated. TPR correspond to the proportion of videos in which the presence of a target emotion is correctly indicated, thereby accounting for the number of misses ($TP/(TP+FN)$). PPV correspond to the proportion of correctly classified videos when accounting for the number of false alarms ($TP/(TP+FP)$). The two metrics which incorporate misses and false alarms together form the basis for an overall harmonic evaluation index called F1 ($2 \times (\text{Sensitivity} \times \text{Precision})/(\text{Sensitivity} + \text{Precision})$). Finally, TNR correspond to the proportion of videos in which the absence of a non-target emotion is correctly indicated.

As detailed in Table 2, there was a clear performance advantage for humans over the machine in terms of sensitivity, precision and F1 for posed expressions. From all computer-based systems, Emotient's Facet and VicarVision's FaceReader achieved the highest scores, thereby approaching human classification levels (Supplementary Figure S1).

Table 1

Correlations of confusion matrices between classifiers separately for posed and spontaneous expressions.

Classifier	id	1	2	3	4	5	6	7	8
Posed Expressions									
Human Observers	1								
Affectiva (Affdex)	2	0.71							
CrowdEmotion (FaceVideo)	3	0.77	0.81						
Emotient (Facet)	4	0.93	0.8	0.91					
Microsoft (Cognitive Services)	5	0.8	0.73	0.82	0.84				
MorphCast (EmotionalTracking)	6	0.94	0.78	0.84	0.93	0.9			
Neurodatalab (EmotionRecognition)	7	0.89	0.8	0.82	0.93	0.89	0.95		
VicarVision (FaceReader)	8	0.96	0.76	0.84	0.97	0.88	0.97	0.97	
VisageTechnologies (FaceAnalysis)	9	0.96	0.84	0.86	0.97	0.92	0.96	0.95	0.99
Spontaneous Expressions									
Human Observers	1								
Affectiva (Affdex)	2	0.4							
CrowdEmotion (FaceVideo)	3	0.38	0.56						
Emotient (Facet)	4	0.58	0.58	0.88					
Microsoft (Cognitive Services)	5	0.38	0.5	0.83	0.84				
MorphCast (EmotionalTracking)	6	0.57	0.45	0.62	0.8	0.65			
Neurodatalab (EmotionRecognition)	7	0.69	0.37	0.73	0.74	0.68	0.71		
VicarVision (FaceReader)	8	0.54	0.58	0.76	0.91	0.83	0.73	0.77	
VisageTechnologies (FaceAnalysis)	9	0.47	0.62	0.8	0.88	0.92	0.65	0.81	0.9

In the context of spontaneous expressions, performance indices were generally lower, except for specificity. Human observers once again exceeded computer-based systems, Affectiva's Affdex being the most accurate of the automatic classifiers; a pattern which was particularly pronounced for measures of precision and F1, and to a moderate extent for sensitivity. Supplementary Table S5 lists performance indices per classifier for posed and spontaneous expressions of all six basic emotions.

To further explore the classifiers' diagnostic ability to discriminate between hits (true positives) and false alarms (false positives), Receiver Operating Characteristic (ROC) curves

Table 2

Average emotion recognition performance indices for each classifier by type of expression. Direct calculation of F1 scores can be obtained from the table presented in Supplementary Materials.

