Title: How emotion is experienced and expressed in multiple cultures: A large-scale experiment

Short Title: Data-driven emotion modeling

One Sentence Summary: Computational analyses of 45 thousand recorded reactions to 2 thousand evocative videos reveal universals in the structure of emotional experience, distinct facial expressions for upwards of 12 specific feelings, and systematic influences of cultural display tendencies.

Authors: Alan S. Cowen^{1,2}, Gautam Prasad³, Misato Tanaka⁴, Yukiyasu Kamitani⁴, Vladimir Kirilyuk³, Krishna Somandepalli^{3,5}, Brendan Jou³, Florian Schroff³, Hartwig Adam³, Jeffrey A. Brooks^{1,2}, & Dacher Keltner^{1,2}

Affiliations: ¹Hume AI, New York NY, USA. ²Department of Psychology, University of California, Berkeley, CA, USA. ³Google Research, Mountain View, CA, USA. ⁴Advanced Telecommunications Research Institute, Kyoto, Japan. ⁵Department of Electrical Engineering, University of Southern California, Los Angeles, CA, USA.

Abstract: Core to understanding emotion are subjective experiences and their embodiment in facial behavior. Past studies have focused on six emotions and prototypical facial poses, reflecting limitations in scale and narrow assumptions about emotion. We examine 45,231 reactions to 2,185 evocative videos, largely in North America, Europe, and Japan, collecting participants' self-reported experiences in English or Japanese and manual/automated annotations of facial movement. We uncover 21 dimensions of emotion underlying experiences reported across languages. Facial expressions predict at least 12 dimensions of experience, despite individual variability. We also identify culture-specific display tendencies—many facial movements differ in intensity in Japan compared to the U.S./Canada and Europe, but represent similar experiences. These results reveal how people actually experience and express emotion: in high-dimensional, categorical, and complex fashion.

Introduction

Suppose a film routinely arouses laughs, cries, or grimaces. What can we infer about how it makes people feel? Answers to this question, although seemingly intuitive, remain scientifically elusive^{1–3}. For the past 50 years, efforts tracing to a classic study by Ekman and Friesen⁴ have sought to illuminate the nature of emotional expression in the search for one-to-one mappings between six “basic” emotions—anger, disgust, fear, happiness, sadness, and surprise—and prototypical facial expressions^{1,5–7}. However, the focus on these six states has little theoretical or empirical rationale, and the world of emotional expression is proving to be far more complex.

Recent studies have established that emotional experience and expression are upwards of four times more complex, in terms of the number of states and their blends that people experience and express, than represented in studies of six emotions^{2,8–16}. People reliably report feeling dozens of other emotions such as “awe”, “excitement”, and “relief” in a wide range of circumstances^{8,11} and perceive most of these emotions in widely varying emotional expressions^{13,17–19}. The six “basic” emotions capture under a third of the variance in people’s diverse, systematic responses to emotion antecedents and expressions². Taken together, these methodological and empirical advances reveal that studies seeking to map prototypical expressions onto experiences of the six “basic” emotions can only have captured a fraction, at best, of what facial

expressions reveal. How people move their faces when they feel emotions is still largely an open question.

Understanding how people actually express emotion requires studying a much wider range of behaviors with open-ended analytic approaches. We introduce the Berkeley Reactions to Affective Video Elicitors (BRAVE) database, a corpus of 45,231 recordings of participants largely in North America, Europe, and Japan reacting to 2,185 evocative short videos and reporting on their emotional experiences. BRAVE is orders of magnitude larger than previously published databases of recorded reactions to emotional stimuli. The videos presented to participants are diverse, naturalistic and emotionally evocative, including scenes from across life--vast landscapes, births, funerals, accidents, historic footage, practical jokes, artwork. Moreover, we compare cultural groups with disparate histories, languages, belief systems, norms, and values^{20–23}. By analyzing thousands of spontaneous reactions to rich, diverse antecedents, we paint a detailed portrait of how people experience and express emotion.

What emotions do people in different cultures experience when viewing evocative videos? What facial experiences do they produce? How are feelings and expressions related? To answer these questions, we approach BRAVE with a recently developed framework for investigating semantic spaces of emotion-related behavior^{2,3,24}, automated methods of quantifying facial expressions²⁵, and data-driven statistical modeling

approaches. Our findings reveal how (a) a rich, high-dimensional space of emotional experience and expression is shared across different cultural groups, (b) a wide range of self-reported emotions can accurately be inferred from facial expressions at the aggregate--but not necessarily individual--level, and (c) culture-specific display tendencies explain dramatic differences in the intensity of facial movements. These results greatly sharpen our understanding of how people in different cultures use their faces to express emotion.

The primacy of specific feelings in emotional experience

Preliminary to our investigation is the question of what people in multiple cultures have in common when it comes to the experience of emotion. When shown emotionally evocative videos--newborn babies laughing, animal carcasses decaying, thunderstorm-filled skies, birds solving puzzles--to what extent do participants in different countries report having similar feelings? How many feelings are preserved both in subjective experience across cultures and in the English and Japanese lexica relied on to communicate these experiences? Understanding cultural commonalities in emotional experience is central to understanding human emotion³ and preliminary to investigating whether emotional experiences are expressed in similar facial movements across cultures.

We collected self-reported emotional responses to each of 2,185 evocative videos

using surveys in English (averaging 18 responses per video, largely in U.S./Canada and Europe) and Japanese (averaging 7.6 responses per video, all in Japan) (Fig. 1). Participants selected from 34 emotion categories⁸ and provided 1-100 intensity ratings for each selection, or rated valence and arousal, each on a 1-9 scale derived from constructivist approaches²⁶. Surveys were translated into Japanese by a team of bilingual researchers (see Methods).

We first asked how reported emotional experiences were preserved across languages and cultural groups. To do so, we computed correlations between cultural groups in ratings of emotion categories and affective features, adjusted for within-culture variation (see Methods). As shown in Figure 2A-E, a wide range of reported emotional experiences were well-preserved across the English and Japanese surveys, many with correlations exceeding .9 (e.g., “aesthetic appreciation”/“美的感動”, “amusement”/“愉快”, “fear”/“恐怖”, “interest”/“興味”, “sadness”/“悲しみ”). As in studies of emotion-related experience, perception, and brain response⁸⁻¹⁵, these specific feeling states were generally better preserved than reports of valence ($r = .74$, SE .035) and arousal ($r = -.065$, SE .062) of evoked experience, consistent with the latter being higher-order, more culture-specific evaluations^{9,11,27}. The pronounced differences in reported arousal are consistent with other studies comparing English and Japanese speakers²⁰⁻²² (see also Extended Data Fig. 1).

Figure 1. The Berkeley Reactions to Affective Video Elicitors (BRAVE) database. 2185 short evocative videos were shown to participants in multiple countries, who recorded their reactions and reported on their emotional experiences. Each of 45,231 resulting reaction videos was annotated by human raters (2 ratings per 2s segment) in terms of 42 perceived emotions and by a deep neural network in terms of 30 predicted human judgments (6 annotations per second).

Next, we ask how many distinct emotional experiences were captured in both the English-language and Japanese survey responses. For instance, “amusement”/“愉快” and “joy”/“喜び” were well-preserved across languages ($r=.92$ [SD .03], $.89$ [SD .05]), but were they used interchangeably? To determine how many different emotions were evoked in both cultures by distinct videos, we applied principal preserved components analysis, or PPCA^{9,11,28}, which extracts the distinct patterns (reported emotions) that covary across equivalent ratings collected from different groups (see Methods and Video S2 to understand why PPCA suits this purpose rather than PCA). At least 21 distinct components of subjective experience were preserved across English- and Japanese-language emotion judgments ($q<0.005$ across all held-out raters, $q<0.05$ across held-out raters in each culture individually, ForwardStop sequential FDR-corrected²⁹ one-tailed Wilcoxon signed-rank test³⁰). Thus, at least 21 different kinds of self-reported emotional experience were evoked by the 2,185 videos in different cultural groups and captured in both languages.

To characterize the 21 kinds of subjective experience that were captured in both English-language and Japanese responses, we applied factor rotation (varimax) to the 21 significant components extracted using PPCA. This method extracts a more interpretable representation of the 21 components by rotating them onto dimensions that each correspond to as few emotions as possible. Each of the 21 resulting

dimensions loaded maximally on a distinct category of emotion reported in response to the 2,185 videos (Fig. 2B). These dimensions can therefore be interpreted as 21 distinct varieties of emotional experience that are evoked by similar antecedents in multiple cultures and are present in both the English and Japanese lexica.

