

Published in final edited form as:

*Curr Biol.* 2014 March 31; 24(7): 738–743. doi:10.1016/j.cub.2014.02.009.

## Automatic decoding of facial movements reveals deceptive pain expressions

Marian Stewart Bartlett<sup>1</sup>, Gwen C. Littlewort<sup>1</sup>, Mark G. Frank<sup>2</sup>, and Kang Lee<sup>3</sup>

<sup>1</sup>Institute for Neural Computation, University of California, San Diego, 9500 Gilman Drive, MC 0440, La Jolla, CA 92093-0440 USA

<sup>2</sup>Department of Communication, University at Buffalo, State University of New York, Baldy Hall, Buffalo, New York 14260-1020 USA

<sup>3</sup>Dr. Erick Jackman Institute of Child Study, University of Toronto, 45 Walmer Road, Toronto, Ontario M5R 2X2 Canada

### Summary

In highly social species such as humans, faces have evolved to convey rich information for social interaction, including expressions of emotions and pain [1–3]. Two motor pathways control facial movement [4–7]. A subcortical extrapyramidal motor system drives spontaneous facial expressions of felt emotions. A cortical pyramidal motor system controls voluntary facial expressions. The pyramidal system enables humans to simulate facial expressions of emotions not actually experienced. Their simulation is so successful that they can deceive most observers [8–11]. Machine vision may, however, be able to distinguish deceptive from genuine facial signals by identifying the subtle differences between pyramidally and extrapyramidally driven movements. Here we show that human observers could not discriminate real from faked expressions of pain better than chance, and after training, improved accuracy to a modest 55%. However a computer vision system that automatically measures facial movements and performs pattern recognition on those movements attained 85% accuracy. The machine system's superiority is attributable to its ability to differentiate the dynamics of genuine from faked expressions. Thus by revealing the dynamics of facial action through machine vision systems, our approach has the potential to elucidate behavioral fingerprints of neural control systems involved in emotional signaling.

---

© 2014 Elsevier Inc. All rights reserved.

Corresponding Authors: Marian S Bartlett Ph.D., Research Professor, Institute for Neural Computation, University of California, San Diego, 9500 Gilman Drive, MC 0440, La Jolla, CA 92093-0440, USA, mbartlett@ucsd.edu, Phone: 858-395-3433, Fax: 858-822-5242, Kang Lee Ph.D., Professor, Dr. Eric Jackman Institute of Child Study, University of Toronto, 45 Walmer Road, Toronto, Ontario M5R 2X2, Canada, kang.lee@utoronto.ca, 416 934-4597.

**Publisher's Disclaimer:** This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

### Supplemental Data

Supplemental Data include technique descriptions of the human experimental procedures and the computational methods including the video analysis process using CERT, and machine learning and cross-validation procedures, and can be found with this article online at <http://www.cell.com/current-biology/supplemental/XXX>

## Results

### Human experiments

To test both human observers' and our computer vision system's ability to discriminate real vs. faked emotional expressions we created two sets of videos. One set contained faces of individuals while experiencing genuine pain as induced through a 'cold pressor' method [12], whereas the other contained faces of the same individuals pretending to be in pain. Expressions of pain were chosen because pain is a universally experienced emotive-physiological state [12–15]. Additionally, both genuine and faked expressions of pain can be readily elicited using the "Cold Pressor" method, a routine experimental procedure used to induce pain for research purposes [12]. Stimulus subjects either experienced genuine pain while submerging their arm in ice water (5 degree C) for 1 minute, or were instructed to fake pain while submerging their arm in warm water (20 degree C) for 1 minute. Facial expressions in the two conditions were video-recorded.

In Experiment 1, we showed 170 human observers video clips of the stimulus subjects individually in a randomized order. The observers judged whether the expression shown in the video clip was real or faked. The observers distinguished genuine from faked pain at rates no greater than guessing ( $M$  accuracy = 51.9%;  $SD=14.6$ ; chance accuracy = 50%).

Experiment 2 examined whether training could improve human observers' detection accuracy. Thirty-five new participants were shown 24 video pairs in a training procedure to match the cross-validation training of the computer vision system described below. Observers were presented with two videos of the same person shown sequentially. In one video the individual was expressing genuine pain, and in the other, faked pain. Observers then judged which video of the pair was the genuine pain or which video was the faked pain. Observers received feedback about their accuracy immediately. After being trained on all 24 pairs, participants saw, in random order, 20 new videos of 20 new stimulus subjects for the test phase. Half of these new videos displayed faked pain and the other half displayed real pain. Observers judged whether the expression shown in each of the 20 videos was real or faked, with no feedback offered. This test phase assessed whether human observers could generalize what they had learned to detect new exemplars of genuine or faked pain expressions. In the first third of the training trials (8 trials), the accuracy was 49.6% ( $SD=11\%$ ). The accuracy rate for the last third of the training trials was 58.6%, ( $SD = 8.5\%$ ) which was significantly above chance ( $t(34) = 2.45, p < .01$ ), and showed a significant albeit small improvement over earlier training trial blocks ( $t(34) = 2.22, p < .02$ ). In the test phase detection accuracy remained just above chance level at 54.6% ( $SD=15.5\%$ ),  $t(34)=1.76, p < .05$ . Thus, results from the two human experiments together suggest that human observers are generally poor at detecting differences between real and faked pain. There was a small improvement with training, but performance remained below 60%. This result is highly consistent with prior research [14].