Expressions	Classifier	Avg. Sensitivity	Avg. Precision	Avg. Specificity	Avg. F1
Posed	Human Observers	0.79	0.74	0.95	0.76
	Affectiva (Affdex)	0.50	0.46	0.90	0.45
	CrowdEmotion (FaceVideo)	0.55	0.51	0.91	0.46
	Emotient (Facet)	0.72	0.68	0.94	0.66
	Microsoft (Cognitive Services)	0.66	0.57	0.92	0.51
	MorphCast (EmotionalTracking)	0.52	0.54	0.91	0.53
	Neurodatalab (EmotionRecognition)	0.60	0.60	0.92	0.57
	VicarVision (FaceReader)	0.69	0.68	0.94	0.67
	VisageTechnologies (FaceAnalysis)	0.62	0.63	0.93	0.61
Spontaneous	Human Observers	0.56	0.60	0.94	0.54
	Affectiva (Affdex)	0.36	0.33	0.90	0.40
	CrowdEmotion (FaceVideo)	0.29	0.29	0.88	0.32
	Emotient (Facet)	0.45	0.44	0.91	0.37
	Microsoft (Cognitive Services)	0.47	0.34	0.91	0.26
	MorphCast (EmotionalTracking)	0.37	0.39	0.88	0.30
	Neurodatalab (EmotionRecognition)	0.49	0.53	0.91	0.34
	VicarVision (FaceReader)	0.43	0.37	0.89	0.29
	VisageTechnologies (FaceAnalysis)	0.34	0.36	0.90	0.26

were plotted. The ROC curve and AUC are obtained by comparing the recoded target emotional label as response and the confidence score as predictor for each video. As illustrated by Figure 5, human observers exhibited the overall highest discrimination accuracy, with Area Under the Curve (AUC) values close to 1, thereby visibly outperforming all computer-based systems. The performance of the latter can be described as fair in the context of posed expressions. Interestingly, AUC scores were elevated for most classifiers when expressions were spontaneous. This is also exemplified by the steeper ROC curve in humans, indicating that the ability to discriminate between true and false positives was facilitated by spontaneous affective displays.

Figure 5. Receiver operating characteristic (ROC) curves and corresponding area under the curve (AUC) depicting the true positive rate (TP) against the false positive rate (FP) for humans and machine separately for posed and spontaneous expressions. The dotted diagonal line in the ROC space indicates chance performance.

Discussion

Following recent advances in automatic affect analysis, there has been a proliferation of commercially available software applications designed to recognize human facial expressions. Surprisingly, the number of independent peer-reviewed validation studies for these systems is small and generally limited to deliberately posed displays. The present study aimed to provide a multi-system evaluation of 8 commercial software tools using two types of stimuli: posed expressions arising from instructions to portray a specific emotion, and spontaneous expressions in response to emotion-eliciting events. On the basis of dynamic stimuli sampled

from two large databases, which differed on the described dimension of comparison, results revealed a recognition advantage for human observers over the machine classifiers. The human recognition accuracy of 73% in the present study is consistent with evidence reported in the literature for dynamic expressions (*e.g.*, Bänziger, Mortillaro, & Scherer, 2012; Battocchi, Pianesi, & Goren-Bar, 2005; Recio, Schacht, & Sommer, 2013). Among the 8 classifiers tested in this work, we observed considerable variance in recognition accuracy (ranging from 49% to 62%), wherein Emotient's Facet and VicarVision's FaceReader appeared to outperform the competing systems. As shown by the results per displayed expression type, variability of recognition patterns was particularly pronounced for posed facial behaviour. Similar to past research (Lewinski et al., 2014; Stöckli et al., 2018), performance indices for FaceReader and Facet approximated those of human observers, suggesting high agreement in the classification of posed expressions. However, accuracy scores of most classifiers were consistently lower for spontaneous facial behaviour. This could be due to the lack of prototypicality, that is, greater expressive variability, inherent in spontaneous affective responses. For example, it has been shown that spontaneously elicited facial actions differ in their temporal and morphological characteristics (*e.g.*, duration, intensity, asymmetry) from posed ones (Schmidt, Ambadar, Cohn, & Reed, 2006). Furthermore, the overall patterns of activity are often heterogeneous, which renders them more difficult to discern because of their more ambiguous emotional content (Dawel et al., 2017; Hess & Blairy, 2001; Zloteanu, Krumhuber, & Richardson, 2018). Results based on instructed and stereotypical facial portrayals may therefore not be directly transferable to those derived from activity occurring in spontaneous situations.