Figures 2C-E visualize the average emotional experiences associated with individual videos and how they are preserved across multiple cultures and languages along the 21 dimensions. We can see visually that the videos are richly diverse in the emotions they evoke, that these emotions are quite similar on average across the English and Japanese surveys (although there was country-level variation within the English survey responses), and that many videos evoke blends of emotion (e.g., “amusement” and “joy”, “anger” and “fear”, or “aesthetic appreciation” and “interest”). This last observation is consistent with recent work establishing that reported emotional experiences do not fall into discrete categories, as is commonly assumed³¹, but occupy smooth gradients of meaning^{8,11}.

The rich variations in emotional experience signaled in aggregate facial expression

We next explored how the 21 emotions uncovered in English-language and Japanese self-report corresponded to patterns of facial movement. The understanding of how facial movements express emotion has been pursued largely in the search for one-to-

one mappings between the six “basic” emotions and prototypical facial expressions^{1,2,4–6}. Here we apply discovery-based statistical methods to learn how the 21 dimensions of emotional experience captured in both English-language and Japanese responses may correspond to dozens of distinct patterns of facial movement.

To investigate the expression of emotion, we collected 45,231 self-recorded reactions to the 2,185 evocative videos. This included, on average, 14.5 reactions to each video from English-language survey participants and 5.6 reactions to each video from Japanese survey participants. The facial expressions in each reaction video were annotated both by human raters in the U.S. (see Discussion for cultural limitations) and using a deep neural network (DNN). Two human raters annotated each 2s video segment from each recording with 42 categories of perceived emotion, including the 34 categories used in self-report and 8 others derived from studies of emotions perceived in the face^{13,18}, for a total of 648,400 multiple choice ratings. Separately, a DNN was used to predict 30 emotion labels attributed to dynamically moving faces (from¹³, plus “boredom” and “neutral”). It was trained on human annotations (by English speakers in India) and relies only on pixels from the face (see Methods and ref.²⁵ for details). To explore how facial expressions relate to emotional experiences evoked by each video, human and DNN annotations were averaged across each recorded reaction. We also computed the maximum of each

DNN annotation across each recording. This resulted in a total of 102 annotations per recording: 42 averaged human annotations, 30 averaged DNN annotations, and 30 maximal DNN annotations. (Results using human or DNN annotations alone were similar; see Extended Data Fig. 2.)

To examine how facial movements related to reported emotional experience, we used the 102 features characterizing facial movements in reaction to each video to predict the emotional experiences it evoked. We did so by applying linear regression in a leave-one-out fashion across the 2185 videos (see Methods). To compare specific cultures, we fit separate models in the U.S. and Japan--the two most represented countries--and examined prediction accuracy at the individual and aggregate levels.

Predictions of reported emotional experience from expression were modestly accurate at the individual level. As indicated in Figure 3A, with single recordings of participants responding to each evocative video, one can predict the average within-culture emotional experience evoked by each video with an overall average correlation of .26 (SE .01) in the U.S. and .14 (SE .02) in Japan (between flattened matrices of actual and predicted self-report ratings projected onto the 21 dimensions; see Methods for details). Individual facial reactions gave us only modest insight into what the average person experienced when watching each video (see Video S1 [<https://is.gd/8GsKm8>] for example predictions). Further analyses explain why: because reactions were highly

Figure 2. The experience of emotion in multiple cultures and languages. **A. Specific feelings evoked by video were well preserved across English-language and Japanese self-report.** To examine whether reported emotional experiences were preserved across multiple cultures and languages, we computed correlations between ratings in English-language and Japanese responses and adjusted for within-language variation (see Methods). Reports of many specific categories of emotion (e.g., “interest”) were better preserved than ratings of valence and arousal. **B. 21 distinct dimensions of emotional experience were preserved across surveys.** To examine how many distinct self-reported emotions were preserved, we applied principal preserved components analysis (PPCA) between ratings across languages. This analysis revealed 21 significant dimensions ($q < 0.005$ across all held-out raters, $q < 0.05$ across held-out raters in each culture individually, ForwardStop sequential FDR-corrected²⁹ one-tailed Wilcoxon signed-rank test³⁰). The dimensions were subjected to factor rotation to characterize the different kinds of self-reported emotional experience evoked by videos, captured in both English-language and Japanese responses. **C-E. The distribution of emotional experience across videos is similar in the U.S. and Japanese responses.** To visualize the distribution of emotional experiences, we applied uniform manifold approximation and projection (UMAP³²) to the concatenated mean ratings in each language. Each video is assigned a letter corresponding to its maximal dimension (from B) and a color based on its average rating across all responses (C) or within the U.S. or Japan only (D/E). Emotional experiences were remarkably similar across different languages and countries, despite lexical and cultural differences. (See also the interactive map: <https://tinyurl.com/yywa7kjf>.) **F. Country-wise similarity of self-report judgments.** Responses in Poland and Portugal are slightly less similar to other countries, possibly due to the survey being administered in a non-native language (though as we will see, findings are consistent for facial expression [Fig. 4C]). Normalization of self-report values (subtracting mean and dividing by standard deviation) within each country slightly increases correlations between English-language and Japanese responses, a pattern we will see reflected more strongly in nonverbal expression (Fig. 4C). CA: Canada, JP: Japan, OE: Other Europe, Pol: Poland, Por: Portugal, UK: United Kingdom, US: United States.

variable across individuals, as is often the case^{6,7}. As indicated in Figure 3B, a recording of a single U.S. participant reacting to an evocative video captured, on average, only 14.5% of the variance in expression relevant to predicting average experience across all U.S. participants. This yields a maximum prediction correlation of .38 ($\sqrt{.145}$) due to individual variability alone (alternatively, the signal-to-noise ratio would be $.145/[1-.145]$, or 17.0%). Single recordings in Japan were even more variable, capturing just 3.1% of the variance in expression relevant to predicting average experience across Japanese survey participants, corresponding to a maximum prediction correlation of .18.

However, our findings indicate that by averaging expressive responses to each video across a sufficient number of participants, average reported emotional experience can be predicted with remarkable accuracy. As the number of participants increases, the overall prediction correlation across emotions converges at .70 (SE .02) in the U.S. and .80 (SE .05) in Japan using linear models trained to predict reported experience from the expression annotations within each country (Fig. 3A). With a model trained on data from all countries, several emotions (“amusement”, “disgust”, “joy”, “surprise”) can be predicted with correlations of over .9 (Fig. 3C). At the aggregate level, the emotional expressions evoked by videos are richly informative of subjective experience.

To what extent is the information represented by facial expression shared across multiple cultures? We assessed the degree of

cross-cultural convergence in experience and expression covariation by training models to predict emotional experience in one country from expression in another. Specifically, we predicted the average experience reported by U.S. participants in response to each video from the average expression annotations of Japanese participants, and vice versa. Using average facial expression in response to each video in Japan to predict average experience in the U.S., we find that the prediction correlation converges at a very similar value ($r = .65$, SE .03) to that achieved using U.S. facial expressions ($r = .70$, SE .02), reflecting similarities in the information captured by facial expression across countries. The same is true when predicting average emotional experience in Japan from average expression in the U.S. ($r = .72$, SE .03) compared to using Japanese expressions ($r = .80$, SE .05). Much of the information about experience conveyed by facial expression is shared across cultures.

Out of the 21 dimensions of emotional experience we uncovered, how many can be differentiated from facial expression? Our evidence, which accounts for 72% of the variance in expression relevant to average emotional experience, can provide a lower bound. By applying PPCA between predicted and reported emotional experiences, we found that at least 12 distinct dimensions of emotional experience were predicted from facial expression with significant accuracy ($q < 0.05$ across held-out raters, ForwardStop sequential FDR-corrected²⁹ one-tailed Wilcoxon signed-rank test³⁰). By examining

these 12 dimensions in Figure 3D we can see which experiences mapped onto distinct facial expressions. For instance, “aesthetic appreciation” and “romance” were combined into one dimension, suggesting that these feelings are expressed in similar facial movements. The same is true for “fear” and “anxiety”, “joy” and “relief”, and “sadness” and “anger”. With more evidence, we might uncover more granular distinctions in how people move their faces when they feel emotions, but for now conclude that some emotions evoked by video are easier to differentiate from facial expression than others. When we visualize the predictions of average experience along 12 broader dimensions captured by facial expression (Fig. 3E-G), we find that they predict reported experience with remarkable accuracy at the aggregate level.

Expression across cultures: Common meaning but varying display tendencies

We next asked how the expression of emotion differs between countries. Previously, we found that similar information about emotional experience was carried by facial expression in the U.S. and Japan (Fig. 3A), but it could have been carried by different patterns of facial movement. Did U.S. and Japanese participants move their faces in similar ways when they felt similar emotions^{1,2}? To answer this question, we trained models to predict experience from expression in one country and applied them to the other country (Fig. 4A). The performance of these models indicates how well the

mapping from facial movements to emotional experiences translates across the U.S. and Japan. We found that a model fit to Japanese responses and applied to U.S. expressions did a poorer job of predicting U.S. experience ($r = .54$, SE .02) than the model that was trained and evaluated within the U.S. ($r = .70$, SE .02). A model fit to U.S. responses and then applied to Japanese expressions did a far poorer job of predicting Japanese experience ($r = .01$, SE .03) than the model that was trained and evaluated within Japan ($r = .80$, SE .05). In other words, there are systematic differences in how people in the U.S. and Japan express emotion with their faces, requiring translation across cultures.