### Machine learning

We then presented these videos to a computer vision system called the Computer Expression Recognition Toolbox (CERT). CERT is a fully automated system that analyzes facial

muscle movements from video in real-time [16]. It automatically detects frontal faces in video and codes each frame with respect to a set of continuous dimensions, including facial muscular actions from the Facial Action Coding System (FACS) [17]. FACS is a system for objectively scoring facial expressions in terms of elemental facial movements, called action units (AUs). This makes FACS fully comprehensive given its basis in functional neuroanatomical actions. CERT can identify 20 out of a complete set of 46 AUs, each with their own movement and appearance characteristics (see Figure 1). FACS was originally developed for manual coding by human experts. Manual coding is laborious, and can take up to 3 hours to manually code 1 minute of behavior, but CERT instantaneously outputs facial movement information in real time (i.e., every 1/30 seconds). Furthermore, the frame-by-frame CERT output provides information on facial expression intensity and dynamics at temporal resolutions that were previously impractical with human coding. CERT was developed at University of California, San Diego and is presently available at Emotient, Inc.

We used a pattern recognition approach to assess CERT's ability to detect falsified pain expressions (see Figure 2). The 60-second videos were input to the computer vision system one at a time. A set of dynamic descriptors was extracted from the output for each of the 20 AUs. The descriptors consisted of 'bags of temporal features' (Figure 3). (See supplement). Two sets of descriptors were employed: One set that describes the dynamics of facial movement events (event descriptors), and another set that describes the intervals between events (interval descriptors). Our methods for constructing bags of temporal features represent a novel approach that can be applied generally to describe signal dynamics for other pattern recognition problems. They build up on the concept of 'bags of features' to provide sensitivity to some aspects of the signal (such as edges or peaks at different scales) while providing invariance to aspects of the signal across which we wish to generalize, such as the specific location or time point at which the signal occurs.

Next, a classifier was trained to discriminate real from faked pain using these descriptors. The classifier was a support vector machine (SVM). The SVM combined information from multiple AU's. This was accomplished by a sequential feature selection procedure. In this procedure, the model began with the AU that gave the best individual classification accuracy. We then added to the model the AU that gave the best performance when combined with the previously selected AU's. This process was repeated until detection performance stopped improving. The selected AU's were the event descriptors for brow lower (AU 4), lip raise (AU 10), lip corner depress (AU 15), and lower lid tighten (AU 7), and the combined event and interval descriptors for mouth opening (AU 26). Because of the post-hoc nature of sequential feature selection, this model was tested with double cross validation (see procedures). Like the human judges, the computer system was tested on each video individually without the ability to compare real and faked data from a given stimulus subject. The system achieved  $A'=.91$  for detection of real vs. faked pain. This corresponds to 85% correct classification at equal error rate (PC), and is significantly better than chance ( $Z=3.7$ ,  $p<.0001$ ) and outperforms human judges even after training ( $Z=3.3$ ,  $p<.001$ , 2-tailed).

## In-depth analysis of computer vision accuracy: Static vs. Dynamic

To test the contribution of temporal information to classification accuracy, we examined the effect of the size of the time window on classification performance. We re-trained the above 5-AU system on sliding time windows of varying length. Classification performance was tested for temporal integration window sizes ranging from 100 frames (3.3 seconds) to 60 seconds. Figure 4 shows that the classification accuracy increased systematically with time.

To ascertain whether this time-related increase in accuracy was due to the nature of the temporal dynamics rather than the simple accumulation of information, an SVM classifier with the 5 selected AU's was trained on the same frame-by-frame CERT outputs but with the order of the frames randomized. It did not perform at above chance levels ( $A' = .58$ ,  $PC = 56\%$ ,  $Z=0.6$ ,  $p=.27$ ). A static system was also tested that did not extract temporal features but instead measured facial actions in each video frame, histogrammed them, and then passed them to an SVM. The static system was able to detect faked expressions of pain marginally above chance levels ( $A'=.72$ ,  $PC = 66\%$ ,  $Z=1.6$   $p<.05$ ). However, this 66% accuracy rate for the static system was significantly poorer than the accuracy rate for the dynamic system reported above (85%;  $Z=2.2$ ,  $p<.05$ ). These findings show that dynamic temporal information contributes substantially to the system's performance.

## In-depth analysis of computer vision accuracy: Most important Action

### Units

Next we attempted to identify the key features for differentiating real from faked pain. The feature selection procedure identified which facial actions *in combination* provided the most information for detecting faked pain. We next examined which *individual* action units can reliably differentiate real from faked pain. Twenty Support Vector Machine (SVM) classifiers, one per action unit, were individually trained to differentiate genuine from faked pain using the dynamic descriptors of a single action unit. The most informative AU for differentiating real from faked pain was the mouth opening (AU 26) using both the event and interval descriptors ( $A' = .85$ ,  $pc=72\%$ ,  $Z=2.3$ ,  $p<.01$ ). Three other AUs individually differentiated genuine from faked pain at above chance levels using just the event descriptors: Lip raise (AU 10), lip press (AU 24), and brow lower (corrugator muscle; AU 4),  $A' = .75, .73, .73$ , respectively ( $pc = 68\%, 66\%, 66\%$ ;  $Z=1.8, 1.7, 1.7$ ;  $p< .05$  respectively). All other AU's were at chance.