This conclusion further appears to be supported by the observed similarity of confusion errors between humans and the machine. While the present results suggested a considerable overlap in the type of confusion errors for posed expressions, these correlations were much weaker in the case of spontaneous expressions. Further analyses showed that the performance indices precision and F1 (both of which take false alarms into account) were markedly

reduced for the machine classifiers compared to humans in the context of spontaneous expressions. When depicting the true positive rate against the false positive rate, we found that discrimination accuracy (*i.e.*, the AUC) was on average lower for all 8 machine classifiers. Thus, the way affective information is extracted by the machine is necessarily and almost certainly not the same as the manner that human observers used to achieve the task (Gunes & Hung, 2016; McDuff, 2016). Such discrepancies can likely be explained by the quality and quantity of data available to train computer-based systems. Although several efforts over the last few years have been targeted towards the automatic analysis of spontaneous displays (*e.g.*, Bartlett et al., 2005; Valstar, Pantic, Ambadar, & Cohn, 2006), most current software tools have typically been trained and tested using posed or acted facial behaviour. Besides their limited ability to transfer to the subtlety and complexity of spontaneous recordings (Pantic & Bartlett, 2007), the highly standardized form of prototypical expressions makes it difficult to generalize beyond specific training sets. At the technical level, the problem of over-fitting is likely to be prevalent, that is, the classifiers may have learned to respond too closely to artificially uniform training sets, thereby losing flexibility when they are applied to unexpectedly subtle and ambiguous expressions. To develop more robust models in the future, it will be important to obtain and train on more comprehensive datasets that display spontaneous and natural behaviour (Martinez et al., 2017). To achieve this aim, metadata in the form of self-reports, behavioural coding (*e.g.*, Ekman & Rosenberg, 1997), and physiological (facial EMG), or neuroscientific measures (EEG, fMRI) are needed to specify the emotional content of recordings. Such annotation of large video sets can help accelerate the progress of affective computing research by providing more naturalistic benchmarks for the training and testing of automated systems on spontaneous expressions.

In the present study, recognition rates for the six basic emotions were generally above chance level, and highest for happiness. This finding applied to human observers as well as to most of the machine classifiers, and is in accordance with prior research (Bänziger et al., 2012; Recio et al., 2013). Variable patterns of correct classifications occurred for all other

emotions, with reasonable classification performance for sadness and surprise, and weaker performance for anger, disgust, and fear. Following standard methods, we adopted a BET perspective by evaluating recognition accuracy on a small set of emotions. According to BET, these six basic categories are reliably communicated by specific configurations of facial expressions (Ekman, 1992; Ekman & Keltner, 1970). The approach fits a classic factorial design structure which requires categorical distinctions in the form of discrete emotions with established labels (*i.e.*, a ground truth). While BET is the most commonly used taxonomy in affective computing, it must be noted that such a perspective is unlikely to reflect the full range of everyday emotions. Typically, emotional behaviour “in the wild” involves a wide variety of affective displays that span a substantial number of emotional states beyond the basic six. Also, one cannot assume a one-to-one correspondence between the experience and expression of emotion (Calvo & D’Mello, 2010). Given that facial expressions fulfil a variety of functions (*e.g.*, appraisals, action tendencies, social motives), it is unlikely that they always signal current emotions in the sense of a “readout” (Kappas et al., 2013; Parkinson, 2005). To account for this complexity, a few tentative efforts in machine analysis have recently started to address non-basic affective and mental states such as interest, pain, boredom, and frustration (Littlewort, Bartlett, & Lee, 2007; Yeasin, Bullot, & Sharma, 2006). By extending the number of emotion categories, we may be able to gain a fuller understanding of the signals and functions of affective phenomena in the future.