All emotion theories predict variations in emotional expression across cultures, but they posit different kinds of variation. According to some theories, expressions differ fundamentally in meaning across cultures--the same facial movements mean very different things¹. Other theories hold that facial expressions have similar meanings across cultures but different norms regulating their use, especially the intensity with which different emotions are expressed^{33,34}. These norms are generally referred to as “dialects”³⁵ or “display rules”³³, but here, we use “display tendencies”, a term that more clearly acknowledges individual variation. To what extent are U.S. and Japanese expressions fundamentally different in meaning, and to what extent are they merely different in terms of display tendencies, such as the intensity of the facial movements with which different expressions tend to be displayed^{33,34}?

To answer this question, we again trained models to predict experience from expression in one culture and applied them to the other culture, but we first normalized the 102 features characterizing facial movements within each culture (subtracted the mean and divided by the standard deviation). This adjusts for cultural variation in the intensity of each pattern of expression. We found that this simple procedure eliminated much of the cultural variation that we had previously observed (Fig. 4B). A model trained on normalized Japanese expressions and applied to normalized U.S. expressions predicted U.S. experience more accurately ($r = .61$, SE .03) than the non-normalized model ($r = .54$, SE .02). A model trained on normalized U.S. expressions and applied to normalized Japanese expressions predicted Japanese experience far more accurately ($r = .67$, SE .04) than the non-normalized model ($r = .01$, SE .03). Based on the performance of these models relative to that of models fit and evaluated within the same culture, the majority of the cultural variation in how expressions map to emotional experience can be attributed to differences in intensity (Fig. 5A). Without accounting for intensity differences, only 16% of the variation in the meaning of facial expressions was shared between the U.S. and Japan. After accounting for intensity differences, a full 81% of the variation in the meaning of facial expressions was shared across cultures. A PPCA analysis revealed that this shared variance was carried by at least 11 distinct preserved dimensions; $p \leq .016$, $q < 0.05$ across held-out raters,

ForwardStop sequential FDR-corrected²⁹ one-tailed Wilcoxon signed-rank test³⁰; see Extended Data Fig. 3).

Are these differences in expression intensity unique to comparisons between the U.S. and Japan, or can they also be found among different countries in North America and Europe? To answer this question with precision (in countries from which we collected less data than the U.S./Japan), we trained a combined model on all survey responses and applied it to expressive responses in each country (in a leave-one-stimulus-out fashion). We then correlated the model predictions across countries before and after normalizing the predictions (de-meaning and dividing by the standard deviation). Note that the model predictions are linearly dependent on the facial expression annotations. Thus, if the relative intensity of different patterns of facial movement is similar in two countries, then normalization will have little to no effect on the correlation in model predictions between the countries. However, if there are systematic differences in expression intensity, then normalization will serve to increase the correlation. The results of this analysis are striking. As shown in Fig. 4C, normalization makes virtually no difference when comparing among the U.S., Canada, U.K., Poland, Portugal, and other European countries. However, when comparing any of these countries to Japan, normalization greatly increases the correlation. Clearly, display tendencies are similar across North American and European countries, but different in Japan.

Figure 4. Intensity normalization is required to translate expressions across cultures. **A. Culture-insensitive models.** To see how well expressions are preserved across the U.S. and Japan, models were trained in one country to predict experience from expression and then applied to responses in the other. Prediction correlations were lower than those of within-country models (Fig. 3A), pointing to cultural differences in the mapping from experience to expression. **B. Normalization improves prediction by accounting for differences in expression intensity.** The same models were trained and tested after normalizing expression intensity within each country (demeaning and dividing by standard deviation). Prediction correlations improved dramatically. Cultural differences in the mapping from experience to expression are largely accounted for by the intensity with which expressions are displayed. **C. Country-wise similarity of raw and normalized expression.** A model was trained with all data to predict experience from raw expression, then applied to data from each of five individual countries. Predictions were correlated pairwise between countries, before and after normalization. Comparisons between North American/European countries and Japan (red) were dramatically more affected by normalization than those among different North American/European countries (blue). Thus, Japan differs systematically in expression intensity. After normalization, we can see to the right that the correlation in expression between countries closely reflected the correlation in experience ($r = .81$, despite noise in each estimate), suggesting that, aside from display tendencies, cultural differences were largely driven by differences in experience rather than differences in what facial expressions mean. CA: Canada, JP: Japan, OE: Other Europe, Pol: Poland, Por: Portugal, UK: United Kingdom, US: United States. **D. Differences in intensity of each emotional expression between Japan and the U.S./Canada/Europe.** A model was trained with all data to predict experience from raw expression, then standard deviations were computed for each prediction across responses in Japan and the U.S./Canada/Europe. Expressions of many emotions (e.g., "joy", "disgust") were much more intense in the North American and European countries than in Japan.

What exactly is the nature of the differences in intensity between expressions in Japan and the North American and European countries? To answer this question, we computed the standard deviation in model predictions within Japan and the U.S./Canada/Europe. This analysis enabled us to compare the average intensity of the facial movements that express each emotion in each group. As expected, the results of this analysis reveal considerable cultural variation in the intensity of certain expressions. As shown in Figure 4D, expressions of “amusement”, “disgust”, “joy”, and “romance” were over 60% more pronounced in the U.S./Canada/Europe. By contrast, only a few expressions were more pronounced in Japan, and none by more than 15%. These findings suggest that many facial expressions in Japan were substantially more nuanced than in North America and Europe^{20,34,36}. (To ensure that these findings could not be explained by cultural bias in facial expression annotations, we confirmed that the average expressions in the U.S./Canada/Europe also involved substantially more physical movement of the face than those in Japan; see Extended Data Fig. 4. We also verified that the annotations were not biased by racial group; see Extended Data Fig. 5.)

Expressions vary as a function of individual, not group, experience

Thus far we have explored how facial expressions vary across cultures in meaning--that is, in the mapping between facial expression and emotional experience (Fig.

5A). We have not explored the extent to which differences in expression arise directly from cultural and individual variations in the experiences evoked by the same videos--for instance, whether different people laugh at different things because they have different senses of humor.

To answer this question, we evaluated how well individuals’ reported experiences could be predicted from their own expressions. In doing so, there is no adjustment for sampling error, as it is unknown how reliably individuals reported on their emotions. Given that there were numerous options--selecting from 34 emotions and rating each on a 1-100 scale--it is likely that individual responses would have low test-retest reliability (even if stimuli could be repeated without being remembered). Hence, to gauge the impact of individual and cultural variation in emotional experience on expression, we assessed how well an individual’s self-reported experiences could be explained by their own expressions versus the expressions of random individuals from within the same country and those of random individuals from the other country. Across these comparisons, reliability is held constant, so it is valid to compare raw prediction correlations in relative terms.

We found that expressions were largely driven by individual variation in emotional experience (Fig. 5B). Among U.S. and Japan participants, only 24% of the variance in reported experience explained by expression was shared, on average, by individuals from the other country. Another

Figure 5. Individual, cultural, and universal variation in expressive response to videos in English-language and Japanese survey participants. A. Intensity normalization largely accounts for variation in the meaning of expressions across cultures. Percentage variance explained by models trained within the same versus other culture (summarizing results in Fig. 4A-B and a within-culture comparison model; see Methods: Model Comparison). Prediction correlations from each model were averaged across cultures and squared to compute explained variance. Just 16% of the variance in experience that was explained by within-culture models was explained by models trained in one culture and then applied to the other. However, after intensity normalization, 81% of the variance was explained. Thus, the meaning of expressions is largely preserved across cultures once intensity differences are accounted for. **B. Individual variability in emotional experience accounts for the bulk of the systematic variation captured in expression.** Fig. 1A is concerned with the mapping from experience to expression. It ignores variations in what different people experience when they view the same videos (by predicting experience from expression within culture and adjusting for sampling error across individuals). Here, we examine the degree to which expressive behavior varied as a function of individual experience by applying a single model trained on all data to individual-level expressive responses. To measure how individual, cultural, and universal variation in experience affected expression, we compared how well individuals' self-reported experiences could be explained by their own expressions versus the expressions of individuals from within the same culture and those of individuals from the other culture. 19% of the individual-level variance in experience explained by expression was shared by individuals in the other culture, 11% was culture-specific, and the remaining 70% corresponded to individuals' unique emotional experiences.