Mouth opening was the single most informative feature for discriminating genuine from faked expressions of pain. This feature contained dynamic information about mouth opening as well as the intervals between mouth openings. This finding led us to explore how mouth opening dynamics differ in genuine versus faked expressions of pain using some simple statistics on the unfiltered CERT output. First, there was no difference in the overall mean CERT output for mouth opening between real vs. faked expressions of pain ( $t(24) = 0.006$ ,  $p = .99$ ), implying that the crucial information was in the dynamics. A measure of the duration of mouth openings,  $\tau$ , was then extracted (see Procedures), as well as an estimate of the temporal intervals between mouth openings. There was a difference in the mean duration of mouth openings for genuine and faked expressions, with faked expressions being 5.4

frames shorter than genuine expressions on average, and the interval between mouth openings lasting 11.5 frames less for faked expressions on average,  $t(24) = 2.23$  and  $t(24) = 2.19$  respectively, both  $p < .05$ . The variance of  $\tau$  was then computed for faked and genuine expressions. A within-subjects comparison showed that the variance of  $\tau$  was 55% less for faked than genuine expressions of pain, ( $t(24) = 2.7$ ,  $p < .01$ ). Similarly, the variance of the interval length between mouth openings was 56% less for faked than genuine expressions of pain,  $t(24) = 2.11$ ,  $p < .05$ .

## Discussion

We show for the first time that a fully automated computer vision system can be trained to detect a deceptive emotional-physiological state, faked expressions of pain, from facial cues alone. The rate of accurate discrimination by the computer vision system was 85%. This is far superior to the accuracy of human observers, regardless of whether they have received training, which is consistently below 60% accuracy. This is a significant milestone for machine vision systems [18] because although computers have long outperformed humans at logic processes (such as playing chess), they have significantly underperformed compared to humans at perceptual processes, rarely reaching even the level of a human child.

Furthermore, our computer vision approach has led to the discovery of new information about facial behavior in falsified pain expressions. The single most predictive feature of falsified expressions of pain is the dynamics of the mouth opening, which alone could differentiate genuine from deceptive expressions of pain at a detection rate of 85%. Faked expressions were associated with a reduction in variance in terms of both the duration of mouth openings, and the duration of the interval between mouth openings. In other words, the action was repeated at intervals that were too regular. The critical feature for faked pain may be this over-regularity of the dynamics of the mouth opening action. Further investigations will explore whether over-regularity is a general feature of faked expressions.

Our findings further support previous research on human facial expressions which has shown that the dynamics of expression are important distinguishing characteristics of emotional expressions, such as the genuine smile versus a faked smile [8]. This difference stems from the fact that expressions are mediated by two distinct neural systems, each one originating in a different area of the brain [4–7]. A genuinely felt or experienced emotion originates in the subcortical areas of the brain, and is involuntarily propelled onto the face via the extrapyramidal motor system [4–7]. In contrast, posed or faked expressions originate in the cortical motor strip, and are voluntarily expressed in the face via the pyramidal motor system. Research documenting these differences was sufficiently reliable to become the primary diagnostic criteria for certain brain lesions prior to modern imaging methods (e.g [4,6–7]). These two systems may correspond to the distinction between biologically driven versus socially learned facial behavior [8]. The facial expressions mediated by these two pathways have been shown to differ in some dynamic properties. Extrapyramidal motor system based expressions have been associated with synchronized, smooth, symmetrical, and ballistic-like facial muscle movements, whereas pyramidal motor system based expressions are subject to volitional real-time control and tend to be less smooth, less synchronized, and less symmetric [3]. Accordingly, smiles that were spontaneously

generated have been shown to have smoother dynamics than smiles that are posed or faked, as well as faster onset and offset velocity [3,8]). Here we show a new difference in variance between the two systems. Pyramidally driven expression of falsified pain showed a reduced variance in the timing of mouth openings relative to the spontaneous expressions of pain driven by the extrapyramidal system. A survey study revealed that humans have an intuitive knowledge of differences between controlled and automatic responses to pain [15]. However, our findings show that despite this understanding, people could not detect differences between controlled and automatic facial responses when presented with them visually.

In highly social species such as humans and other primates, the face has evolved to convey a rich array of information for social interaction. Although facial expressions are mainly evolved as cooperative social signals to communicate one's genuinely felt emotions to others, and hence behavioral intentions [1], sometimes individuals may wish to control their expressions to mislead. Indeed, deceptions are a part of everyday life [8, 9, 21] and there are considerable adaptive advantages to deliberately manipulating, suppressing, and dissembling emotional expressions, including social acceptance [9]. Such voluntary facial control may have been refined for adaptive purposes, to be polite, to facilitate interaction, so much so as to make it very difficult for observers to discern honest signals from controlled or falsified ones. In studies of deception, untrained human judges are typically only accurate at or near chance levels when detecting deceptive facial behaviors [11]. This inaccuracy persists despite the fact that (albeit imperfect) diagnostic signals exist [12]. In some domains genuine and faked expressions of emotion have not only shown morphological differences, but also dynamic differences. While human judges were better than chance at detecting these morphological markers, they were unable to detect spatiotemporal dynamic markers [8].

