Using evolutionary algorithms to uncover individual differences in how humans represent facial emotion

Christina O. Carlisi, PhD^{1*}, Kyle Reed, MSc², Fleur G.L. Helmink, MSc³, Robert Lachlan, PhD⁴, Darren P. Cosker, PhD², Essi Viding, PhD^{1†}, Isabelle Mareschal, PhD^{5†}

¹Division of Psychology and Language Sciences, Developmental Risk and Resilience Unit, University College London, 26 Bedford Way, London, WC1H 0AP, UK

²Department of Computer Science, University of Bath, 1 West, Claverton Down, Bath, BA2 7AY, UK

³Erasmus University Medical Center, 's-Gravendijkwal 230, 3015 CE, Rotterdam, Netherlands

⁴Department of Psychology, Royal Holloway University of London, Wolfson Building, Egham, TW20 0EX, UK

⁵School of Biological and Chemical Sciences, Department of Psychology, Queen Mary University of London, G. E. Fogg Building, Mile End Rd, London, E1 4DQ, UK

[†]Authors contributed equally to this manuscript

****Corresponding author:***

Christina Carlisi, PhD
University College London
Psychology and Language Sciences
26 Bedford Way
London, WC1H 0AP
+44(0)2027 6795367
carlisi.christina@gmail.com

Keywords: Emotion processing, facial emotion, social cognition, individual differences, social psychology

Abstract

How humans perceive and represent emotional facial expressions critically guides social interactions and has a profound impact on social cognition. However, much emotion research to date is potentially flawed by relying on the assumption that people represent categorical emotions in the same way, using standardised stimulus sets and therefore overlooking important individual differences in emotion processing. To resolve this problem, we developed a task using *Genetic Algorithms* and derived participant-generated emotional expressions from 105 individuals. A separate group (N=108) was then asked to identify the expression on these faces. Taken together, results showed relative population-level consistency in the representation of happy faces alongside a high degree of variability in the representation of fear and sadness. High test-retest reliability was observed. This novel method enables the efficient quantification of variation in how people represent emotional expressions, presenting promising advances for the study of individual differences in social cognition and emotion processing.

1 Introduction

2 Facial expressions are ubiquitous in our everyday environments, conveying information
3 about our internal emotional states, thoughts and intentions. It has been debated whether basic
4 emotional expressions represented by faces are universally recognized(1–4), but what is certain is
5 that the ways in which humans represent and interpret such expressions shape relationships and
6 guide social interactions(5, 6). However, much of the psychological research investigating emotion
7 representation has relied on well-validated yet over-generalised stimulus sets of emotional
8 expressions, potentially limiting a complete understanding of the nuanced features underpinning
9 emotion recognition at the individual level. Being able to precisely quantify such features would
10 address this challenge and represent a step change in how we conduct and interpret emotion
11 research.

12 Identifying and interpreting emotions from facial expressions involves processing
13 information that is communicated by morphological facial changes, e.g. opening of the eyes or
14 mouth, frowning, smiling, etc.(7). These morphological facial changes are thought to convey the
15 expresser’s internal feelings or sentiment(8). Emotion recognition is thus a calculated process by
16 which perceptual and consequent affective information must be interpreted using acquired
17 semantic knowledge and contextual information that can vary widely among individuals. There are
18 profound individual differences in the interpretation of emotional facial expressions(9, 10), but it is
19 unclear what underlies these differences. They could reflect a lack of sensitivity to emotional
20 expressions (i.e. a face looks sad to an individual, but their response to it is not typical for a sad
21 face), a different internalised representation of the emotional expression (i.e. what an individual
22 thinks “sad” should look like, in their mind’s eye, when represented on a face), or individual
23 differences in low-level visual processing which might compromise the perception of certain
24 emotional expressions more than others. Distinguishing between these different sources of emotion
25 processing and how they vary among individuals across a population is critical; at the extreme end,
26 such individual differences can result in failures to appropriately interpret emotional expressions

and can significantly increase someone's risk of developing pathological behaviour such as anxiety(11) and psychopathy(12).