6% was shared by individuals within the same country. The remaining 70% corresponded to each individual's unique emotional experiences. Ultimately, how an individual moves their face is diagnostic of their own experience (i.e., people have different senses of humor). This partly explains why single recordings of individual participants' responses to each video predict the population average reported experience with modest accuracy (Fig. 3A). However, the question of how accurately expressions predict individual experience remains difficult, given that we cannot estimate the reliability of individual self-report judgments. While further work is needed to investigate how facial expressions vary at the individual level, the present study establishes that reported experience can be accurately inferred from facial expression at the aggregate level.

Discussion

The thesis that emotional experiences can be inferred from facial movements, so central to theorizing within the field of emotion science¹⁻³, has been only sparsely investigated, and largely with a focus on six prototypical expressions in small samples⁵⁻⁷. By analyzing thousands of spontaneous reactions to diverse emotion antecedents, here we paint a detailed portrait of how people in different cultures actually express emotion. In response to 2185 evocative videos, people in diverse cultures report at least 21 distinct varieties of emotional experience, which are best conceptualized using specific emotion categories (Fig. 2). Average experiences can be predicted with remarkable accuracy from

expressive responses to each video at the aggregate level (Fig. 3). Facial movements largely have common meaning across multiple cultures, but are subject to differing display tendencies, with many facial expressions in Japan being more nuanced than those in North America and Europe (Fig. 4, 5A). At the single-subject level, expressions are highly variable, partly reflecting individuals' varied emotional experiences (Fig. 5B).

The findings of the present investigation speak directly to competing theories regarding how experiences and expressions of emotion are preserved across cultures. Constructivist theories prioritize appraisals of valence and arousal, holding that commonalities in the experience and expression of emotion across cultures will be differentially explained by these broad dimensions^{1,26,37,38}. By contrast, in Basic Emotion Theory and related approaches, specific feelings such as “interest” and “fear” are more primary in the experience and expression of emotion^{3,31,39}. Our findings support the primacy of specific emotional experiences in emotion-related behavior. Feelings such as “sadness” and “amusement” were better preserved in self-reported experience across cultures than evaluations of valence and arousal. These feelings were in turn associated with spontaneous facial movements, which were largely preserved in meaning across cultures (Fig. 5A; Extended Data Fig. 3) and conveyed not two but at least 12 dimensions of emotional experience (Fig. 3D-F).

The present findings also address hypotheses regarding the manner in which emotional expressions vary across cultures. Constructivist approaches hold that cultural variation in expression is observed because expressions are shaped by learned emotion concepts, which are assumed to be highly variable across cultures⁴⁰. By contrast, Basic Emotion Theory holds that cultural variation will be explained in large part by differential tendencies to display an expression, but that the expression itself can be conceptualized similarly across cultures^{33–35}. Our findings more strongly support the Basic Emotion view, in that differences in facial movement between cultures were in large part driven by variations in the intensity with which different expressions were displayed (Fig. 5B). With intensity normalized, the mapping from expression to experience was largely preserved across Japan and the U.S. (Fig. 5A). Across these cultural groups, facial expressions are subject to differing display tendencies, but have similar meaning.

Finally, the above findings are made possible in part by similarities in the English and Japanese lexica relied on to communicate emotional experiences (Fig. 2). These similarities between English and Japanese emotion concepts may seem puzzling, given past studies have found differences in how emotion-related words are abstractly related to one another in different languages^{41,42}. However, few studies have compared the meanings of emotion-related concepts across languages based on how they are actually used to describe experiences in different

contexts. By doing so, we find that the English and Japanese lexica are aligned well enough to convey at least 21 shared dimensions of emotional experience, many with correlations between cultures exceeding .9 across the 2185 videos (Fig. 2A). These findings suggest that abstract studies of broad lexical associations may overlook the more specific meanings conveyed by emotion-related concepts, grounded in their associations with actual referents in the world.

This study had several limitations. To characterize facial expressions, we used human and DNN annotations that were based on English-language categorizations of perceived emotion. Although these annotations were broadly predictive of aggregate emotional experience in both English-speaking countries and Japan (Fig. 3A, 4B), it is still important to note that they were constrained by the culture-specific knowledge of the raters used for manual annotation and for training the DNN (see Methods). For instance, Japanese categorizations of perceived emotion could potentially reveal additional dimensions of emotional response, potentially with different display tendencies. Such research could build on our findings and provide further insight into how people in different cultural groups express (and perceive) emotion. On a similar note, it is worth acknowledging that cross-cultural similarities in emotion-related behavior may be explained in part by cultural globalization. To further examine human universals in expression, it would be fruitful

to apply similar methods in remote, small-scale cultures, although such work would likely be more constrained in scale. Finally, we relied on videos to evoke emotion, but more personal forms of emotion elicitation (e.g., social interaction) would likely reveal further dimensions of emotional expression²⁵.

The present study focused on facial expressions formed in isolation, but our everyday understanding of facial expression is sensitive to context. Other studies could incorporate vocal and bodily expression to further characterize the complex ways in which people across cultures express emotion. More broadly, social context matters. Here, participants viewed evocative videos alone, and although they were encouraged to use their face expressively, emotional expression would likely vary in the presence of others⁴³, an important topic for further study.

The defining moments of our lives--romance, friendship formation, loss, discovery, betrayal--are rife with emotional expression. Yet we have surprisingly little understanding of the meaning of expressions. A narrow focus on whether six emotions map to prototypical facial movements has led to skepticism among some researchers that emotions can be inferred from facial expression¹. Open-ended methods and large-scale evidence paint a more comprehensive picture--a detailed portrait of the complex ways in which people actually move their faces in response to thousands of emotionally evocative scenes. Across diverse cultural groups in North America, Europe, and Japan,

facial expressions reflect a broad array of specific feelings, are similar in meaning, and are subject to varying display tendencies.

Methods

Participants. Participants of the English-language survey were recruited via Amazon Mechanical and Prolific. Japanese participants were recruited via Crowdfunder. On average, a total of 18.7 English-language and 7.7 Japanese survey participants provided ratings and/or reactions for each of the 2,185 videos. English-language survey participants (n=1332, 678 male [124 unreported], mean age 31.8) responded to 30.7 videos on average and Japanese survey participants (n=61, 39 male [6 unreported], mean age 38.4) responded to 274.3 videos on average. Given that the average Japanese participant responded to a greater number of videos, it is important to note that the intensity of their expressions did not decline over time (Extended Data Fig. 6). English-language survey participants originated from the U.S. (n=568), U.K. (n=158), Portugal (n=96), Poland (n=94), Canada (n=55), Mexico (n=35), Italy (n=29), Greece (n=28), Spain (n=26), France (n=11), Nigeria (n=10), Australia (n=8), New Zealand (n=8), India (n=6), Israel (n=3), other Europe (n=62), other Latin America (n=16), and other Africa (n=6) (113 unreported; 9 from East or Southeast Asia were excluded). The U.S./Canada/Europe sample used in certain analyses comprises 92% (n=1127) of English-language survey participants who reported their country of origin. All Japanese survey participants originated from Japan (6 unreported).

Surveys. In a given survey, participants viewed 15-30 randomly selected videos and used a webcam plugin to record their reactions. Each video was repeated so that its final duration was at least 8 seconds. Recording began at the start of each video and was stopped 1s after the end of each video. Participants were encouraged to use their face expressively. In the emotion category survey, participants selected from 34 emotion categories to describe their response to each video and reported the intensity of each selection on a 1-100 scale. In the valence/arousal survey, participants rated the valence and arousal evoked by each video on 1-9 Likert scales. Participants on Amazon Mechanical Turk and Crowdfunder could perform as many surveys as desired with different videos in each. Participants on Prolific were limited to one survey. The experimental procedures were approved by the Institutional Review Board at the University of California, Berkeley. All participants gave their informed consent to participate in experiment and consent to publish their identifiable images. See Supporting Information for further details regarding the surveys. Surveys were translated into Japanese by a team of bilingual researchers (see

Acknowledgements, with particular care provided in translating emotion terms and questions about valence and arousal (see Table S1 for references).

Correlations between cultural groups. We computed the correlations between ratings of each category and affective feature across cultural groups (countries or survey languages). To do so, for each category or affective feature, we first drew individual ratings at random from each cultural group. One rating or score was drawn for each of the 2185 videos. For instance, when comparing ratings of “disgust” one rating of each video was drawn from each cultural group (each ranging from 0-100). We then computed the correlation in the ratings between the cultural group, across the 2185 videos. Finally, we divided by the maximum attainable correlation, or noise ceiling, based on within-culture variability of the ratings. As explained further below, the maximum attainable correlation is estimated by computing the geometric mean of the correlations between individual ratings drawn at random from within the same cultural group⁴⁴.