Specifically with regard to pain, lay adults and even experienced physicians cannot reliably differentiate real from faked expressions of pain [13,14, 22–24]. As shown in Experiment 2 and by others [13], immediate feedback might enable perceptual learning and improve detection accuracy to above chance levels. However, accuracy remains modest. Previous research using a laborious manual coding method has shown that there is no telltale facial action that can indicate faked pain by its presence or absence as real and faked expressions of pain include the same set of facial actions. However, these earlier studies hinted differences in their dynamics [14]. In the current study, the computer vision system was able to analyze facial expression dynamics at a much higher temporal resolution, and with richer description, than was feasible with manual coding methods. Thus it revealed aspects of a pain expression that have been previously unavailable to observers.

Our findings taken together suggest that in spite of the pyramidal motor system's sophisticated voluntary control over facial expressions, its control is imperfect; the system cannot fully replicate the genuine expressions of pain driven by the extrapyramidal motor system, particularly in their dynamics. Thus, our findings support the hypothesis that information exists, particularly in facial dynamics, that can differentiate experienced spontaneous expressions of emotion driven by the extrapyramidal motor system from posed or falsified expressions controlled by the pyramidal motor system [4]. Although the present



study addressed one psychophysiological state - pain - the approach presented here may be generalizable to the comparison of other genuine and faked emotional states, which may differentially activate the cortical and subcortical facial motor pathways. Thus, our automated facial movement coding system provides a new paradigm for the study of facial dynamics, and has the potential to elucidate behavioral fingerprints of neural control systems involved in emotional signaling.

There are some practical implications of the present findings. Falsified pain can be a lie told by patients to their physicians for insurance fraud or to receive unneeded prescription narcotics. Some healthcare professionals perceive such lies to be common [22], while in others perceive them to be relatively rare. Our findings suggest that it might be possible to train physicians to specifically attend to mouth opening dynamics to improve their ability to differentiate real from faked pain. In addition to detecting pain malingering, our computer vision approach may be used to detect other real world deceptive actions in the realm of homeland security, psychopathology, job screening, medicine, and law. Like pain, these scenarios also generate strong emotions, along with attempts to minimize, mask, and fake such emotions [21], which may involve dual control of the face. In addition, our computer vision system can be applied to detect states in which the human face may provide important clues as to health, physiology, emotion, or thought such as drivers' expressions of sleepiness [16, 18], students' expressions of attention and comprehension of lectures [18], or to track response to treatment of affective disorders [19].

Several limitations to our findings should be noted. First, the pain manipulation task – cold pressor – is a good, but not perfect analogue to all varieties of clinical pain. Thus future research will be needed to use CERT to analyze expressions of various pain experiences (e.g., back pain) collected in clinical settings. Second, pain is a complicated concept involving attitudes, movements elsewhere in the body, and so forth [25]. This paper addresses just one element - facial expression – and shows proof of principle. Future studies will need to address the other elements of the pain phenomenon.

## Main findings and implications

In summary, the present study demonstrated the effectiveness of a computer vision and pattern recognition system for detecting faked pain from genuine expressions. The computer system outperformed human observers, achieving significantly better accuracy. Moreover, the automated system revealed new information about facial dynamics that differentiate real from faked expressions of pain. Our findings demonstrated the ability of the computer system to extract information from spatiotemporal facial expression signals that humans either cannot or do not. Automated systems such as CERT may bring about a paradigm shift by analyzing facial behavior at temporal resolutions previously not feasible with manual coding methods. This novel approach has succeeded in illuminating basic questions pertaining to the many social situations in which the behavioral fingerprint of neural control systems may be relevant.

## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

## Acknowledgments

Support for this work was provided by NSF grants SBE-0542013, CNS-0454233, NSF ADVANCE award 0340851, NIH grants R01 HD047290 and NR013500 and NSFC 30528027 and 31028010. Marian Bartlett and Gwen Littlewort are founders, employees, and shareholders of Emotient Inc., a company which may potentially benefit from the research results. The terms of this arrangement have been reviewed and approved by the University of California, San Diego in accordance with its conflict of interest policies. Mark G. Frank and the University at Buffalo are minority owners of the software used in this paper and may potentially benefit from this research. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or National Institute of Health.