Prospective approaches to machine recognition of human affect should further aim to integrate relevant contextual information, as well as learn to better suppress irrelevant information. Both databases used in this work comprised stimuli recorded under relatively controlled conditions, and depicted full frontal shots with neutral backgrounds and steady head poses. While these databases have kept contextual variations across senders constant, information about the wider physical environment and situational factors is likely to be critical to human perception outside the laboratory. Past research, for example, has shown that the same facial expression is interpreted differently depending on the social context in

which it occurs (Barrett & Kensinger, 2010; Wieser & Brosch, 2012). Also, context helps to disambiguate between various exemplars of an emotion category (Aviezer, Ensenberg, & Hassin, 2017). Failures to address the relative role of context may therefore lead to difficulties in generalizing to real-world settings with natural expressions. It will fall into the responsibility of future research to train and test computer systems on more ecologically valid and meaningful materials that are representative of a wider range of emotional and situational contexts. The present study is a first attempt to provide a systematic multi-system evaluation of current commercial software tools using the basic six emotions. By doing so, we hope to help pave the way for the development of more robust machine classifiers in the future.

Acknowledgement

Blank for review

Additional Information

D.D., E.G.K., and D.K. conceived the experiment. E.G.K. and D.K. conducted the experiment. D.D., D.K., and E.G.K. analysed the results. E.G.K., D.K., D.D., and G.M. wrote the manuscript. All authors reviewed the manuscript and approved the submitted version.

The authors declare no competing interests.

Data Availability

The code written in R resulting of the present paper as well as the data used are available in the following GitHub repository <https://github.com/blank-for-review> (set as private until publication acceptance).

References

- Aviezer, H., Ensenberg, N., & Hassin, R. R. (2017). The inherently contextualized nature of facial emotion perception. *Current Opinion in Psychology*, 17, 47–54.
- Barrett, L. F., & Kensinger, E. A. (2010). Context is routinely encoded during emotion perception. *Psychological Science*, 21(4), 595–599.
- Barrett, L. F., & Wager, T. D. (2006). The structure of emotion: Evidence from neuroimaging studies. *Current Directions in Psychological Science*, 15(2), 79–83.
- Bartlett, M. S., Littlewort, G., Frank, M., Lainscsek, C., Fasel, I., & Movellan, J. (2005). Recognizing facial expression: Machine learning and application to spontaneous behavior. In *Proceedings of the conference on computer vision and pattern recognition* (Vol. 2, pp. 568–573).
- Battocchi, A., Pianesi, F., & Goren-Bar, D. (2005). A first evaluation study of a database of kinetic facial expressions (dafex). In *Proceedings of the international conference on multimodal interfaces* (pp. 214–221).
- Bänziger, T., Mortillaro, M., & Scherer, K. R. (2012). Introducing the geneva multimodal expression corpus for experimental research on emotion perception. *Emotion*, 12(5), 1161–1179.
- Buck, R. (1984). *The communication of emotion*. Guilford Press.
- Calvo, M. G., Fernández-Martín, A., Recio, G., & Lundqvist, D. (2018). Human observers and automated assessment of dynamic emotional facial expressions: KDEF-dyn database validation. *Frontiers in Psychology*, 9, 1–12.
- Calvo, M. G., & Nummenmaa, L. (2016). Perceptual and affective mechanisms in facial expression recognition: An integrative review. *Cognition and Emotion*, 30(6), 1081–1106.

Calvo, R. A., & D'Mello, S. (2010). Affect detection: An interdisciplinary review of models, methods, and their applications. *IEEE Transactions on Affective Computing*, 1(1), 18–37.

Cohn, J. F., & De la Torre, F. (2014). Automated face analysis for affective. In R. Calvo, S. D'Mello, J. Gratch, & A. Kappas (Eds.), *The oxford handbook of affective computing* (pp. 131–150). New York, NY: Oxford University Press.

Darwin, C. (1872). *The expression of the emotions in man and animals*. London, UK: John Murray.