Principal preserved component analysis (PPCA). PPCA extracts shared dimensions that maximize the covariance between two parallel datasets (e.g., emotions ratings). To do so, PPCA first seeks a unit vector α_1 that maximizes the objective function

$$\text{Cov}(X\alpha_1, Y\alpha_1)$$

in which we call α_1 the first principal preserved component. Subsequent components are obtained by seeking additional unit vectors α_i that maximize the objective function $\text{Cov}(X\alpha_i, Y\alpha_i)$ subject to the constraint that α_i is orthogonal to the previous components, $\alpha_1, \dots, \alpha_{i-1}$.

In the special case that $X = Y$, PPCA is equivalent to PCA, given that the latter method maximizes the objective function

$$\text{Var}(X\alpha_i) = \text{Cov}(X\alpha_i, X\alpha_i)$$

(substituting another X for Y in $\text{Cov}[X\alpha_1, Y\alpha_1]$). See Video S2 for illustration of why the PCA objective is ill-suited to our aims. Also note the similarity to the PLSC objective, which seeks to find two separate bases and β to maximize $\text{Cov}(X\alpha_i, Y\beta_i)$, as well as the CCA objective, which seeks to maximize $\text{Corr}(X\alpha_i, Y\beta_i)$. However, given our aim of finding *preserved* dimensions across ratings (between cultures, or predicted vs. reported), PPCA derives only one basis, α , that applies to both datasets. (In PPCA, therefore, the data matrices must be commensurate: observations in both datasets must be of the same dimension; i.e. the number of rows in X and Y must be equal.)

To solve the PPCA objective and find an α_1 we apply eigendecomposition to the addition of the cross-covariance matrix between datasets and its transpose: $\text{Cov}(X, Y)/2 + \text{Cov}(Y, X)/2$. We claim that the principal eigenvector of this symmetric matrix maximizes $\text{Cov}(X\alpha_1, Y\alpha_1)$. To derive this, first recall a

general property of cross-covariance, $\text{Cov}(Xa, Yb) = b^T \text{Cov}(X, Y)a$. Thus,

$$\text{Cov}(X\alpha_1, Y\alpha_1) = \alpha_1^T \text{Cov}(X, Y) \alpha_1$$

(Property 1)

In addition, because both $X\alpha_1$ and $Y\alpha_1$ are vectors, $\text{Cov}(X\alpha_1, Y\alpha_1) = \text{Cov}(Y\alpha_1, X\alpha_1)$. Thus,

$$\text{Cov}(X\alpha_1, Y\alpha_1) = \text{Cov}(X\alpha_1, Y\alpha_1)/2 + \text{Cov}(Y\alpha_1, X\alpha_1)/2$$

(Property 2)

Combining these two properties, we can see that

$$\begin{aligned} \text{Cov}(X\alpha_1, Y\alpha_1) &= \text{Cov}(X\alpha_1, Y\alpha_1)/2 + \text{Cov}(Y\alpha_1, X\alpha_1)/2 \\ &\quad \text{(By property 2)} \\ &= \alpha_1^T \text{Cov}(X, Y) \alpha_1/2 + \alpha_1^T \text{Cov}(Y, X) \alpha_1/2 \\ &\quad \text{(By property 1)} \\ &= \alpha_1^T [\text{Cov}(X, Y)/2 + \text{Cov}(Y, X)/2] \alpha_1 \end{aligned}$$

Now, letting $R = [\text{Cov}(X, Y)/2 + \text{Cov}(Y, X)/2]$, we see that maximizing $\alpha_1^T R \alpha_1$ is equivalent to maximizing $\text{Cov}(X\alpha_1, Y\alpha_1)$, the originally stated PPCA objective. (Note that if $X = Y$, we are applying eigendecomposition to $\text{Var}[X\alpha_i] = \text{Cov}[X\alpha_i, X\alpha_i]$, which performs PCA.) Finally, the min-max theorem dictates that the principal eigenvector of R maximizes $\alpha_1^T R \alpha_1$ subject to α_1 being a unit vector ($\|\alpha_1\|=1$)

We have thus found a unit vector α_1 that maximizes $\text{Cov}(X\alpha_1, Y\alpha_1)$ —the covariance between the projections of X and Y projected onto the first component. Based on the min-max theorem, subsequent eigenvectors α_i will maximize $\text{Cov}(X\alpha_i, Y\alpha_i)$ subject to their orthogonality with previous components α_1 through α_{i-1} and to each α_i also being a unit vector ($\|\alpha_i\|=1$).

We note that the min-max theorem also provides that the last eigenvector, α_n , will minimize $\text{Cov}(X\alpha_n, Y\alpha_n)$, equivalent to maximizing $-\text{Cov}(X\alpha_n, Y\alpha_n)$. Hence, if there are dimensions of negative covariance between the two datasets, then some eigenvectors will maximize the negative covariance.

With respect to the corresponding eigenvalues, each eigenvalue λ_i will be equal to $\text{Cov}(X\alpha_i, Y\alpha_i)$. To see this, note that:

$$[\text{Cov}(X, Y)/2 + \text{Cov}(Y, X)/2] \alpha_i = \lambda_i \alpha_i$$

(Eigenvalue equation)

$$\begin{aligned} \alpha_i^T [\text{Cov}(X, Y)/2 + \text{Cov}(Y, X)/2] \alpha_i &= \alpha_i^T \lambda_i \alpha_i \text{Cov}(X\alpha_i, Y\alpha_i) \\ &= \lambda_i \alpha_i^T \alpha_i \quad \text{(By property 1)} \end{aligned}$$

Now $\alpha_i^T \alpha_i = 1$ because the α_i are orthonormal. Hence,

$$\text{Cov}(X\alpha_i, Y\alpha_i) = \lambda_i$$

This also entails that there will be negative eigenvalues corresponding to negative covariance.

We performed PPCA between the averaged self-report ratings from each culture (Fig. 2) or

between the averaged self-report ratings across cultures and the predicted self-report ratings from our model (Fig. 3D-F). To ascertain the number of significant dimensions of covariance, we performed a leave-one-subject-out analysis, in which PPCA was iteratively performed on data from all but one participant and then the held-out ratings were projected onto the extracted dimensions and correlated. Partial Spearman correlations were used, controlling for projections onto previous dimensions, to account for possible curvilinear relationships. Wilcoxon signed-rank tests were then applied to the correlations for held out raters to test each dimension for significance. See Cowen et al.^{9,11} for results of repeated Monte Carlo simulations further validating these methods.

Facial expression deep neural network (DNN). A DNN is a machine learning algorithm that learns to apply multiple layers of transformation to input data in order to predict complex outputs. We used a supervised DNN that processed the temporal trajectory of the RGB pixel values making up a video of a face over the course of up to one second, at a rate of 6 frames per second, to predict the proportions of raters who would attribute each label to each expression. To extract faces from video, each face was detected within a video frame using a deep convolutional neural network, similar to the method provided by the Google Cloud Face Detection API. Faces were tracked across each video using traditional optical flow methods. **DNN architecture.** Face-based visual features were extracted using layers from the NN2 FaceNet architecture³⁶ to characterize mid-level attributes involved in early stages of face perception. These layers consisted of an Inception (5a) block with a 7x7 feature map comprising 1,024 channels, which was fed into a 7x7 average pooling layer, generating a 1,024 dimensional vector representing face image features within a given frame of the video. The resulting features were then fed into two long short-term memory (LSTM) layers, each with 64 recurrent cells, to capture the dependence of facial expression recognition on temporally unfolding patterns of facial movement. Finally, the output of the last LSTM layer was fed into a mixture-of-experts model (two experts, plus a dummy expert). A cross entropy loss with a sigmoid activation function was used for the final layer with thirty nodes.

Training. The DNN was trained on a total of 342,546 ratings of 247,292 clips of faces on YouTube. Raters were English speakers in India. The task was to select all facial expression categories that applied to each face. The videos from which the clips were extracted were largely collected manually by raters. These raters had been instructed to conduct a broad search for videos likely to contain emotional expressions.

Manual annotations of reaction videos. We collected manual annotations on Amazon Mechanical Turk of

each 2s segment extracted from the reaction videos (the final segment for each reaction was 1-3s). A total of 3293 raters participated (1599 male, mean age approx. 34.7 [based on decade ranges]). Each segment was rated by 2 raters, for a total of 581,393 ratings, in terms of 42 emotion categories (“Admiration”, “Adoration”, “Aesthetic Appreciation”, “Amusement”, “Anger”, “Anxiety”, “Awe”, “Boredom”, “Calmness”, “Concentration”, “Confusion”, “Contemplation”, “Contempt”, “Contentment”, “Desire”, “Disappointment”, “Disgust”, “Distress”, “Doubt”, “Ecstasy”, “Elation”, “Embarrassment”, “Empathic Pain”, “Entrancement”, “Fear”, “Horror”, “Interest”, “Love”, “Neutral”, “None”, “Nostalgia”, “Pain”, “Pride”, “Realization”, “Relief”, “Romance”, “Sadness”, “Satisfaction”, “Shame”, “Surprise”, “Sympathy”, “Triumph”). Raters were asked to describe the emotions felt by the person in the video by selecting all that applied. Annotations were averaged across each video.