## References

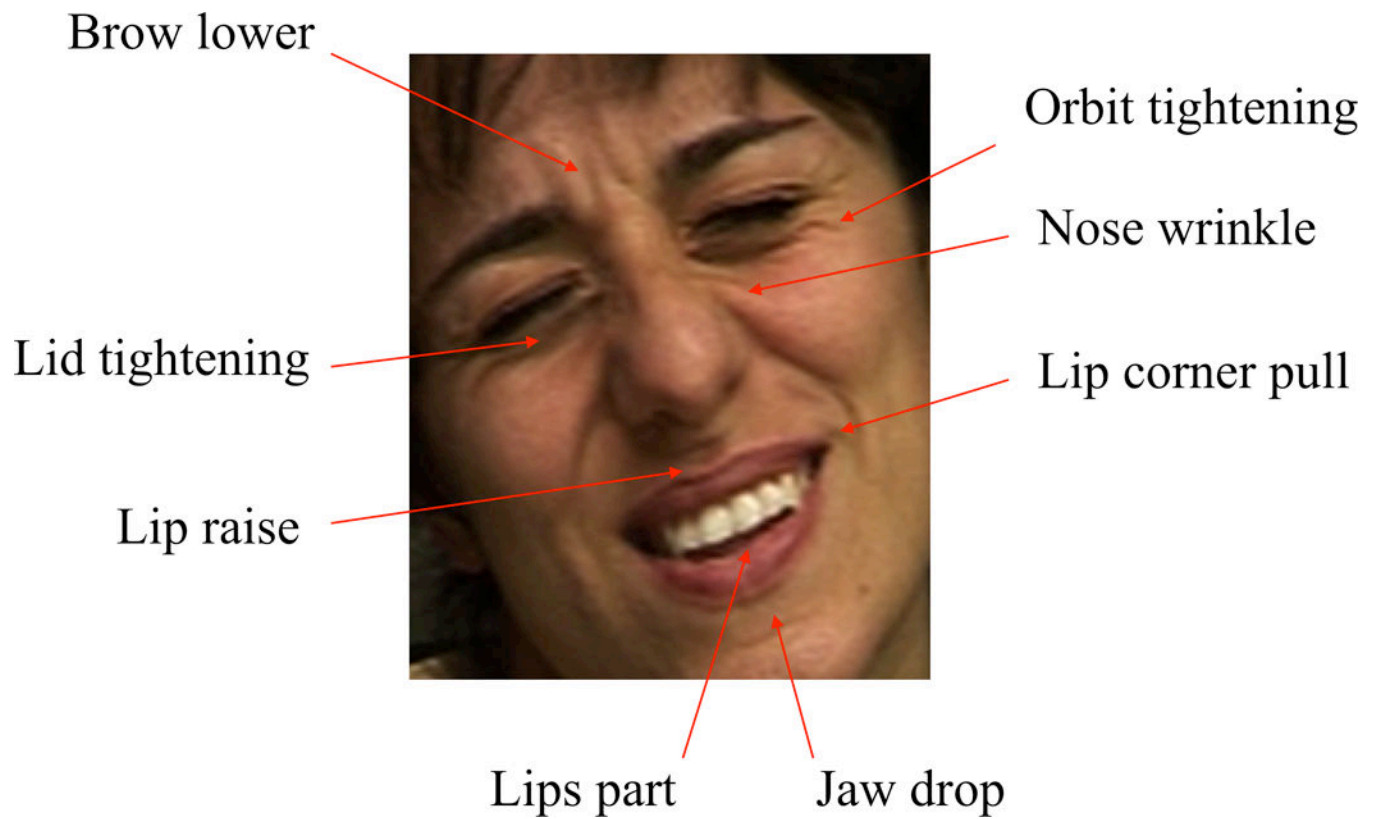
1. Darwin, C. *The Expression of the Emotions in Man and Animals*. London: Murray; 1872.
2. Ekman, P. The argument and evidence about universals in facial expressions of emotion. In: Raskin, DC., editor. *Psychological Methods in Criminal Investigation and Evidence*. New York: Springer Publishing Co, Inc.; 1989. p. 297-332.
3. Frank M, Ekman P, Friesen W. Behavioral markers and recognizability of the smile of enjoyment. *J. Pers. Soc. Psychol.* 1993; 64:83–93. [PubMed: 8421253]
4. Rinn WE. The neuropsychology of facial expression: a review of the neurological and psychological mechanisms for producing facial expression. *Psychol. Bull.* 1984; 95:52–77. [PubMed: 6242437]
5. Kunz M, Chen JI, Lautenbacher S, Vachon-Preseu E, Rainville P. Cerebral regulation of facial expressions of pain. *J. Neurosci.* 2011; 31:8730–8738. [PubMed: 21677157]
6. Brodal, A. *Neurological anatomy: In relation to clinical medicine*. New York: Oxford University Press; 1981.
7. Tschiasny K. Eight syndromes of facial paralysis and their significance in locating the lesion. *Ann. Otol. Rhinol. Laryngol.* 1953; 62:677–691. [PubMed: 13092757]
8. Ekman P, Friesen W. Felt, false, and miserable smiles. *J. Nonverbal Behav.* 1982; 6:238–252.
9. DePaulo BM, Kashy DA, Kirkendol SE, Wyer MM, Epstein JA. Lying in everyday life. *J. Pers. Soc. Psychol.* 1996; 70:979–995. [PubMed: 8656340]
10. Bond CF Jr, DePaulo BM. Accuracy of deception judgments. *Pers. Soc. Psychol. Rev.* 2006; 10:214–234. [PubMed: 16859438]
11. Frank MG, Ekman P. The ability to detect deceit generalizes across different types of high-stake lies. *J. Pers. Soc. Psychol.* 1997; 72:1429–1439. [PubMed: 9177024]
12. Hadjistavropoulos HD, Craig KD, Hadjistavropoulos T, Poole GD. Subjective judgments of deception in pain expression: accuracy and errors. *Pain.* 1996; 65:251–258. [PubMed: 8826514]
13. Hill M, Craig K. Detecting deception in facial expressions of pain: accuracy and training. *Clin. J. Pain.* 2004; 20:415–422.
14. Hill M, Craig K. Detecting deception in pain expressions: the structure of genuine and deceptive facial displays. *Pain.* 2002; 98:135–144. [PubMed: 12098625]
15. McCrystal KN, Craig KD, Versloot J, Fashler SR, Jones DN. Perceiving pain in others: Validation of a dual processing model. *Pain.* 2011; 152:1083–1089. [PubMed: 21388739]
16. Bartlett MS, Littlewort G, Frank M, Lainscsek C, Fasel I, Movellan J. *Recognizing Facial Expression: Machine Learning and Application to Spontaneous Behavior*. IEEE International Conference on Computer Vision and Pattern Recognition. 2005:568–573.
17. Ekman, P.; Friesen, W. *Facial Action Coding System: A Technique for the Measurement of Facial Movement*. Palo Alto, CA: Consulting Psychologists Press; 1978.
18. Bartlett, M.; Whitehill, J. Automated facial expression measurement: Recent applications to basic research in human behavior, learning, and education. In: Rhodes, G.; Calder, Andrew; Haxby,



