Facial movements have over twenty dimensions of perceived meaning that are only partially captured with traditional methods

Alan S. Cowen^{1,2}, Kunalan Manokara³, Xia Fang⁴, Disa Sauter³, Jeffrey A. Brooks^{1,5}, and Dacher Keltner²

1. Hume AI
2. University of California, Berkeley
3. University of Amsterdam
4. Zhejiang University
5. New York University

Address correspondence to:
Alan S. Cowen
alan@hume.ai

Abstract

Central to science and technology are questions about how to measure facial expression. The current gold standard is the facial action coding system (FACS), which is often assumed to account for all facial muscle movements relevant to perceived emotion. However, the mapping from FACS codes to perceived emotion is not well understood. Six prototypical configurations of facial action units (AU) are sometimes assumed to account for perceived emotion, but this hypothesis remains largely untested. Here, using statistical modeling, we examine how FACS codes actually correspond to perceived emotions in a wide range of naturalistic expressions. Each of 1456 facial expressions was independently FACS coded by two experts ($r = .84$, $\kappa = .84$). Naive observers reported the emotions they perceived in each expression in many different ways, including emotions ($N = 666$); valence, arousal and appraisal dimensions ($N = 1116$); authenticity ($N = 121$), and free response ($N = 193$). We find that facial expressions are much richer in meaning than typically assumed: At least 20 patterns of facial muscle movements captured by FACS have distinct perceived emotional meanings. Surprisingly, however, FACS codes do not offer a complete description of real-world facial expressions, capturing no more than half of the reliable variance in perceived emotion. Our findings suggest that the perceived emotional meanings of facial expressions are most accurately and efficiently represented using a wide range of carefully selected emotion concepts, such as the Cowen & Keltner (2019) taxonomy of 28 emotions. Further work is needed to characterize the anatomical bases of these facial expressions.

Keywords: facial expressions, facial actions, FACS, emotion, emotion perception, facial emotion, face perception, psychology

Introduction

The very thought of being sneered at can induce a shudder; the remembrance of a smile can provide visceral comfort. The potency of facial expression figures prominently in the work of great thinkers from Aristotle to Darwin and has inspired a centuries-long tradition of scientific inquiry. Empirical study confirms that facial expressions are central to our social lives: infants rely on mothers' facial expressions for guidance on how to approach strangers and navigate obstacles (1); adults are more likely to forgive those who express embarrassment (2), mimic those who express pride (3), and concede to those who express anger (4, 5); a CEO's facial expression can even affect the confidence of shareholders (6). Such findings have inspired investment in technologies that analyze facial expression, with applications ranging from psychiatric diagnosis (7, 8) to teaching human preferences to robots (9).

Critical to these scientific and technological efforts are questions about how to classify facial expressions. Competing schools of thought have proposed that facial expressions convey either six discrete emotions (10) or degrees of valence (pleasantness or unpleasantness) and arousal (11). More recently, large-scale experiments have revealed that real-world facial expressions are much richer than these approaches assume, reliably conveying upwards of 28 distinct dimensions of perceived emotion (12, 13). However, the gold standard in measuring facial expression has remained consistent for upwards of 40 years: the facial action coding system (FACS). Developed in the 1970s, FACS is intended to be a comprehensive system for annotating all visible facial muscle movements (14–16). As such, it is often assumed to account for the perceived meanings of facial expressions (15–20).

However, surprisingly few studies have broadly explored how FACS codes relate to the meanings perceived in face movements. Instead, Ekman and Friesen (20) offered narrow predictions of how combinations of FACS codes would be expected to relate to six so-called “basic” emotions—anger, disgust, fear, happiness, sadness, and surprise—which have been the subject of confirmatory emotion recognition studies (21–23). Most subsequent studies have adopted this focus on six emotions without exploring the myriad other meanings that facial expressions may convey (17, 18, 23). One recent exception, and perhaps the most comprehensive study to date on this topic, used FACS to explore the occurrence of specific facial behaviors in 2,600 posed expressions of 22 emotions across five cultures (24). This study revealed that the associations between posed emotions and FACS codes are far more complex than the traditional focus on six prototypical expressions would imply. However, it did not examine the degree to which specific emotions can actually be differentiated based on FACS alone, nor did it address the meanings that facial behaviors convey to perceivers. The extent to which posed facial expressions are representative of the expressions found in everyday life is also unclear (25, 26). In the present study, we characterize how facial behaviors coded using FACS relate to perceived emotion in a wide range of real-world expressions.

Perhaps most importantly, no study has addressed the basic question of the degree to which FACS codes can account for the emotions perceived in facial movement. Of course, this would be self-evident if FACS codes accounted for all visible facial muscle contractions, as intended (15)—but no study has established this definitively. FACS has been difficult to

validate in part because it is difficult to track the full array of patterns of facial muscle activity more directly. One method has been facial electromyography (EMG), which allows for the measurement of a small subset of facial muscles (particularly the corrugator and the zygomaticus major). EMG is however unable to isolate activity to individual muscles, since an electrode placed over any given region of the face typically conflates the activation of four or five muscles (19). The extent to which FACS accounts for variation in perceived emotion remains largely an open question, which is surprising given its centrality to the study of facial expression. Here, we address this question by using statistical modeling to determine comprehensively the degree to which FACS captures perceived emotion in a rich variety of natural facial expressions.

The present study

In the present study, we move beyond traditional assumptions about prototypical expressions of a limited number of emotions to explore the many possible ways FACS codes might actually correspond to perceptions of a high dimensional space of emotion (27). To do so, we asked two independent FACS-trained experts to code 1456 images of wide-ranging expressions captured in highly diverse contexts (e.g., protests, sports, massage, birth) for 53 facial-bodily actions (a standard set of 45 FACS codes plus 8 additions). We compare these FACS codes to emotion judgments gathered in a separate experiment, in which participants reported on the meanings they perceived in the same expressions using theoretically relevant methods: free response; selecting among the Cowen-Keltner 28 (CK28) taxonomy of emotions perceived in facial expression (12); selecting among six “basic” emotions; rating valence, arousal, and 11 other scales of affect and appraisal; and rating

authenticity (12). Finally, separate participants provided ratings of perceived gender, race, and age as well as measures of head pose and gaze direction, factors that could potentially influence perceived emotion. (See Methods 1-6 for details of data collection).

By correlating FACS codes with perceived dimensions of emotion, we find that facial behavior is much richer in meaning than typically assumed (17–23). In all, we find that at least 20 different configurations of facial-bodily muscle actions are perceived as having distinct emotional meanings. Importantly, however, we find that FACS codes do not fully account for all facial muscle movements relevant to perceived emotion. Overall, they capture no more than half of the reliable variance in the emotions people associate with different facial expressions, even when combined with facial demographics, gaze and head direction. These findings have fundamental implications for the broad range of scientific and technological endeavors that base their measurements of facial expression on FACS.