Model comparison. Each model predicting experience from expression was cross-validated using a leave-one-video-out approach. Overall prediction correlations were computed by flattening and correlating the matrices of predicted and actual self-reported emotional experiences projected onto the 21 dimensions. Aside from the analysis in Fig. 5B, predicted experiences based on a given participant’s expressive response were always compared with actual self-reported experiences from a different participant, so that we could adjust for explainable variance (see below). Bootstrap resampling of participants was applied to estimate standard errors in prediction correlations.

In applying culture-specific models across cultures (Fig. 4A,B; Fig. 5A) an extra calibration was applied to the experience ratings in the culture where the model was evaluated. Specifically, we used linear regression in a leave-one-out fashion to calibrate experience ratings in the culture where the model was evaluated to the experience ratings in the culture where the model was originally fit. This was necessary because a particular English category (e.g., “sadness”) may not translate perfectly to its Japanese counterpart, thus lowering model fit not because the expression has a different meaning but because the predicted experience rating has a different meaning. The information present in the ratings will be weighted somewhat differently after this calibration step, so the maximum prediction correlation may change. Thus, to compare across-culture models to within-culture models (Fig. 5A), we trained new within-culture models after performing the same calibration.

Proportion of explainable variance and sampling error adjustment. When two sets of sample means are correlated, the correlation is biased downward by

sampling error^{44–47}. Thus, more ratings would lower the standard error and yield a higher correlation, even though “unbiased” estimates should not be biased by sample size. To correct for this bias, one can divide the sample correlation by the approximate maximum correlation that could be obtained given the sampling error. The maximum attainable correlation is the square root of the proportion of the variance that is not attributable to sampling error (the explainable variance).

Here, we adjust for sampling error at the level of individual responses (self-report ratings and/or facial expression annotations for a given video). In doing so, the maximum attainable correlation is given by the product of the index of reliability for each response, which is equal to the square root of the interrater reliability coefficient⁴⁴. The interrater reliability coefficient was estimated by randomly pairing two responses from separate participants to each video and computing the Pearson correlation across videos. This process was repeated 100 times and the average interrater reliability coefficient was computed.

Prediction correlations for models predicting average experience were adjusted for sampling error in experience by dividing by explainable variance in experience. Prediction correlations for models that used average expression as a predictor were adjusted for sampling error in average expression by dividing by explainable variance in the prediction based on individual expressive responses to the same video, to estimate the maximum prediction correlation that could be achieved with large participant samples.

Explainable variance in the average response is estimated for varying sample sizes in Fig. 3B. These were estimated using the formula $r_N^2 = ((1-r_1^2)/(N*r_1^2) + 1)^{-1}$, where r_1^2 represents the explainable variance of individual-level responses (the square of the interrater reliability coefficient) and r_N^2 represents the explainable variance of an average of N responses. This formula can be derived from the fact that the signal-to-noise ratio is given by $r_N^2/(1-r_N^2)$ and is directly proportional to sample size, such that $r_N^2/(1-r_N^2) = N*r_1^2/(1-r_1^2)$. By solving for r_N^2 we arrive at a simple formula for explainable variance as a function of N .

References and Notes

1. Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M. & Pollak, S. D. Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. *Psychol. Sci. Public Interes.* **20**, 1–68 (2019).
2. Cowen, A., Sauter, D., Tracy, J. L. & Keltner, D. Mapping the Passions: Toward a High-Dimensional Taxonomy of Emotional Experience and Expression. *Psychol. Sci. Public Interes.* **20**, 69–90 (2019).
3. Keltner, D., Sauter, D., Tracy, J. & Cowen, A. Emotional Expression: Advances in Basic Emotion Theory. *J. Nonverbal Behav.* **43**, 133–160 (2019).
4. Ekman, P. & Friesen, W. V. Constants across cultures in the face and emotion. *J. Pers. Soc. Psychol.* **17**, 124–129 (1971).
5. Elfenbein, H. A. & Ambady, N. On the universality and cultural specificity of emotion recognition: A meta-analysis. *Psychol. Bull.* **128**, 203–235 (2002).
6. Lench, H. C., Flores, S. A. & Bench, S. W. Discrete emotions predict changes in cognition, judgment, experience, behavior, and physiology: A meta-analysis of experimental emotion elicitations. *Psychol. Bull.* **137**, 834–855 (2011).
7. Durán, J. I., Reisenzein, R. & Fernández-Dols, J.-M. Coherence Between Emotions and Facial Expressions. in *The Science of Facial Expression* (eds. Fernandez-Dols, J.-M. & Russell, J. A.) **1**, 107–129 (Oxford U Press, 2017).
8. Cowen, A. S. & Keltner, D. Self-report captures 27 distinct categories of emotion bridged by continuous gradients. *Proc. Natl. Acad. Sci. U. S. A.* **114**, E7900–E7909 (2017).
9. Cowen, A. S., Laukka, P., Elfenbein, H. A., Liu, R. & Keltner, D. The primacy of categories in the recognition of 12 emotions in speech prosody across two cultures. *Nat. Hum. Behav.* **3**, 369–382 (2019).
10. Cowen, A. S., Elfenbein, H. A., Laukka, P. & Keltner, D. Mapping 24 Emotions Conveyed by Brief Human Vocalization. *Am. Psychol.* **22**, 274–276 (2018).
11. Cowen, A. S., Fang, X., Sauter, D. A. & Keltner, D. What music makes us feel: Uncovering 13 kinds of emotion evoked by music across cultures. *Proc. Natl. Acad. Sci.* (2020).
12. Cowen, A. S. & Keltner, D. Universal emotional expressions uncovered in art of the ancient Americas: A computational approach. *Sci. Adv.* (2020).
13. Cowen, A. S. & Keltner, D. What the Face Displays: Mapping 28 Emotions Conveyed by Naturalistic Expression. *Am. Psychol.* (2019). doi:10.1037/amp0000488
14. Horikawa, T., Cowen, A. S., Keltner, D. & Kamitani, Y. The neural representation of emotion is high-dimensional, categorical, and distributed across transmodal brain regions. *bioRxiv* (2019). doi:10.1101/872192
15. Saarimäki, H. *et al.* Distributed affective space