- James V.; Johnson, Mark H., editors. Handbook of Face Perception. Oxford University Press; 2010.
19. Ekman, P.; Matsumoto, D.; Friesen, W. Facial expression in affective disorders. In: Rosenberg, PEEL., editor. What the Face Reveals. Oxford: New York; 1997. p. 331-341.
  20. Ashraf AB, Lucey S, Cohn JF, Chen T, Ambadar Z, Prkachin KM, Solomon PE. The painful face - Pain expression recognition using active appearance models. *Image Vision Comput.* 2009; 27:1788–1796.
  21. Ekman, P. Telling Lies: Clues to Deceit in the Marketplace, Politics, and Marriage. New York, NY: W.W. Norton; 2001.
  22. Siaw-Asamoah, DM. North American Primary Care Physicians Research Group. Banff, Alberta Canada: 2011. Patients telling lies to physicians: Perception of deception by physician & patients in the United States and Ghana.
  23. Jung B, Reidenberg MM. Physicians being deceived. *Pain Med.* 2007; 8:433–437. [PubMed: 17661857]
  24. Poole GD, Craig KD. Judgments of genuine, suppressed, and faked facial expressions of pain. *J. Pers. Soc. Psychol.* 1992; 63:797–805. [PubMed: 1447693]
  25. Turk DC, Okifuji A. Assessment of patients' reporting of pain: An integrated perspective. *The Lancet.* 1999; 353:1784–1788.