Our findings reveal that to capture the facial movements that convey distinct meanings, there is currently no substitute for high-dimensional human perceptual judgments. Judgments of facial expressions in terms of a broad array of carefully selected emotions have the most explanatory power, capturing almost all of the information available in affect, appraisal, and free response judgments. Of course, this leaves open how facial movements unaccounted for by FACS contribute to facial expression, a question that we suggest requires further study.

Results

The meaning of specific facial actions

To examine how specific facial behaviors coded using FACS correspond to perceived emotions, we correlated each emotion judgment with each FACS code across the 1456 facial expressions. We found that a wide range of distinct emotion judgments corresponded with highly specific facial movements (Fig. 1). This includes the lip wipe for desire, the raised upper lip for disgust, the closing of the eyes for ecstasy, the head droop for shame, and the arm thrust for triumph, among many other associations.

These associations are generally inconsistent with notions that the meanings conveyed by facial behavior reduce to six commonly studied emotions or to valence and arousal (17, 18, 20, 23). For example, the lip corner puller (AU 12), which has been proposed to correspond with happiness or positive emotion (23, 28), was associated with perceptions of only some positive emotions—amusement, contentment, elation, love—but not awe or triumph (29). Other facial actions were associated with multiple specific, but seemingly contradictory emotion judgments; for example, the cheek raise (AU 6) was associated with both amusement and pain (30). More nuanced associations also emerged, such as that of the lip pucker (AU 18) with relief, and nostril compression (AU 39) with concentration. Clearly, the emotions conveyed by facial muscle movements do not reduce to the basic six.

Indeed, when FACS codes are compared only to judgments of the six “basic” emotions or affect and appraisal dimensions (Fig. S1), only a few broad associations emerge. Moreover, these associations are misleading. For example, considering these comparisons in isolation would give the impression that the lip corner puller (AU 12) corresponds with all

perceptions of “happiness” and positive valence, even though we have already seen that it is not associated with all positive emotions. Furthermore, many of the facial behaviors that were revealed to have specific meanings have little to no association with any of the six “basic” emotions, “core” affect (valence and arousal), or the 11 other appraisal dimensions. Clearly, these traditional taxonomies of perceived emotion overlook much of what we see in the face (27).

At least 20 dimensions of expressive facial behavior

Given that emotion perception did not reflect mappings between facial actions and six emotions, as is commonly assumed (17, 18, 20, 23), we next asked how many patterns of facial behavior actually had distinct meanings. To do so, we applied canonical correlations analysis (CCA) between emotion judgments and FACS codes. CCA finds independent patterns in one set of variables that are correlated with patterns in another. In this case, CCA was used to extract independent patterns of emotion perception that were maximally correlated with independent patterns of facial behavior. We found 25 significant canonical correlations ($p \leq .039$, $q[\text{FDR}] < .05$; bootstrap test) when using all codes (face- and body-related), implying that at least 25 independent patterns of facial-bodily behavior have distinct emotional meanings (Fig. 2A). Clearly, the world of facial behavior is much richer than traditionally assumed.

In the above CCA analysis, we included behaviors that go beyond just facial muscle movement, including codes for head pose (AU 53 and 54), as well as hand and arm movement visible near the face (Fig. 1). To determine how many facial muscle movements alone have distinct meanings, we repeated the CCA analysis with only

AUs numbered 46 and below, which correspond specifically to *facial* muscle movements. Here, we found 20 significant canonical correlations ($p \leq .036$, $q[\text{FDR}] < .05$; Fig 3A). This means that at least 20 distinct combinations of facial muscle movements convey distinct meanings. The anatomy of the human face enables an incredibly rich and complex vocabulary of emotion.

To characterize how these 20 different combinations of facial muscle movements corresponded to perceived emotion, we performed factor rotation on the dimensions extracted using CCA. Specifically, we applied varimax rotation to the CCA loadings on the FACS codes,

which extracts a set of independent dimensions that are more interpretable because they have fewer loadings. By applying the same rotation matrix to the corresponding loadings on the CK28 emotions, we visualized how each factor relates FACS codes to perceived emotions. We found that each of the 20 combinations of facial muscle movements with distinct perceived meanings corresponded to a small and often heterogeneous set of highly specific emotions (Fig. 2C). For example, the lip suck (AU 28; factor number 9) often signals shame or concentration, but not sadness or contemplation.

Fig. 1. The emotions perceived in specific facial behaviors.

Correlations between judgments of the CK28 facial emotions and individual FACS codes, including AUs and additional custom codes (e.g., tears, hand position on or near the face). FACS codes are sorted alphabetically according to their most strongly correlated emotion judgment. The emergence of a red diagonal indicates that many FACS codes have highly specific associations with perceived emotions. These associations are often specific to newly studied emotions such as ecstasy, relief, or concentration, and fail to appear in judgments of the six “basic” emotions, valence and arousal, or appraisal dimensions (Fig. S1). Example images corresponding to specific FACS codes are shown for illustration. * $q[\text{FDR}] < .05$ (bootstrap test; see Methods 7a).

Fig. 2. At least twenty facial muscle movements have distinct perceived meaning (a lower bound).

A. Canonical correlations analysis (CCA) reveals 20-25 dimensions. To determine the number of facial behaviors with distinct meaning, we performed canonical correlations analysis (CCA) between the FACS codes and CK28 emotion judgments. We found 25 significant dimensions ($p \leq .039$, $q[FDR] < .05$, bootstrap test; see Methods 7b) when including all codes, and 20 significant dimensions ($p \leq .036$, $q[FDR] < .05$) when including only codes related to facial muscle movement. At least 20 individual facial muscle movements have distinct meanings. **B. Characterizing 20 facial muscle movements with distinct perceived meanings.** To characterize the 20 distinct facial muscle movements, we applied factor rotation (varimax) to the CCA loadings on the FACS codes, and also applied the same rotation matrix to the CCA loadings on the CK28 emotions (see Fig. S2 for original CCA loadings). This extracted the most isolated patterns of facial muscle movement with distinct meaning and determined their associated emotion judgments. Shown here for illustration are facial expressions scoring among the highest on each factor (based on the product of factor scores for FACS and the emotions). **C. Factor loadings of the 20 facial muscle movements with distinct meanings.** For ease of interpretation, we ordered the AUs and emotions in terms of their loadings on each factor. We can see that the emotional meaning of each facial muscle movement is highly distinct and specific, as evidenced by the red diagonals.