Davis, J., & Goadrich, M. (2006). The relationship between precision-recall and roc curves. In *Proceedings of the international conference on machine learning* (pp. 233–240).

Dawel, A., Wright, L., Irons, J., Dumbleton, R., Palermo, R., O'Kearney, R., & McKone, E. (2017). Perceived emotion genuineness: Normative ratings for popular facial expression stimuli and the development of perceived-as-genuine and perceived-as-fake sets. *Behavior Research Methods*, 49(4), 1539–1562.

Dente, P., Küster, D., Skora, L., & Krumhuber, E. (2017). Measures and metrics for automatic emotion classification via facet. In *Proceedings of the conference on the study of artificial intelligence and simulation of behaviour* (pp. 160–163).

Den Uyl, M., & Van Kuilenburg, H. (2005). The facereader: Online facial expression recognition. In *Proceedings of the international conference on methods and techniques in behavioral research* (pp. 589–590).

Deshmukh, R. S., & Jagtap, V. (2017). A survey: Software api and database for emotion recognition. In *Proceedings of the international conference on intelligent computing and control systems* (pp. 284–289).

D'mello, S. K., & Kory, J. (2015). A review and meta-analysis of multimodal affect detection systems. *ACM Computing Surveys (CSUR)*, 47(3), 43.

Dupré, D., Andelic, N., Morrison, G., & McKeown, G. (2018). Accuracy of three commercial automatic emotion recognition systems across different individuals and their facial expressions. In *Proceedings of the international conference on pervasive computing and communications* (pp. 627–632).

Ekman, P. (1972). Universal and cultural differences in facial expression of emotion. In J. Cole (Ed.), *Nebraska symposium on motivation* (Vol. 19, pp. 207–284). Lincoln, NE: University of Nebraska Press.

Ekman, P. (1992). An argument for basic emotions. *Cognition & Emotion*, 6(3-4), 169–200.

Ekman, P., & Friesen, W. V. (1969). The repertoire of nonverbal behavior: Categories, origins, usage, and coding. *Semiotica*, 1(1), 49–98.

Ekman, P., & Keltner, D. (1970). Universal facial expressions of emotion. *California Mental Health Research Digest*, 8(4), 151–158.

Ekman, P., & Rosenberg, E. L. (1997). *What the face reveals: Basic and applied studies of spontaneous expression using the facial action coding system (facs)*. New York, NY: Oxford University Press.

Frank, M. G., & Stennett, J. (2001). The forced-choice paradigm and the perception of facial expressions of emotion. *Journal of Personality and Social Psychology*, 80(1), 75.

Fridlund, A. J. (1994). *Human facial expression: An evolutionary view*. San Diego, CA: Academic Press.

Gudi, A., Tasli, H. E., Den Uyl, T. M., & Maroulis, A. (2015). Deep learning based

562 face action unit occurrence and intensity estimation. In *Proceedings of the international*
563 *conference on automatic face and gesture recognition* (pp. 1–5).

564 Gunes, H., & Hung, H. (2016). Is automatic facial expression recognition of emotions
565 coming to a dead end? The rise of the new kids on the block. *Image and Vision Computing*,
566 55(1), 6–8.

567 Gunes, H., & Pantic, M. (2010). Automatic, dimensional and continuous emotion
568 recognition. *International Journal of Synthetic Emotions*, 1(1), 68–99.

569 Hess, U., & Blairy, S. (2001). Facial mimicry and emotional contagion to dynamic
570 emotional facial expressions and their influence on decoding accuracy. *International Journal*
571 *of Psychophysiology*, 40(2), 129–141.

572 Kappas, A., Krumhuber, E., & Küster, D. (2013). Facial behavior. In J. A. Hall & M.
573 L. Knapp (Eds.), *Handbook of communication science* (Vol. 2, pp. 131–166). New York, NY:
574 Mouton de Gruyter.