- represents multiple emotion categories across the human brain. *Soc. Cogn. Affect. Neurosci.* **13**, 471–482 (2018).
16. Shiota, M. N. *et al.* Beyond happiness: Building a science of discrete positive emotions. *Am. Psychol.* **72**, 617–643 (2017).
 17. Cordaro, D. T. *et al.* Universals and cultural variations in 22 emotional expressions across five cultures. *Emotion* **18**, 75–93 (2018).
 18. Cordaro, D. T. *et al.* The Recognition of 18 Facial-Bodily Expressions Across Nine Cultures. *Emotion* **20**, 1292–1300 (2020).
 19. Sauter, D. A. The Nonverbal Communication of Positive Emotions: An Emotion Family Approach. *Emot. Rev.* **9**, 222–234 (2017).
 20. Scherer, K. R., Wallbott, H. G., Matsumoto, D. & Kudoh, T. Emotional experience in cultural context: A comparison between Europe, Japan, and the USA. in *Facets of emotion: Recent research* (ed. Scherer, K. R.) 5–30 (Erlbaum, 1988).
 21. Kitayama, S., Markus, H. R. & Kurokawa, M. Culture, emotion, and well-being: Good feelings in Japan and the United States. *Cognition and Emotion* **14**, 93–124 (2000).
 22. Lim, N. Cultural differences in emotion: differences in emotional arousal level between the East and the West. *Integr. Med. Res.* **5**, 105–109 (2016).
 23. Rothbaum, F., Weisz, J., Pott, M., Miyake, K. & Morelli, G. Attachment and culture: Security in the United States and Japan. *Am. Psychol.* **55**, 1093–1104 (2000).
 24. Cowen, A. S. & Keltner, D. Clarifying the Conceptualization, Dimensionality, and Structure of Emotion: Response to Barrett and Colleagues. *Trends in Cognitive Sciences* **22**, 274–276 (2018).
 25. Cowen, A. S. *et al.* Sixteen facial expressions occur in similar social contexts recorded around the world. *Nature*
 26. Russell, J. A. Core Affect and the Psychological Construction of Emotion. *Psychol. Rev.* **110**, 145–172 (2003).
 27. Tsai, J. L. Ideal Affect: Cultural Causes and Behavioral Consequences. *Perspect. Psychol. Sci.* **2**, 242–259 (2007).
 28. Demszky, D. *et al.* GoEmotions: A Dataset of Fine-Grained Emotions. 4040–4054 (2020).
 29. G'Sell, M. G., Wager, S., Chouldechova, A. & Tibshirani, R. Sequential selection procedures and false discovery rate control. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **78**, 423–444 (2016).
 30. Wilcoxon, F. Individual Comparisons by Ranking Methods. *Biometrics Bull.* **1**, 80 (1945).
 31. Ekman, P. & Cordaro, D. What is meant by calling emotions basic. *Emot. Rev.* **3**, 364–370 (2011).
 32. McInnes, L. PCA, t-SNE, and UMAP: Modern Approaches to Dimension Reduction. *PyData Conference 2018* (2018).
 33. Ekman, P. *et al.* Universals and Cultural Differences in the Judgments of Facial Expressions of Emotion. *J. Pers. Soc. Psychol.* **53**, 712–717 (1987).
 34. Matsumoto, D. & Ekman, P. American-Japanese cultural differences in intensity ratings of facial expressions of emotion. *Motiv. Emot.* (1989). doi:10.1007/BF00992959
 35. Elfenbein, H. A., Beaupré, M., Lévesque, M. & Hess, U. Toward a dialect theory: Cultural differences in the expression and recognition of posed facial expressions. *Emotion* **7**, 131–146 (2007).
 36. Sato, W., Hyniewska, S., Minemoto, K. & Yoshikawa, S. Facial expressions of basic emotions in Japanese laypeople. *Front. Psychol.* **10**, (2019).
 37. Barrett, L. F. Valence is a basic building block of emotional life. *J. Res. Pers.* **40**, 35–55 (2006).
 38. Gendron, M., Crivelli, C. & Barrett, L. F. Universality Reconsidered: Diversity in Making Meaning of Facial Expressions. *Curr. Dir. Psychol. Sci.* **27**, 211–219 (2018).
 39. Scherer, K. R. & Wallbott, H. G. Evidence for Universality and Cultural Variation of Differential Emotion Response Patterning. *J. Pers. Soc. Psychol.* **66**, 310–328 (1994).
 40. Hoemann, K. & Feldman Barrett, L. Concepts dissolve artificial boundaries in the study of emotion and cognition, uniting body, brain, and mind. *Cogn. Emot.* **33**, 67–76 (2019).
 41. Jackson, J. C. *et al.* Emotion semantics show both cultural variation and universal structure. *Science (80-.)*. **366**, 1517–1522 (2019).
 42. Russell, J. A. Culture and the categorization of emotions. *Psychol. Bull.* **110**, 426–450 (1991).
 43. Jakobs, E., Manstead, A. S. R. & Fischer, A. H. Social Context Effects on Facial Activity in a Negative Emotional Setting. *Emotion* **1**, 51–69 (2001).
 44. Lord, F. M. & Novick, M. R. The Construction of True and Error Scores (Chapter 2). in *Statistical Theories of Mental Test Scores* 358–393 (1968).
 45. Spearman, C. The proof and measurement of association between two things. By C. Spearman, 1904. *Am. J. Psychol.* **100**, 441–471 (1987).
 46. Benjamini, Y. & Yu, B. The shuffle estimator for explainable variance in fMRI experiments. *Ann. Appl. Stat.* **7**, 2007–2033 (2013).
 47. Csárdi, G., Franks, A., Choi, D. S., Airolidi, E. M. & Drummond, D. A. Accounting for

Experimental Noise Reveals That mRNA Levels, Amplified by Post-Transcriptional Processes, Largely Determine Steady-State Protein Levels in Yeast. *PLoS Genet.* **11**, (2015).

Acknowledgements: We thank Mitsuaki Tsukamoto and Tomoyasu Horikawa, along with authors M.T. and Y.K., for translating surveys into Japanese. This work was supported in part by Google Research.

Author contributions: A.S.C. formulated and designed the study. A.S.C., M.T., Y.K., V.K., and D.K. developed survey tools and collected the data. A.S.C., G.P., K.S., B.J., F.S., and H.A. contributed analytic tools. A.S.C. analyzed the data. A.S.C., G.P., J.A.B., and D.K. wrote the paper with input from all other authors.

Extended Data

Extended Data Fig. 1. Distribution of valence and arousal ratings in U.S./Canada/Europe and Japan. Histogram of valence and arousal ratings in each cultural group, colored according to the highest average intensity emotion category in judgments of videos represented by each bin (within the culture). Close inspection of these histograms reveals why English-language and Japanese arousal ratings had little to no correlation. For example, videos that evoked “disgust” were often assigned low arousal ratings in the U.S./Canada/Europe but had almost uniformly high arousal ratings in Japan (whereas videos evoking “fear” were assigned high arousal ratings in both cultures). Videos that evoked “romance”, “aesthetic appreciation”, and most other positive emotions were assigned low arousal ratings in Japan. Generally, negative emotions were higher arousal in Japan than in the U.S./Canada/Europe, and positive emotions were higher arousal in U.S./Canada/Europe than in Japan. These findings are consistent with research on affect valuation in so-called “Western” and “Eastern” cultures²⁷ and more focused cultural comparisons, which consistently find that high arousal positive emotions are more common in North American and European countries than in Japan^{20–22,34}.

Extended Data Fig. 2. Human and DNN annotations alone each predicted emotional experience with only slightly lower accuracy than the combined model. Models were trained using linear regression to predict average experience in response to each of the 2185 videos from average expression annotations combined across cultures. Models based on human or DNN annotations alone performed only slightly worse than the full model, suggesting that human and DNN annotations largely overlap in the information they capture. Note that performance of the human annotation model is not adjusted for sampling error and could potentially improve with additional annotations.

Extended Data Fig. 3. Experience-expression association across the U.S. and Japan. A. Raw individual-level correlations between 21 dimensions of emotional experience and 102 DNN and human expression annotations in each country. Annotations are ordered by maximally-correlated dimension of emotional experience. Correlations are similar in the two cultural groups, albeit more nuanced in Japan. **B. Dimensions required to explain preserved experience-expression correlations across cultures.** PPCA was applied to the 102x21 matrices of correlations between cultures. This analysis was performed in a 13-fold cross-validated fashion across U.S. participants, iteratively leaving out data from 77 of the 1001 participants that contributed both experience judgments and self-recorded expressions. Correlations were separately computed between experience and expression in the held-in and held-out data. PPCA was then performed between the held-in U.S. correlations and the correlations in Japan. Finally, the held-out U.S. correlations and the Japanese correlations were projected onto the extracted components and correlated across cultures (partial Spearman correlations, controlling for previous dimensions). Eleven dimensions were significant ($*p \leq .016$, $q < 0.05$ across folds, ForwardStop sequential FDR-corrected²⁹ one-tailed Wilcoxon signed-rank test³⁰). Japanese data were not subdivided during cross-validation, because the data were more limited, but cross-

validation within the U.S. responses is sufficient to guarantee independence. With additional data, it is likely that further dimensions would be discovered.

Extended Data Fig. 4. Differences in average facial movement between U.S./Canada/Europe and Japanese participants. To measure physical movement of the face, we used a face tracking algorithm to locate 29 facial landmarks in each of the 45,231 reaction recordings on a frame-by-frame basis (2 frames per second). Across each reaction, we then computed the standard deviation in the average distance between each facial landmark and the three other facial landmarks to which it was most often positioned closest. We divided by the average distance in the video, yielding proportional measures of local facial movement. These measures were averaged within each cultural group across reactions to videos that most evoked each of the 21 dimensions of emotional experience. Facial movements in response to most videos were more physically pronounced in U.S./Canada/Europe participants, on average (the only minor exception being responses to “anger”-evoking videos). The difference is present in every region of the face, especially near the cheeks, mouth, and nose, where physical facial movements in response to “disgust”-evoking stimuli were more than 50% more pronounced in the U.S./Canada/Europe than in Japan. These findings confirm that the greater intensity we observe in average U.S./Canada/Europe expressions (Fig. 4C) cannot be explained by a cultural bias in facial expression annotations.