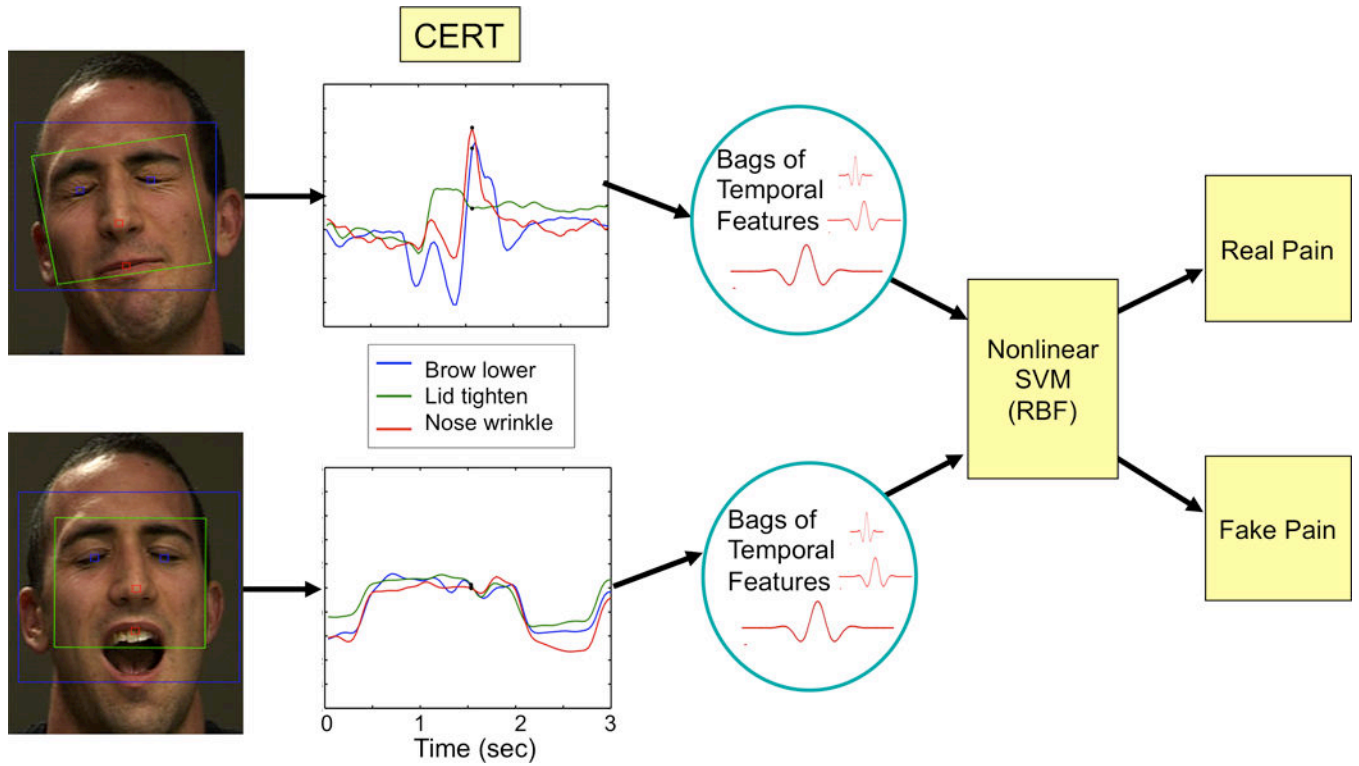
### Highlights

- Untrained human observers cannot differentiate faked from genuine pain expressions
- With training, human performance is above chance but remains poor.
- A computer vision system distinguishes faked from genuine pain better than humans.
- The system detected distinctive dynamic features of expression missed by humans.



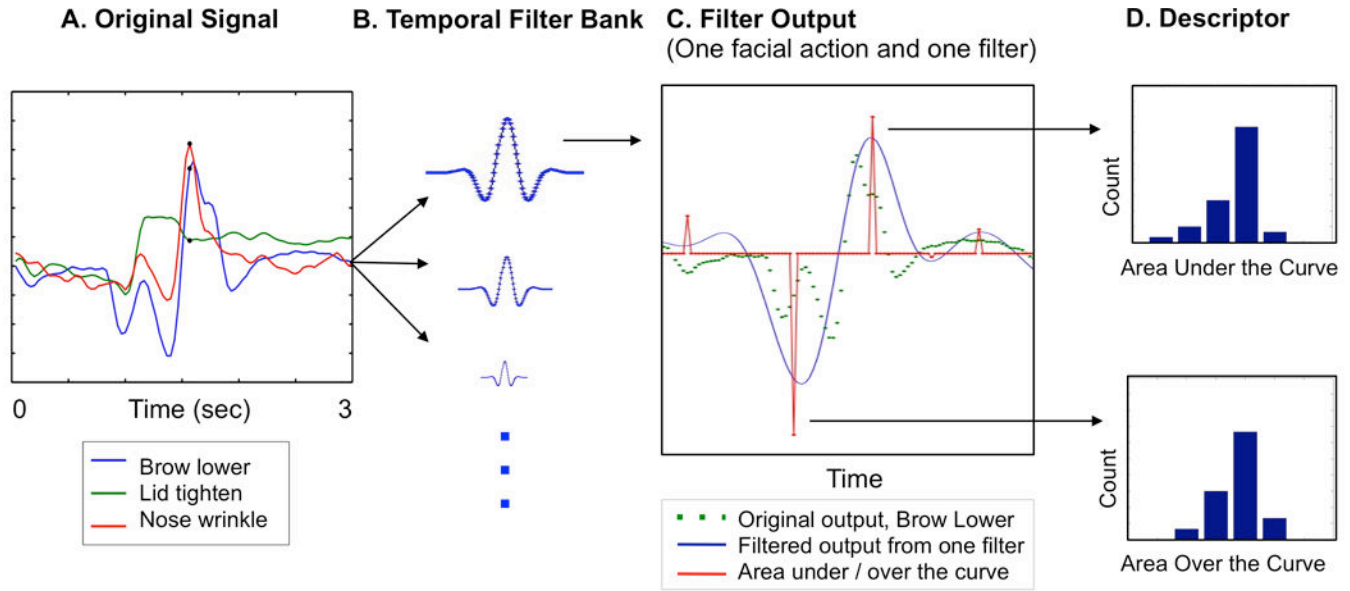
**Figure 1. Example of facial action coding**

Here, a facial expression of pain is coded in terms of eight component facial actions based on the Facial Action Coding System (FACS).