What FACS misses

It is safe to say that FACS codes capture meaningful patterns of facial behavior. However, this does not necessarily mean that they fully account for all visible facial muscle movements, or, consequently, for the full meaning of facial expressions. Nevertheless, the assumption that FACS codes are exhaustive of human facial expression (i.e., all facial muscle contractions relevant for perceiving emotion) continues to inspire FACS-based scientific research and technology (17, 18,

23, 31, 32), so it is critical that it be evaluated.

To determine the extent to which FACS codes explain the perception of emotion in facial expression, we used statistical modeling methods. Specifically, we used regression methods to determine the degree to which FACS can explain the reliable variance in a wide range of emotion judgments: the CK28 facial emotions (12); the six “basic” emotions (10, 20, 33); “core” affect (valence and arousal) (11); 13 appraisal and affect dimensions proposed by appraisal theories (34–36); free response (16, 37); and authenticity (see Methods 4 for

details regarding judgment surveys). We used both linear regression and nonlinear regression models (see Methods 7c for details). The results from linear models are directly relevant to psychological inquiry, which generally relies on linear methods such as correlation, analysis of variance, and factor analysis. The results from nonlinear models approximate the overall explanatory value of each rating system. In evaluating each model, we adjust for explainable variance (see Methods 7d). This allows us to estimate what FACS and each set of perceptual judgments could explain in principle, in the absence of sampling error (we will then examine what they explain in practice, when including sample error). We find that FACS misses about half of what people reliably perceive in facial expression. FACS codes explain just 47.2% of the systematic variance in judgments of the CK28 emotions, 57.0% in the six “basic”

emotions, 65.7% in “core” affect, 54.9% in the 13 affect and appraisal dimensions, 35.3% in free response, and 15.9% in authenticity judgments (greater of variance captured by linear or nonlinear models) (Fig. 1A). Clearly, FACS does not fully account for the perceived meaning of facial expressions.

By contrast, judgments of facial expressions in terms of the CK28 emotions explain the vast majority of the variance in other judgments of perceived emotion (92.9% for “core” affect, 92.4% for the six “basic” emotions, 91.2% for appraisal judgments, and 84.4% for free response; Fig. 1A, greater of variance captured by linear or nonlinear models). These results confirm that facial behavior can be mapped onto a few dozen dimensions that largely capture the perception of emotion, even if this is not accomplished by using FACS.

Fig. 3. What FACS and other ratings capture and what they miss.

A. What FACS and other ratings capture in principle. Using linear and nonlinear models with explainable variance adjustment, we determined how much of the reliable variance in various emotion judgments can be captured by FACS, free response, and the CK28 emotions in lieu of sampling error. We found that FACS captures only a fraction of the systematic variance in emotion judgments. By contrast, free response and the CK28 emotions largely capture the same information, and they capture the vast majority of the variance in other kinds of emotion judgments. Neither FACS nor free response nor the CK28 emotions capture much of the reliable variance in judgments of authenticity (especially not FACS), in spite of arguments that FACS specifically captures facial movements relevant to perceived authenticity (such as activation of the orbicularis oculi, or cheek raiser [AU 6]) (38). **B. What FACS and other ratings capture in practice.** We determined what FACS and other ratings capture in practice, with limited numbers of raters, using a calculation derived from the explainable variance (see Methods 7e). Here, we show how much of the systematic variance can be captured in free response, which is considered the broadest conceptualization of emotion. To make different kinds of ratings commensurable, we plot the prediction correlation in free response as a function of the average

amount of person-time it would take raters to provide the ratings. We find that ratings of the CK28 emotions are consistently the most efficient in predicting free response.

Demographics, viewpoint, and gaze

In interpreting the above results, a potential concern is that emotion judgments may also be influenced by factors unrelated to facial muscle movement, including facial demographics, viewpoint (that is, head orientation relative to the camera), and gaze direction. However, additional analyses suggest that these factors explain little additional variance when combined with FACS. The addition of demographics (perceived age, gender, and race) to FACS explained 41.8% of the variance in judgments of the CK28 emotions, barely more than FACS alone (41.0%), using linear regression. A nonlinear model performed slightly worse than FACS alone (45.1% vs. 47.3%), due presumably to demographics washing out more relevant features. Similarly, combining FACS with viewpoint and gaze direction (see Methods 6) explained less of the variance in the CK28 emotions (linear: 41.0%, nonlinear: 40.4%) than FACS alone, perhaps because FACS already includes the important feature of upward and downward head orientation [AU 53 and 54], and the other features effectively add noise). Thus, neither the influence of demographics nor that of viewpoint and gaze can account for the remaining variance in perceived emotion. Having ruled out these confounds, it seems likely that FACS fails to offer a complete description of facial muscle movements related to perceived emotion.

Accounting for intensity

The FACS codes we collected were not graded for levels of intensity, which are difficult to code reliably. This raises the question of whether variation in intensity could potentially account for the unexplained variance in perceived emotion.

To answer this question, we determined the most variance that could possibly be explained through the addition of intensity scores using a linear optimization procedure (see Methods 7f). We found that the upper bound for variance explained in the CK28 emotions with intensity codes would be 69.1% (note that this number is almost certainly much higher than would be achieved with actual intensity scores, given that it is based on the mathematically optimal solution). Thus, even with idealized intensity scores, FACS would be unable to fully account for perceived emotion.

Efficiently conceptualizing perceived emotion

Thus far we have focused on what FACS and perceived emotion judgments can capture in principle, in lieu of sampling error. What do these methods capture in practice? To answer this question, we analyzed how well FACS codes and emotion judgments from varying numbers of human raters can predict free response, which is considered to offer the broadest conceptualization of perceived emotion (16, 37). In Fig. 3B, we plot the accuracy with which free response can be explained by FACS codes and perceptual judgments as a function of the amount of person-time it takes people to provide those codes or judgments (see Methods 7e).

We find that FACS is far from the most efficient way to capture perceived emotion. A single coder takes roughly 3 minutes to code a face, and these codes predict only about 16% of the variance in the average free response (correlation of .4). Because FACS coders are well-trained and highly reliable, the gain from additional coders is limited. Even after about a

person-hour of coding, the explained variance asymptotes at 35% (correlation of .59), which is the most that FACS codes can capture in principle.

We found that judgments of the CK28 emotions are more efficient than the other kinds of ratings we examined (the six basics, core affect, etc.). They captured the average free response rating more accurately in less time, explaining about 67% of the variance (correlation of .82) after 3 person-minutes of rating a face. They were even more efficient in this regard than free response ratings themselves, which captured about 50% of the explainable variance after 3 person-minutes of rating (Fig. 3B). While this may seem counterintuitive, it is due to the tradeoff between precision and sampling error. An individual free response rating is a more precise predictor of the average free response rating, but it is less reliable—it is subject to broad sampling error due to the wide range of possible responses. By contrast, an individual selection using the CK28 emotions predicts a range of free response selections while being more reliable. We found that it takes about 10 person-minutes of rating a single face before free response is more accurate than the CK28 emotions at predicting the average free response. The CK28 emotions, which again explain about 67% of free response with 3 person-minutes of data, are more efficient in this regard than the six “basic” emotions (42%), “core” affect (14%), and the 13 appraisal dimensions (10%). In practice, this suggests that forced choice among a sufficiently wide range of carefully selected emotion concepts is the most accurate and efficient method available to capture perceived emotion.