575 Knowledge-Sourcing-Intelligence-LLP. (2017). *Global affective computing market -*
576 *forecasts from 2017 to 2022*. Retrieved May 1, 2019, from [https://www.researchandmarkets.](https://www.researchandmarkets.com/reports/4396321/global-affective-computing-market-forecasts)
577 [com/reports/4396321/global-affective-computing-market-forecasts](https://www.researchandmarkets.com/reports/4396321/global-affective-computing-market-forecasts)

578 Krumhuber, E. G., Kappas, A., & Manstead, A. S. (2013). Effects of dynamic aspects
579 of facial expressions: A review. *Emotion Review*, 5(1), 41–46.

580 Krumhuber, E. G., & Skora, L. (2016). Perceptual study on facial expressions. In B.
581 Müller & S. Wolf (Eds.), *Handbook of human motion* (pp. 1–15). Springer.

582 Krumhuber, E. G., Skora, L., Küster, D., & Fou, L. (2017). A review of dynamic
583 datasets for facial expression research. *Emotion Review*, 9(3), 280–292.

584 Kuhn, L. K., Wydell, T., Lavan, N., McGettigan, C., & Garrido, L. (2017). Similar

representations of emotions across faces and voices. *Emotion*, 17(6), 912–937.

Lewinski, P., Uyl, T. M. den, & Butler, C. (2014). Automated facial coding: Validation of basic emotions and faces in facereader. *Journal of Neuroscience, Psychology, and Economics*, 7(4), 227–236.

Littlewort, G. C., Bartlett, M. S., & Lee, K. (2007). Faces of pain: Automated measurement of spontaneous facial expressions of genuine and posed pain. In *Proceedings of the international conference on multimodal interfaces* (pp. 15–21).

Martinez, B., Valstar, M. F., Jiang, B., & Pantic, M. (2017). Automatic analysis of facial actions: A survey. *IEEE Transactions on Affective Computing*, 1–1.

McDuff, D. (2016). Discovering facial expressions for states of amused, persuaded, informed, sentimental and inspired. In *Proceedings of the international conference on multimodal interaction* (pp. 71–75).

McDuff, D., Mahmoud, A., Mavadati, M., Amr, M., Turcot, J., & Kaliouby, R. el. (2016). AFFDEX sdk: A cross-platform real-time multi-face expression recognition toolkit. In *Proceedings of the conference on human factors in computing systems* (pp. 3723–3726).

Murgia, M. (2016). *Affective computing: How “emotional machines” are about to take over our lives*. *The Daily Telegraph*. Retrieved May 1, 2019, from <https://www.telegraph.co.uk/technology/news/12100629/Affective-computing-how-emotional-machines-are-about-to-take-over-our-lives.html>

O’Toole, A. J., Harms, J., Snow, S. L., Hurst, D. R., Pappas, M. R., Ayyad, J. H., & Abdi, H. (2005). A video database of moving faces and people. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(5), 812–816.

Palermo, R., & Coltheart, M. (2004). Photographs of facial expression: Accuracy,

response times, and ratings of intensity. *Behavior Research Methods, Instruments, & Computers*, 36(4), 634–638.

Pantic, M., & Bartlett, M. S. (2007). Machine analysis of facial expressions. In K. Delac & M. Grgic (Eds.), *Face recognition* (pp. 377–416). InTech.

Parkinson, B. (2005). Do facial movements express emotions or communicate motives? *Personality and Social Psychology Review*, 9(4), 278–311.

Perez, A. (2018). *Recognizing human facial expressions with machine learning*. Retrieved May 1, 2019, from <https://www.thoughtworks.com/insights/blog/recognizing-human-facial-expressions-machine-learning>

Picard, R. W. (1997). *Affective computing*. Boston, MA: MIT Press.

Picard, R. W. (2011). Measuring affect in the wild. In *International conference on affective computing and intelligent interaction* (pp. 3–3).