Extended Data Fig. 5. The expression DNN predicts human judgments and is largely invariant to demographics (from Cowen et al., *Nature*). a-c. Accuracy of the expression DNN in emulating human judgments. Human judgments (a) and the expression DNN annotations (b-c) have been projected onto a map of 1456 facial expressions adapted from Cowen & Keltner¹³. Human judgments and the expression DNN annotations are represented using colors, according to the color scheme used in Cowen & Keltner¹³. Likely because the expression DNN was trained on dynamic faces, it can in some cases make systematic errors in predicting judgments of static faces (c). For example, a number of static faces of “surprise” were more strongly annotated by the expression DNN as “awe” (c, bottom left), likely because dynamic faces that convey “surprise” are distinguished in part by dynamic movement. This problem is mitigated when the expression DNN is calibrated for still images (B). To calibrate the DNN,

multiple linear regression is applied in a leave-one-out-fashion to predict human judgments of the still images from the DNN annotations. After calibration on still images, we can see that the expression DNN annotations are fairly accurate in emulating human judgments (overall $r = .69$ between calibrated expression DNN annotations and human judgments after adjusting for explainable variance in human judgments). **d. The expression DNN can emulate human judgments of individual emotions and valence/arousal with moderate to high accuracy.** Individual expression DNN predictions are correlated with human judgments across the 1456 faces. Valence and arousal judgments (also from Cowen & Keltner¹³) were predicted using multiple linear regression in a leave-one-out fashion from the 16 facial expression DNN annotations. **e. The expression DNN is reliable for different demographic groups.** By correlating the expression DNN predictions (calibrated for static images) across subsets of the 1456 faces from Cowen & Keltner¹³ we can see that the expression DNN is accurate for faces from different demographic groups (adjusted for explainable variance in human judgments). **f. The expression DNN has little bias across demographic groups.** To assess demographic bias, the expression DNN annotations of each face were predicted by averaging the expression DNN annotations across all other faces from the same demographic group. The variance in the expression DNN annotations explained by demographic group in this dataset was quite low, even though no effort was originally made to balance expressions in this dataset across demographic groups. Gender explained 0.88% of the total variance, race explained 0.28%, and age explained 2.4%. Results for individual expressions were generally negligible, although age did explain more than 4% of the variance for three expressions—“contempt”, “disappointment”, and “surprise” (max = 6.2% for “surprise”). Note that these numbers only provide a ceiling for the demographic bias, given that explained variance may also derive from systematic associations between expression and demographics in this naturalistic dataset—for example, because elderly people are less often pictured playing sports, they are less likely to be pictured with expressions that occur during sports. We can conclude that the expression DNN is largely unbiased by race and gender, with age possibly having at most a minor influence on certain annotations.

Extended Data Fig. 6. Expression intensity by number of videos previously rated. Variation in expression was computed as the average standard deviation across expression-based emotion predictions using a linear model trained on all data. Model predictions were first normalized by subtracting the mean across all Japanese responses and dividing by the standard deviation across all Japanese responses. Because Japanese participants responded to many videos, it is critical to note that their expressions did not decline in intensity (operationalized as variation) over time. Variation in expression in the U.S./Canada/Europe is shown for comparison (most English-language survey participants responded to 30 videos or less).

Extended Data Table 1. Translating English-language emotion terms into Japanese.

English	Japanese	Reference(s)
Admiration	称賛	48
Adoration	崇拜	Bulk search of English-Japanese dictionaries via https://ejje.weblio.jp/ . Note that the term “崇拜” is not used in response to “cute” stimuli as is the English “adoration” (“崇拜” is closer to literal “worship”), which explains why ratings were not consistent across cultures (Fig. 2).
Aesthetic appreciation	美的感動	Bulk search of English-Japanese dictionaries via https://ejje.weblio.jp/
Amusement	愉快	49
Anger	怒り	49
Anxiety	不安	50
Awe	畏敬	49
Awkwardness	ばつの悪さ	51
Boredom	退屈	52
Calmness	落ち着き	Based on “落ち着いた気持ち” from ⁴⁹

Confusion	困惑	Bulk search of English-Japanese dictionary via https://ejje.weblio.jp/
Contempt	輕蔑	53
Craving	渴望	Bulk search of English-Japanese dictionary via https://ejje.weblio.jp/
Desire	性欲	49
Disappointment	失望	54
Disgust	嫌惡	50
Empathic Pain	痛みの共感	Bulk search of English-Japanese dictionaries via https://ejje.weblio.jp/
Entrancement	没入感	The most direct translations of "entrancement" – "忘我の境地", "恍惚状態", "有頂天", or "狂喜" – are not easily understood as an emotional state, or have a more pleasant connotation. The term "没入感" was selected because it more directly conveys how a video can capture one's awareness.
Envy	妬み	Selected from ⁵⁴
Excitement	興奮	49
Fear	恐怖	Adapted from "恐れ" from ⁵⁰
Guilt	罪悪感	50

Horror	戦慄	Both “horror” and “fear” are most directly translated to “恐怖”. The term “戦慄” was selected for horror because it is more intense and visceral/physical than “恐怖”.
Interest	興味	49
Joy	喜び	49
Nostalgia	郷愁	49
Pride	自尊心	49
Relief	安心	49
Romance	情愛	Since the most direct translation (“ロマンス”) is not easily understood as an emotional state, chose “情愛” (roughly, “affection”)
Sadness	悲しみ	49
Satisfaction	満足	49
Surprise	驚き	49
Sympathy	同情	49
Triumph	勝利感	Bulk search of English-Japanese dictionaries via https://ejje.weblio.jp/

Subdued (arousal)	気持ちを 穏やかに する	Direct translation, confirmed by bulk search of English-Japanese dictionaries via https://ejje.weblio.jp/ . Term chosen in part based on neutral valence connotation.
Stimulated (arousal)	刺激的な	Direct translation, confirmed by bulk search of English-Japanese dictionaries via https://ejje.weblio.jp/ . Term chosen in part based on neutral valence connotation.
Unpleasant (valence)	不快	Direct translation, confirmed by bulk search of English-Japanese dictionaries via https://ejje.weblio.jp/ .
Pleasant (valence)	快	Direct translation, confirmed by bulk search of English-Japanese dictionaries via https://ejje.weblio.jp/ .

Extended Data References

48. Muto, S. Development of the trait respect-related emotions scale for late adolescence. *Shinrigaku Kenkyu* **86**, 566–576 (2015).
49. Muto, S., Sugawara, D. & Sugie, M. The conceptual structure of positive emotions in Japanese university and graduate students. *Shinrigaku Kenkyu* **89**, 479–489 (2018).
50. Nakai, A. Ikari to ikari no kinjigainen no sousatekiteigi no idou oyobi ikari no sousatekiteigi ni eikyo wo ataeta youin [Differences in operational definitions of anger and the approximate concept of anger, and factors affecting the operational definition of anger. *Meiji Gakuin Univ. Bull. Psychol. Bull. Psychol.* **22**, 13–30 (2012).
51. Higuchi, M., Kuranaga, H., Fukada, H. & Teruya, Y. Mediating mechanisms of embarrassment in non-negative situations. *JAPANESE J. Res. Emot.* (2012). doi:10.4092/jsre.19.90
52. Kusumi, T. & Komeda, H. Kanjo to gengo [Emotion and language]. in *Affective Science* 55–84 (Kyoto University Press, 2007).
53. Morishita, A. & Maiya, K. A Study on Japanese Facial Expressions : Analyzing Facial Expressions of Contempt, Disgust and Anger. *Tech. Rep. IEICE. HCS* **100**, 1–6 (2000).
54. Yamamoto, K. SNOW D18: Japanese Emotional Expression Dictionary [Data set]. <http://www.jnlp.org/SNOW/D18>. (2018).

Supplementary Material

Video S1 [<https://is.gd/8GsKm8>]. **Example reactions, self-report, and predicted emotion for four stimuli.** Reactions on the left side of each stimulus are from the U.S. (stimuli 1, 3, and 4) and U.K. (stimulus 2). Reactions on the right side of each stimulus are from Japan. Self-report ratings have been projected onto the 21 dimensions extracted using PPCA (Fig. 2) for comparison across cultures/languages. Emotions are predicted via models trained and evaluated on data specific to each survey language (English or Japanese). Explore more reactions at <https://tinyurl.com/yywa7kjf>.

Video S2. Illustrating why principal component analysis (PCA) or factor analytic methods are not designed to account for the reliability of judgments of individual features. Each scatterplot represents hypothetical ratings from two sets of raters, with one set of ratings represented in blue and the other in orange. Each dot represents a hypothetical stimulus. Its position represents its average rating in terms of two hypothetical features. In the scatterplots on the left, there is no consistency in ratings across the two sets of raters (blue and orange). In the middle, the consistency in ratings can be captured by one dimension, with ratings orthogonal to this dimension being inconsistent across the two sets of raters. On the right, both dimensions are required to explain the consistency in ratings across the two sets of raters, given that the dots are stable in position along both axes. The dimensions extracted by PCA and factor analytic methods (first row) account only for the distribution of ratings – concatenated or averaged across datasets – and extract the same dimensions regardless of the reliability of judgments of each feature. Thus, PCA is not equipped to identify when an individual category, like fear, is reliably distinguished from every other rated category. By contrast, dimensions extracted by PPCA (second row) are designed to account for the reliability of ratings of each individual feature.