Discussion

The science of emotion continues to reveal ways in which facial expressions are

central to our social lives (2–6) and inspire technologies that could address pressing psychological needs by analyzing the face (7–9). However, these endeavors hinge on assumptions about how best to classify facial expressions. Following Ekman and Friesen’s (20) prediction that the expression of emotion can be mapped onto six patterns of facial movement (with some allowances for variability in each pattern), many studies have focused narrowly on these six states (17, 18, 23). Studies attempting to look beyond these six expressions have relied on FACS, which is widely assumed to capture all visible facial muscle movements (15–24). These assumptions have shaped the course of the science of emotional expression (17–19, 23, 24, 28, 31, 39), development approaches to emotion (23, 31, 40), and mental health (8, 31, 41), as well as the pursuit of related technologies (8, 18, 32, 41). In the present study, however, we find that human facial expression perception can neither be reduced to six emotions nor to FACS codes. At least 20 different facial muscle movements identified using FACS have distinct meanings, contrary to more reductive typologies of facial expression (Fig. 1-2). Even so, FACS captures no more than 50% of the variance in the perception of emotion in facial behavior (Fig. 3A). A more exhaustive characterization of the meanings conveyed by facial expression can be found in a taxonomy of 28 emotion concepts (12), which largely explain other kinds of judgments of facial expression (Fig. 3).

The meanings conveyed by individual facial movements are sometimes at odds with Ekman and Friesen’s predictions, and are altogether more complex (Fig. 1). The lip corner puller (AU 12), predicted to be an indicator of “happiness”, is associated with some positive emotions, but not all—“amusement”, “contentment”, and

“elation” but not “awe”, “pride”, or “triumph”. The cheek raise (AU 6), also mapped to “happiness”, is associated with “amusement” but also “pain”. The lip tightener (AU 23), mapped to “anger”, is more strongly associated with “contemplation” and “contempt”. And other associations fall outside of the purview of the six “basic” emotions: “desire” with the lip wipe (AU 37), “relief” with the lip pucker (AU 18), “triumph” with the neck tightener (AU 21), and more. In all, facial muscle movements have at least 20 dimensions of meaning (25 when including codes for head pose and hand position) (Fig. 2).

Despite the power of FACS to reveal these associations, on the whole it captures only around 50% of what we perceive in emotional expression. This raises the question of what FACS is missing. Neither intensity codes, nor facial demographics, nor viewpoint and gaze direction can explain the remaining variance in emotion perception. It is possible that subtle factors such as skin coloration, facial sweating, pupil dilation, muscle tone, rashes, and permanent facial characteristics explain some of the remaining variance (42). However, these characteristics tend to have somewhat subtle effects on perceived emotion, even when taken to the extreme (43–45). Furthermore, a qualitative examination of the emotion ratings of the facial expressions, visualized within an interactive map (<https://s3-us-west-1.amazonaws.com/face28/nocontext.html>), reveals faces rated similarly (neighbors within the map) tend to have broadly similar facial configurations, but few discernible similarities in characteristics such as skin coloration and sweating. It is unlikely, then, that such subtle characteristics could explain half of the variance in perceived emotion.

Thus, we ought to consider the possibility that there are inherent limitations

in FACS coders’ capacity to gauge facial muscle contractions based solely on visual inspection of a face. This would run counter to longstanding assumptions, but it should not be controversial. After all, typically-functioning humans are, in a sense, experts in understanding facial expression—a largely innate skill (1, 46, 47) that is sharpened by a lifetime of social interaction (48). By contrast, we are not born to be experts in tracking individual facial muscles, and 100 hours of FACS training is still no match for a lifetime of learning. When we perceive emotion in a face, our brain draws on specialized configural processing capabilities (49) to make sense of complex and interwoven configurations of facial muscle activation. These complex combinations of facial muscle contractions may be too visually convoluted to manually parse into underlying movements of individual facial muscles, even with a carefully designed system like FACS. Perhaps, then, FACS is simply unable to account for all of the nuanced variations in facial behavior that influence the perception of emotion.

Although FACS was unable to fully account for perceived emotion, we found that facial behavior could be represented along other dimensions that capture a broad range of perceptual judgments. In particular, we found that the Cowen & Keltner (2019) taxonomy of 28 emotions (CK28) perceived in facial expression was able to account for the vast majority of the variance in other kinds of judgments, including “core” affect, the six “basic” emotions”, 13 appraisal dimensions, and free response (Fig. 3). These findings indicate that the feelings conveyed by facial expressions are most precisely conceptualized in terms of a wide range of carefully selected emotions (Fig. 3B).

It is important to note that not all English-language emotion concepts have

direct translations in every language (50, 51). Nevertheless, studies of the facial muscle movements that English-speakers associate with the CK28 emotions indicate that they are used across the modern world in similar circumstances, regardless of language (24, 52, 53). While this does not guarantee that these facial behaviors correspond to the same emotional experiences everywhere, another recent study showed that they are in fact associated with similar self-reported emotional experiences in English- and non-English-speaking countries, including the US, UK, Japan, Poland, and Portugal (54) (though the average intensity with which people tend to form each expression differed across cultures, consistent with well-documented differences in perception of high-intensity negative facial expressions between Eastern and Western cultures, as well as notions of cultural display rules (55–57)). Thus, judgments of the CK28 emotions delineate patterns of facial behavior that carry meaning across a wide range of cultures.

Further research is needed to characterize the structural nuances of facial expression not captured by FACS. To do so, two approaches could be taken, both which are based on machine learning and require much more data. One approach would be to train deep neural networks that explain open-ended judgments of perceived emotion, but include a “bottleneck layer” (58) within the network wherein the representation is reduced to the smallest possible dimensionality that maintains optimal model performance. The bottleneck layer should in theory learn to represent the different facial muscle contractions relevant to perceived emotion, since these comprise the most distilled basis for the information conveyed by the face. Another approach would be to train a generative graphical model, such as a deep belief network, to

compress footage of a single face carrying out the widest possible range of facial behaviors into as few dimensions as possible. Once again, the algorithm should in theory learn to represent facial muscle contractions (though additional work would be required to align these representations across different people). In both cases, interpreting the extracted dimensions will also require careful study, which will likely involve the use of FACS as well as further examination of facial anatomy.