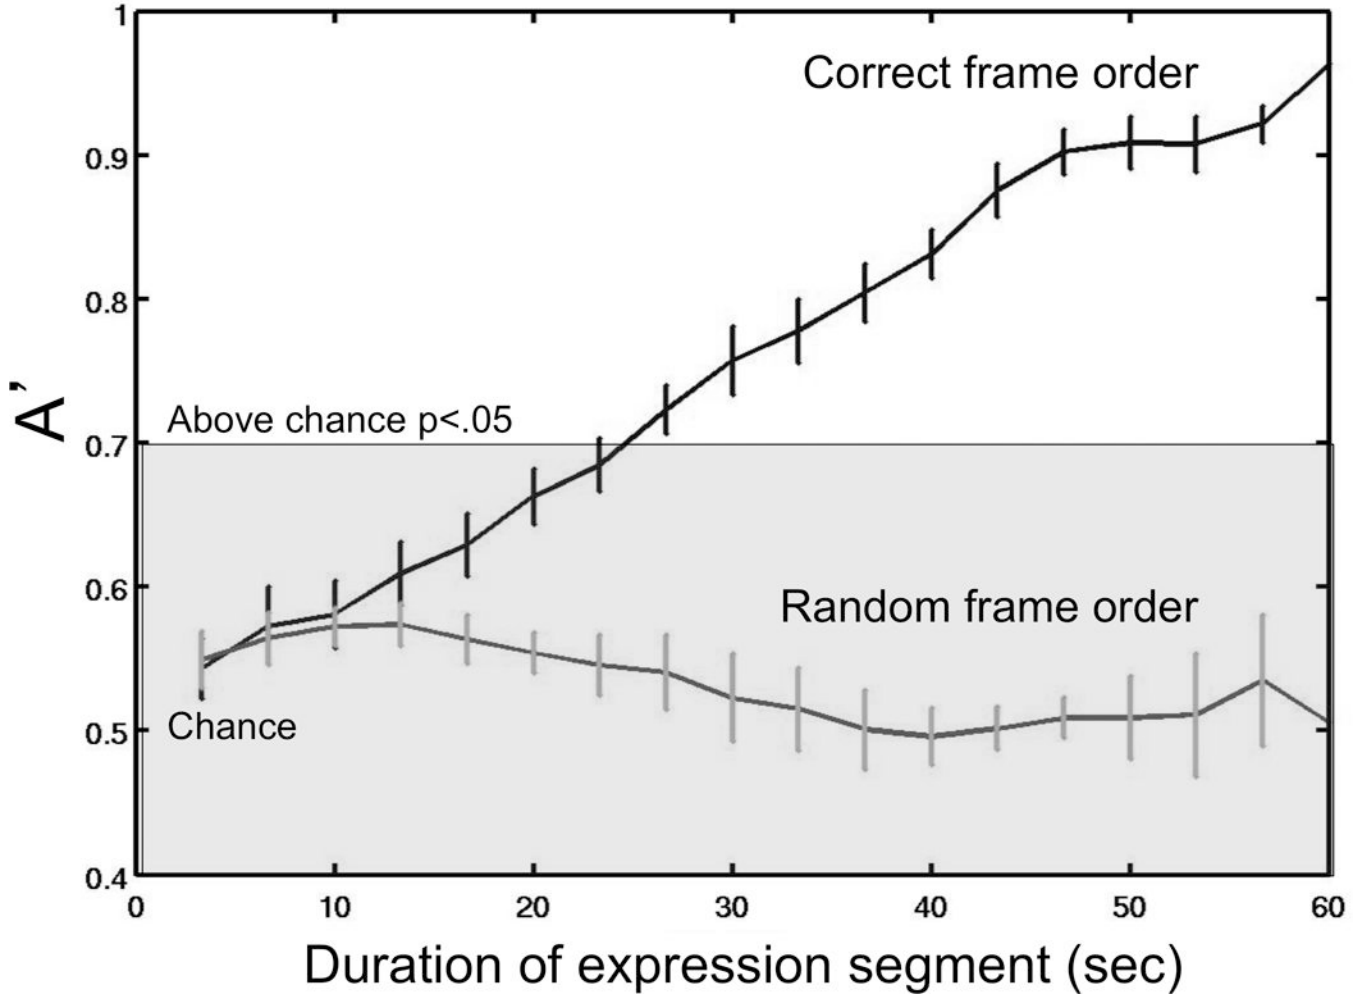
**Figure 2. System Overview**

Face video is processed by the computer vision system, CERT, to measure the magnitude of 20 facial actions over time. The CERT output on the top is a sample of real pain, while the sample on the bottom shows the same three actions for faked pain from the same subject. Note that these facial actions are present in both real and faked pain, but their dynamics differ. Expression dynamics were measured with a bank of 8 temporal Gabor filters and expressed in terms of ‘bags of temporal features.’ These measures were passed to a machine learning system (nonlinear support vector machine) to classify real versus faked pain. The classification parameters were learned from the 24 one-minute examples of real and faked pain.



**Figure 3. Bags of Temporal Features**

Here we illustrate an exemplar of one stimulus as it is processed by each step. A. Sample CERT signals from one subject (Black circles indicate the time point of the face image shown in Figure 2). Three seconds of data are illustrated, but processing is performed on the full 60 seconds of video. B. The CERT signals were filtered by temporal Gabor filters at eight frequency bands. C. Filter outputs for one facial action (brow lower) and one temporal frequency band (the highest frequency). D. Zero crossings are detected, and area under / over the curve calculated. The descriptor consists of histograms of area under the curve for positive regions, and separate histograms for area over the curve for negative regions. (Negative output is where evidence indicates absence of the facial action.) E. Full bag of temporal features for one action (brow lower). Consists of eight pairs of histograms, one per filter.



**Figure 4. Contribution of temporal information**

Classification performance (A') is shown for temporal integration window sizes ranging from 10 seconds to 60 seconds. Windows were sliding and then performance was averaged across temporal position. Performance is shown for the 5-feature system. The region above the shaded region is statistically significant at the  $p < .05$  level. Error bars are one standard error of the mean.