Picard, R. W., & Klein, J. (2002). Computers that recognise and respond to user emotion: Theoretical and practical implications. *Interacting with Computers*, 14(2), 141–169.

Poria, S., Cambria, E., Bajpai, R., & Hussain, A. (2017). A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion*, 37, 98–125.

Recio, G., Schacht, A., & Sommer, W. (2013). Classification of dynamic facial expressions of emotion presented briefly. *Cognition & Emotion*, 27(8), 1486–1494.

Russell, J. A., & Fernández-Dols, J.-M. (1997). What does facial expression mean? In J. A. Russell & J.-M. Fernández-Dols (Eds.), *The psychology of facial expression* (pp. 3–30). New York, NY: Cambridge University Press.

Sariyanidi, E., Gunes, H., & Cavallero, A. (2015). Automatic analysis of facial affect:

A survey of registration, representation, and recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(6), 1113–1133.

Schmidt, K. L., Ambadar, Z., Cohn, J. F., & Reed, L. I. (2006). Movement differences between deliberate and spontaneous facial expressions: Zygomaticus major action in smiling. *Journal of Nonverbal Behavior*, 30(1), 37–52.

Schröder, M., Bevacqua, E., Cowie, R., Eyben, F., Gunes, H., Heylen, D., & Pelachaud, C. (2012). Automatic analysis of facial affect: A survey of registration, representation, and recognition. *IEEE Transactions on Affective Computing*, 3(2), 165–183.

Sokolova, M., & Lapalme, G. (2009). A systematic analysis of performance measures for classification tasks. *Information Processing & Management*, 45(4), 427–437.

Stöckli, S., Schulte-Mecklenbeck, M., Borer, S., & Samson, A. C. (2018). Facial expression analysis with affdex and facet: A validation study. *Behavior Research Methods*, 50(4), 1446–1460.

Tomkins, S. S. (1962). *Affect, imagery, consciousness: Vol 1. Positive affects*. New York, NY: Springer.

Valstar, M. F., Mehu, M., Jiang, B., Pantic, M., & Scherer, K. (2012). Meta-analysis of the first facial expression recognition challenge. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 42(4), 966–979.

Valstar, M. F., Pantic, M., Ambadar, Z., & Cohn, J. F. (2006). Spontaneous vs. Posed facial behavior: Automatic analysis of brow actions. In *Proceedings of the international conference on multimodal interfaces* (pp. 162–170).

Wieser, M. J., & Brosch, T. (2012). Faces in context: A review and systematization of contextual influences on affective face processing. *Frontiers in Psychology*, 3, 471–471.

653 Yeasin, M., Bulot, B., & Sharma, R. (2006). Recognition of facial expressions and
654 measurement of levels of interest from video. *IEEE Transactions on Multimedia*, 8(3),
655 500–508.

656 Yin, L., Chen, X., Sun, Y., Worm, T., & Reale, M. (2008). A high-resolution 3D
657 dynamic facial expression database. In *Proceedings of the international conference on*
658 *automatic face and gesture recognition* (pp. 1–6).

659 Yitzhak, N., Giladi, N., Gurevich, T., Messinger, D. S., Prince, E. B., Martin, K., &
660 Aviezer, H. (2017). Gently does it: Humans outperform a software classifier in recognizing
661 subtle, nonstereotypical facial expressions. *Emotion*, 17(8), 1187–1198.

662 Zeng, Z., Pantic, M., Roisman, G. I., & Huang, T. S. (2009). A survey of affect
663 recognition methods: Audio, visual, and spontaneous expressions. *IEEE Transactions on*
664 *Pattern Analysis and Machine Intelligence*, 31(1), 39–58.

665 Zloteanu, M., Krumhuber, E. G., & Richardson, D. C. (2018). Detecting genuine and
666 deliberate displays of surprise in static and dynamic faces. *Frontiers in Psychology*, 9(1184).