Questions about how to classify facial expression are central to affective science and technology. In these endeavors, it is essential to represent facial expression comprehensively (16), in a manner that fully accounts for their emotional meaning and how it is perceived. Only a comprehensive representation of facial expressions can fully account for their powerful role in our everyday lives: in the formation of romantic partnerships (59), the forgiveness of transgressions (2), and even the allocation of investments (6). By looking beyond FACS and six prototypical expressions to consider the much wider range of subtle and nuanced expressions used in the real world (52), the science of emotion can pinpoint what these signals really mean. Such scientific progress will open the door to new technologies that incorporate measures of facial expression to address pressing psychological needs.

Methods

Participants

Emotion judgments of the 1456 facial expressions were obtained for a previous study (12) using Amazon Mechanical Turk. Informed by effect sizes in a previous study of emotional experience (60), we recruited 3096 US English-speaking raters ages 18–76 (1091 female, mean age = 35.2). We used five separate surveys: (i) categorical judgments,

(ii) appraisal judgments, (iii) free response judgments, (iv) authenticity judgments, and (v) categorical judgments of more tightly cropped versions of the facial expressions. See Methods 3: Judgments Formats for details. The experimental procedures were approved by the Institutional Review Board at the University of California, Berkeley.

Collection of 1456 naturalistic expressions

In our previous study (12), we aimed to collect widely varying natural stimuli to depart from a longstanding focus on a narrow range of prototypical expressions. In particular, we originally collected 1500 naturalistic photographs centered on the face but also including hand and body movements near the face. In the present study, we focus on a subset of 1456 images with little to no occlusion of the face. The photographs were found using online search engines by searching for contexts and phrases (e.g., “roller coaster ride”, “flirting”) associated with a wide range of affect- and appraisal-related features and specific emotions. The appraisal features that inspired the search queries were compiled from theories of the appraisal and construction of emotion (11, 34, 61–64). The specific emotions were derived from taxonomies of prominent theorists (see (65) for review), along with more focused studies of positive emotions such as amusement, awe, love, desire, elation, and sympathy (66, 67); subtle emotions that have been found to occur in daily interactions, such as confusion, concentration, doubt, and interest (68, 69); and other recently studied states such as distress, disappointment, and shame (24, 70). See Cowen & Keltner (12) for further details of how the expressions were collected, including search queries. The expressions were selected for authenticity and diversity, but we do not claim that they are exhaustive of all facial-bodily signaling.

FACS coding

Each facial expression was FACS coded by two expert coders, authors K.M. and X.F., in terms of 45 action units (AUs) and 8 additional codes (hand over face, hand over mouth, hand over eyes, hand over forehead, hand touching face, arms raised, tears, neck visible). Five codes went unused (AUs 29, 31, 33, 35 and 72). The codes were highly reliable, with an overall interrater reliability correlation (r) of .84 (also known as Matthew’s correlation coefficient) and a Cohen’s kappa (κ) also equal to .84, demonstrating the expertise and skill of the coders.

Judgment formats

In each of the five different judgment formats in our previous study (12), we obtained repeated (18-20) judgments of each expression from a different group of participants. This sample size was found to be sufficient to approximate more than 90% of the systematic variance in judgments of each image in previous studies that relied on similar methods (60, 71).

One group of participants ($N = 364$, 190 female, mean age = 34.3, mean facial expressions rated = 82.4) judged each image in terms of 28 specific emotions derived from an extensive review of past studies of emotional expression (24, 65–70). The emotions included: AMUSEMENT, ANGER, AWE, CONCENTRATION, CONFUSION, CONTEMPLATION, CONTEMPT, CONTENTMENT, DESIRE, DISAPPOINTMENT, DISGUST, DISTRESS, DOUBT, ECSTASY, ELATION, EMBARRASSMENT, FEAR, INTEREST, LOVE, PAIN, PRIDE, REALIZATION, RELIEF, SADNESS, SHAME, SURPRISE, SYMPATHY, and TRIUMPH. To our knowledge, this taxonomy encompasses the emotions for which facial-bodily expressive signals have been proposed in past literature, but we do not claim that it is exhaustive of all facial-bodily expressions.

A second group of participants ($N = 1116$, 584 female, mean age = 34.3, mean faces judged = 24.2) evaluated each expression in terms of 13 appraisals derived from constructivist and appraisal theories (11, 34, 61–64). To ensure that perceived emotion was represented comprehensively (16, 72), a third set of participants ($N = 193$, 110 female, mean age = 37.1, mean faces judged = 69.9) provided free response descriptions of each expression. In this method, as participants typed, a drop-down menu appeared, displaying items from a corpus of 600 emotion terms containing the currently typed substring. For example, typing the substring “lov-” caused the following terms to be displayed: love, “brotherly love,” “feeling loved,” “loving sympathy,” “maternal love,” “romantic love,” and “self-love.” Participants selected as many emotion terms as desired. Another group of participants ($N = 121$, 56 female, mean age = 36.8, mean faces judged = 37.2) rated each expression in terms of its perceived authenticity. To rule out effects of the immediate context surrounding the face (e.g., clothing), a final group of participants ($N = 302$, 151 female, mean age = 37.6, mean faces judged = 86.9) completed the first survey (28 emotion categories) but with more tightly cropped versions of the face images. See (12) for more details and precise questions from each survey.

All told this amounted to 425,208 individual judgments (56,208 forced-choice categorical judgments, 351,000 nine-point appraisal scale judgments, 13,500 free response judgments, and 4,500 authenticity judgments).

Demographics

To examine the extent to which the perceived demographic group of each face influenced perceived emotion, we collected separate ratings of the perceived gender, age, and race/ethnicity of each face. One

group of participants ($N = 39$, 17 female, mean age = 33.2, mean faces judged = 115.4) rated each face in terms of its perceived gender (‘Male’, ‘Female’, or ‘Other’). Another group of participants ($N = 66$, 31 female, mean age = 32.9, mean faces judged = 45.5) rated each face in terms of its perceived age in years (0–110). A final group of participants ($N = 66$, 28 female, mean age = 33.0, mean faces judged = 45.5) rated each face in terms of its perceived race/ethnicity (‘Asian’, ‘Black or African Descent’, ‘Indian American or Desi’, ‘Hispanic or Latino’, ‘Middle Eastern’, ‘Native American’, ‘White’, ‘Other’). For purposes of modeling the effects of perceived age on perceived emotion, age ratings were divided into ten bins with a fairly balanced distribution of faces: ‘Infant’ (ages 0–1), ‘Toddler’ (ages 2–4), ‘Child’ (ages 5–12), ‘Teenager’ (ages 13–18), ‘Young Adult’ (ages 19–22), ‘20s’ (ages 23–29), ‘30s’ (ages 30–39), ‘40s’ (ages 40–49), ‘Middle Age’ (ages 50–64), and ‘Senior’ (ages 65+).

Viewpoint and gaze

To examine the extent to which facial viewpoint and gaze direction influenced perceived emotion, we estimated the viewpoint and gaze direction of each face using algorithms developed at Google. The viewpoint algorithm processed each face and estimated the pitch, raw, and roll of the face relative to the camera. The gaze direction algorithm processed each face and estimated the X and Y coordinates of the gaze on the geometric plane that encompassed the camera lens (we excluded faces where the pupils were not visible). For purposes of modeling the effects of viewpoint and gaze direction on perceived emotion, the dimensions of viewpoint and gaze direction were divided into quantile bins and intersected, yielding 12 features in total: three bins of magnitude in “yaw”

(ignoring direction, given leftward and rightward yaw likely have equivalent effects on perceived emotion) plus three bins of “pitch” (down, neutral, or up) plus three bins of X gaze extremity (again ignoring direction) plus three bins of Y gaze extremity (down, neutral, up). For analysis involving viewpoint and gaze direction, we excluded 226 faces where the pupils were not visible or where gaze direction could not be estimated.

Statistical analyses

Bootstrapping and false discovery rate. For all significance tests, we used bootstrap resampling across the face images. The number of iterations was set to 1000. We controlled the false discovery rate (FDR) using the Benjamini-Hochberg procedure (73).

Significance testing for CCA. To examine the number of significant dimensions of emotion that corresponded to different combinations of FACS codes, we applied CCA in a leave-on-out fashion. Specifically, CCA was iteratively applied on all but one face and the FACS codes and emotion ratings of the remaining face were then projected onto the extracted CCA variates. We then correlated the held-out variate scores across all 1456 faces, using a Spearman partial rank correlation to partial out previous variates and thereby control for any monotonic nonlinear associations between the extracted variates. We used bootstrapping to test for significance of the held-out canonical correlations.

Linear and nonlinear models. To partition the variance among any two kinds of ratings, we used linear and nonlinear models to predict one from the other. The linear models were ordinary least squares (OLS) models, trained and tested using 10-fold

cross-validation. The nonlinear models were k -nearest neighbors with $k = 5$.

Accounting for sampling error. To compute unbiased prediction correlations, it is necessary to control for interrater agreement levels by dividing by the noise ceiling, which is the maximum raw prediction correlation that can be obtained given the sampling error in individual human ratings (74–77). To compute the noise ceiling in any rating, we randomly assigned the ratings provided for each face into two groups. We then took the mean of the ratings of each face within each group, correlated the means between the groups – across all 1456 faces – and took the square root of the correlation (74). This represents the maximum correlation that can be obtained in predicting the ratings in either group.

Similarly, we also computed the noise ceiling in model predictions generated using any human ratings (FACS codes or emotion judgments) by randomly assigning the ratings into two groups and training and evaluating the model within each group. The square root of the correlation in the predicted ratings across the two groups represents the maximum correlation that can be obtained using those predictions.

We divided prediction correlations by the noise ceiling both in the prediction and in the predicted variable. This procedure estimates the prediction correlation that can be obtained in principle, in lieu of sampling error.

Given that these procedures involved splitting the ratings randomly into two groups, the procedures repeated 100 times and the resulting correlations were averaged. The standard deviations in the correlation across repetitions of the procedure were exceedingly small (generally <0.5% of the correlation) so the average after 100 repetitions was extremely stable.

What FACS and other ratings capture in practice. To determine what FACS and other ratings capture with limited numbers of raters, we used a calculation derived from the explainable variance. In particular, we know that the signal-to-noise ratio (explainable variance divided by unexplainable variance) is directly proportional to the number of ratings obtained for each face (74). We can thus compute the explainable variance and noise ceiling in the average rating of each face as a function of the number of ratings obtained. The formula for the noise ceiling of the prediction correlation as a function of the number of ratings is $\sqrt{1/(\frac{1}{SNR} + 1)}$ with $SNR = N * EXPVAR / (1 - EXPVAR)$ where N is the number of ratings per face and EXPVAR is the explainable variance of a prediction with a model trained and tested on one rating per face.

To plot the noise ceiling as a function of person-time rating each face, we approximated the amount of time required to contribute each kind of rating by experimenting with each survey by hand. We determined that FACS rating a single face required about 3 minutes, a choice among the CK28 emotions required about 9 seconds, ratings of “core” affect (valence and arousal) required 8 seconds, ratings of 13 different appraisal dimensions required 50 seconds, a free response rating required 14 seconds, and a choice among the basic emotions required 6 seconds.

Accounting for intensity. During FACS coding, we opted not to collect intensity scores, which are generally low in reliability. Instead, we analyzed the maximum amount of additional variance that could be

explained by the addition of intensity scores. To do so, we used a constrained least squares linear optimization procedure.

The first step in this procedure was to train a linear model to predict the CK28 emotion ratings from FACS. Then, we used linear optimization to choose the optimal intensity levels for each FACS code that would explain the greatest possible variance in the linear mapping from FACS to emotion ratings of each face. The intensity levels were constrained to between .2 and 5, given that FACS codes are typically assigned five levels of intensity (A-E), which means the coded intensity could be between one fifth and five times the intensity level of other FACS codes. Where the two FACS coders disagreed, intensity levels were instead constrained to between 0 and 5. The optimization was performed using the interior-point-convex quadratic programming algorithm with a 10^{-8} tolerance on the constraint violation, a 10^{-8} tolerance on the first-order optimality, and a 10^{-12} termination tolerance. Finally, to simulate ordinal intensity values, we divided the non-zero intensity codes into five quantile bins and replaced them with the averages from each bin.

After computing the optimal intensity scores, we performed least squares linear regression in a leave-one-face-out manner to determine the variance explained by the intensity-coded FACS scores. This procedure overestimates the variance that could be explained by intensity codes, due to overfitting. Nevertheless, we found that FACS with optimal hypothetical intensity codes explained 69.1% of the variance in the CK28 emotion judgments, still leaving much of the variance unexplained.

References

1. Sorce JF, Emde RN, Campos J, Klinnert MD (1985) Maternal Emotional Signaling. Its Effect on the Visual Cliff Behavior of 1-Year-Olds. *Dev Psychol* 21(1):195–200.
2. Feinberg M, Willer R, Keltner D (2012) Flustered and faithful: Embarrassment as a signal of prosociality. *J Pers Soc Psychol* 102(1):81–97.
3. Martens JP, Tracy JL (2013) The Emotional Origins of a Social Learning Bias: Does the Pride Expression Cue Copying? *Soc Psychol Personal Sci* 4(4):492–499.
4. Van Kleef GA, De Dreu CKW, Manstead ASR (2004) The Interpersonal Effects of Anger and Happiness in Negotiations. *J Pers Soc Psychol* 86(1):57–76.
5. Reed LI, DeScioli P, Pinker SA (2014) The Commitment Function of Angry Facial Expressions. *Psychol Sci* 25(8):1511–1517.
6. ten Brinke L, Adams GS (2015) Saving face? When emotion displays during public apologies mitigate damage to organizational performance. *Organ Behav Hum Decis Process* 130:1–12.
7. Grabowski K, et al. (2019) Emotional expression in psychiatric conditions: New technology for clinicians. *Psychiatry Clin Neurosci* 73(2):50–62.
8. Girard JM, Cohn JF, Mahoor MH, Mavadati S, Rosenwald DP (2013) Social risk and depression: Evidence from manual and automatic facial expression analysis. *2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition, FG 2013*. doi:10.1109/FG.2013.6553748.
9. Hemminghaus J, Kopp S (2017) Towards Adaptive Social Behavior Generation for Assistive Robots Using Reinforcement Learning. *ACM/IEEE International Conference on Human-Robot Interaction*, pp 332–340.
10. Ekman P (1992) An Argument for Basic Emotions. *Cogn Emot* 6(3–4):169–200.
11. Russell JA (2003) Core Affect and the Psychological Construction of Emotion. *Psychol Rev* 110(1):145–172.
12. Cowen AS, Keltner D (2019) What the Face Displays: Mapping 28 Emotions Conveyed by Naturalistic Expression. *Am Psychol*. doi:10.1037/amp0000488.
13. Cowen AS, Keltner D (2021) Semantic space theory: A computational approach to emotion. *Trends Cogn Sci*.
14. Hjortsjö C-H (1969) *Man's face and mimic language* (Studentlitteratur).
15. Ekman P, Friesen W V. (1978) *Facial Action Coding System* (Consulting Psychologist Press, Palo Alto).
16. Matsumoto D, Ekman P, Fridlund A (1991) Analyzing nonverbal behavior. *Practical Guide To Using Video in the Behavioral Sciences*, pp 153–166.
17. Jack RE, Sun W, Delis I, Garrod OGB, Schyns PG (2016) Four not six: Revealing culturally common facial expressions of emotion. *J Exp Psychol Gen* 145(6):708–730.
18. Srinivasan R, Martinez AM (2018) Cross-Cultural and Cultural-Specific Production and Perception of Facial Expressions of Emotion in the Wild. *IEEE Transactions on Affective Computing* doi:10.1109/TAFFC.2018.2887267.
19. Cohn JF, Ambadar Z, Ekman P (2007) Observer-based measurement of facial expression with the Facial Action Coding System. *The Handbook of Emotion Elicitation and Assessment*, eds Coan JA, Allen JJB (Oxford University Press), pp 203–221.
20. Ekman P, Friesen W V. (1978) *Investigator's guide to the facial action coding system, parts I and III* (Consulting Psychologist Press).
21. Wiggers M (1982) Judgments of facial expressions of emotion predicted from facial behavior. *J Nonverbal Behav* 7(2):101–116.
22. Suzuki K, Naitoh K (2003) Brief report useful information for face perception is described with FACS. *J Nonverbal Behav* 27(1):43–55.
23. Barrett LF, Adolphs R, Marsella S, Martinez AM, Pollak SD (2019) Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. *Psychol Sci Public Interes* 20(1):1–68.
24. Cordaro DT, et al. (2018) Universals and cultural variations in 22 emotional expressions across five cultures. *Emotion* 18(1):75–93.
25. Ekman P, Friesen W V., Ellsworth P (1972) Emotion in the human face: Guidelines for research and an integration of findings. *J Pers Soc Psychol* 74(1):80–85.
26. Bartlett MS, et al. (2006) Automatic recognition of facial actions in spontaneous expressions. *J Multimed* 1(6):22–35.
27. Cowen AS, Sauter D, Tracy JL, Keltner D (2019) Mapping the Passions: Toward a High-Dimensional Taxonomy of Emotional Experience and Expression. *Psychol Sci Public Interes* 20(1):69–90.
28. Crivelli C, Carrera P, Fernández-Dols JM (2015) Are smiles a sign of happiness? Spontaneous expressions of judo winners. *Evol Hum Behav* 36(1):52–58.
29. Sauter DA (2017) The Nonverbal Communication of Positive Emotions: An Emotion Family Approach. *Emot Rev* 9(3):222–234.
30. Mattson WI, Cohn JF, Mahoor MH, Gangi DN, Messinger DS (2013) Darwin's Duchenne: Eye constriction during infant joy and distress. *PLoS One* 8(11). doi:10.1371/journal.pone.0080161.
31. Ekman P, Rosenberg EL (2012) *What the Face Reveals: Basic and Applied Studies of Spontaneous Expression Using the Facial Action Coding System (FACS)* doi:10.1093/acprof:oso/9780195179644.001.0001.

32. Girard JM, Cohn JF, Jeni LA, Sayette MA, De la Torre F (2014) Spontaneous facial expression in unscripted social interactions can be measured automatically. *Behav Res Methods* 47(4):1136–1147.
33. Elfenbein HA, Ambady N (2002) On the universality and cultural specificity of emotion recognition: A meta-analysis. *Psychol Bull* 128(2):203–235.
34. Smith CA, Ellsworth PC (1985) Patterns of Cognitive Appraisal in Emotion. *J Pers Soc Psychol* 48(4):813–838.
35. Shaver P, Schwartz J, Kirson D, O'Connor C (1987) Emotion Knowledge: Further Exploration of a Prototype Approach. *J Pers Soc Psychol* 52(6):1061–1086.
36. Fontaine JRJ, Scherer KR, Roesch EB, Ellsworth PC (2007) The world of emotions is not two-dimensional. *Psychol Sci* 18(12):1050–1057.
37. Russell JA (1994) Is there universal recognition of emotion from facial expression? A review of the cross-cultural studies. *Psychol Bull* 115(1):102–141.
38. Gunnery SD, Ruben MA (2016) Perceptions of Duchenne and non-Duchenne smiles: A meta-analysis. *Cogn Emot* 30(3):501–515.
39. Ekman P, Friesen W V., Ancoli S (1980) Facial signs of emotional experience. *J Pers Soc Psychol* 39(6):1125–1134.
40. Camras LA, Shutter JM (2010) Emotional facial expressions in infancy. *Emot Rev* 2(2):120–129.
41. Xu X, et al. (2019) Automated pain detection in facial videos of children using human-assisted transfer learning. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, pp 162–180.
42. Kret ME (2015) Emotional expressions beyond facial muscle actions. A call for studying autonomic signals and their impact on social perception. *Front Psychol* 6. doi:10.3389/fpsyg.2015.00711.
43. Nakajima K, Minami T, Nakauchi S (2017) Interaction between facial expression and color. *Sci Rep* 7. doi:10.1038/srep41019.
44. Hess EH (1975) The role of pupil size in communication. *Sci Am* 233(5):110–112, 116.
45. De Melo CM, Gratch J (2009) Expression of emotions using wrinkles, blushing, sweating and tears. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, pp 188–200.
46. Bertin E, Striano T (2006) The still-face response in newborn, 1.5-, and 3-month-old infants. *Infant Behav Dev* 29(2):294–297.
47. Wu Y, Muentener P, Schulz LE (2017) One- to four-year-olds connect diverse positive emotional vocalizations to their probable causes. *Proc Natl Acad Sci U S A* 114(45):11896–11901.
48. Rodger H, Vizioli L, Ouyang X, Caldara R (2015) Mapping the development of facial expression recognition. *Dev Sci* 18(6):926–939.
49. Calder AJ, Keane J, Young AW, Dean M (2000) Configural information in facial expression perception. *J Exp Psychol Hum Percept Perform* 26(2):527–551.
50. Kollareth D, Russell JA (2017) The English word disgust has no exact translation in Hindi or Malayalam. *Cogn Emot* 31(6):1169–1180.
51. Jackson JC, et al. (2019) Emotion semantics show both cultural variation and universal structure. *Science (80-)* 366(6472):1517–1522.
52. Cowen AS, et al. Sixteen facial expressions occur in similar contexts worldwide. *Nature*.
53. Cordaro DT, et al. (2020) The Recognition of 18 Facial-Bodily Expressions Across Nine Cultures. *Emotion* 20:1292–1300.
54. Cowen AS, et al. How emotion is experienced and expressed in multiple cultures: A large-scale experiment. *under Rev Nat Hum Behav*.
55. Ekman P, Sorenson ER, Friesen W V. (1969) Pan-cultural elements in facial displays of emotion. *Science (80-)* 164(3875):86–88.
56. Elfenbein HA (2013) Nonverbal dialects and accents in facial expressions of emotion. *Emot Rev* 5(1):90–96.
57. Matsumoto D, Olide A, Schug J, Willingham B, Callan M (2009) Cross-cultural judgments of spontaneous facial expressions of emotion. *J Nonverbal Behav* 33(4):213–238.
58. McLaren M, Ferrer L, Lawson A (2016) Exploring the role of phonetic bottleneck features for speaker and language recognition. *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pp 5575–5579.
59. Gonzaga GC, Keltner D, Londahl EA, Smith MD (2001) Love and the commitment problem in romantic relations and friendship. *J Pers Soc Psychol* 81(2):247–262.
60. Cowen AS, Keltner D (2017) Self-report captures 27 distinct categories of emotion bridged by continuous gradients. *Proc Natl Acad Sci U S A* 114(38):E7900–E7909.
61. Scherer KR (2009) The dynamic architecture of emotion: Evidence for the component process model. *Cogn Emot* 23(7):1307–1351.
62. Roseman IJ (1991) Appraisal Determinants of Discrete Emotions. *Cogn Emot* 5(3):161–200.
63. Mehrabian A, Russell JA (1974) *An approach to environmental psychology* (Cambridge, MA: MIT Press).
64. Osgood CE (1966) Dimensionality of the Semantic Space for Communication Via Facial Expressions. *Scand J Psychol* 7(1):1–30.
65. Keltner D, Tracy JL, Sauter D, Cowen AS (2019) What Basic Emotion Theory Really Says for the Twenty-First Century Study of Emotion. *J Nonverbal Behav* 43(2):195–201.
66. Campos B, Shiota MN, Keltner D, Gonzaga GC, Goetz JL (2013) What is shared, what is different? Core relational themes and expressive displays of eight positive emotions. *Cogn Emot* 27(1):37–52.

67. Shiota MN, et al. (2017) Beyond happiness: Building a science of discrete positive emotions. *Am Psychol* 72(7):617–643.
68. Rozin P, Cohen AB (2003) High Frequency of Facial Expressions Corresponding to Confusion, Concentration, and Worry in an Analysis of Naturally Occurring Facial Expressions of Americans. *Emotion* 3(1):68–75.
69. Benitez-Quiroz CF, Wilbur RB, Martinez AM (2016) The not face: A grammaticalization of facial expressions of emotion. *Cognition* 150:77–84.
70. Tracy JL, Matsumoto D (2008) The spontaneous expression of pride and shame: Evidence for biologically innate nonverbal displays. *Proc Natl Acad Sci U S A* 105(33):11655–11660.
71. Cowen AS, Elfenbein HA, Laukka P, Keltner D (2018) Mapping 24 Emotions Conveyed by Brief Human Vocalization. *Am Psychol* 22(4):274–276.
72. DiGirolamo MA, Russell JA (2017) The emotion seen in a face can be a methodological artifact: The process of elimination hypothesis. *Emotion* 17(3):538–546.
73. Benjamini Y, Hochberg Y (1995) Controlling the False Discovery Rate - a Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society Series B-Methodological* 1995.pdf. *J R Stat Soc Ser B* 57(1):289–300.
74. Lord FM, Novick MR (1968) The Construction of True and Error Scores (Chapter 2). *Statistical Theories of Mental Test Scores*, pp 358–393.
75. John OP, Soto CJ (2007) The Importance of Being Valid: Reliability and the Process of Construct Validation (Chapter 27). *Handbook of Research Methods in Personality Psychology*, p 461.
76. Benjamini Y, Yu B (2013) The shuffle estimator for explainable variance in FMRI experiments. *Ann Appl Stat* 7(4):2007–2033.
77. Csárdi G, Franks A, Choi DS, Airoidi EM, Drummond DA (2015) Accounting for Experimental Noise Reveals That mRNA Levels, Amplified by Post-Transcriptional Processes, Largely Determine Steady-State Protein Levels in Yeast. *PLoS Genet* 11(5). doi:10.1371/journal.pgen.1005206.

Supplementary Figures

Fig. S1. The emotions perceived in specific facial behaviors.

Correlations between FACS codes and judgments of the six “basic” emotions (above) and 14 affect and appraisal judgments including valence and arousal as well as authenticity (below). FACS codes are sorted alphabetically according to their most strongly correlated affect judgment.

Fig. S2. Original loadings from canonical correlations analysis (CCA) applied between the FACS codes and CK28 emotion judgments.

For purposes of visualizing the loadings, a small amount of Gaussian noise ($SD = .01$) was added to the category judgments to make them full rank before performing CCA.