Most common methods for testing emotion recognition rely on the assumption that emotions are discrete categories, which requires agreement across observers. However, evidence suggests that emotional facial expressions may instead be represented dimensionally(8), varying subtly from person to person. This has been shown in studies where basic emotion categories such as sadness and fear are commonly confused with one another by different individuals during identification or recognition tasks(13). One possible source of this variation among individuals is the fact that most experiments investigating differences in the representation and identification of facial emotional expressions are limited by their use of established pre-determined sets of facial stimuli designed to represent basic emotion categories (e.g.(14, 15)) when in fact, these categories may not be represented in the same way by all individuals. Additionally, it has been argued that many studies of emotion processing which use a (two alternative) forced-choice response to identify a given facial stimulus artificially constrain the detection of individual differences because responses are constrained to singular categories(16). Finding new ways to generate accurate, individualised representations of emotional facial expressions would represent a step towards solving some of these issues in the way we capture individual differences in emotion processing, thus allowing a more comprehensive investigation of individual differences which are .

Facial expressions and refinement of their representation has been a topic of great interest in the computer graphics and animation community for some time(17). Williams first introduced the method of performance-driven animation in 1990(18), and there have since been numerous implementations of this method to create realistic, performance (i.e. participant) driven graphics(19, 20). However, traditional performance capture techniques rely heavily on the judgement of a skilled artist. There is a lack of participant-driven animation approaches that unskilled participants can interact with to shape the graphic output. Therefore, user-driven graphic evolution heuristics such as

the Genetic Algorithm (GA), thus termed because of its procedural similarities to evolutionary processes and natural selection, provide a unique solution to the longstanding problem of how to optimally yet efficiently capture an individual's internal representation of a facial expression. GAs efficiently explore a parameter space that exceeds the limits tested using standard stimulus sets and thus strike a balance between two conflicting requirements: on the one hand the need to exhaustively search the stimulus space to examine all possible combinations of attributes required for a specific stimulus (e.g. nuanced facial features comprising an emotional expression); and on the other hand, the need to efficiently test the space of available attributes to rapidly converge on the optimal solution (i.e. an ideal representation of an emotional expression). As such, GAs are particularly well-suited to multi-dimensional problems such as searching 'face-space' where a number of different features (e.g. eye shape, mouth curvature, etc.) contribute to a particular facial expression(21, 22). GAs have already been used in ecology research and forensic settings, where they enable eyewitnesses to evolve the facial identity of a suspect and generate e-fits with greater accuracy compared to traditional sketches(23, 24). However, GAs have not been widely applied to facial emotion research because of technological difficulties around creating realistic expressions. Prior studies using more traditional methods (i.e. not using performance-driven algorithms) have predominantly focused on static, pre-determined sets of emotional face stimuli or dynamic stimuli recorded from actors, but which do not quantify the participant's own perceptual representation of an emotional expression. While some studies have used data-driven approaches with non-static stimuli (e.g.(6, 20, 25–27)), these methods are time consuming, usually requiring a very large number of trials and therefore ill-suited for large scale testing (or testing child or clinical populations).

To address this methodological gap, the present study developed advanced computer graphics methods to allow participants to quickly and efficiently create their own facial expressions, without any prior (experimenter-driven) constraints, but rather based on their internal representation of what a given emotion looks like. Our approach combined computer animation

with visual psychophysics to develop a novel technique inspired by genetic algorithms exploring participants' internal perception of emotional faces. We applied this new method to quantify a visual approximation of individual differences not typically captured by traditional facial emotion paradigms. This method employs a global optimisation framework using participants' choices as a heuristic to drive stimulus creation and exploits humans' ability to process minute differences in facial expression. Moreover, we tested the reliability of this measure as a valid probe of these individual differences, examining the stability of within-subject responses across three time points 30-45 days apart. Ultimately, our method allows us to precisely examine and quantify individual differences that exist in emotion recognition at the level of nuanced perceptual changes in the signaler's facial features. This approach has not previously been possible and opens up new possibilities for the study of human emotion perception, social cognition and affective processing.

Methods

1. Paradigm specifications

1.1 Facial expressions and the use of blendshapes

We sought to obtain individualized facial expressions corresponding to four basic emotions: happy, angry, sad, and fear, in accordance with each participant's own internal representation of these emotions. Expressions were generated from a face rig(28) (<https://www.camera.ac.uk/>) in Autodesk Maya 2018 (<https://www.autodesk.co.uk/>) by combining weights from 42 preselected *blendshapes* within a rig which can be conceptualised as facial features (e.g. lip curl, jaw openness, eye squint) loosely governed by the Facial Action Coding System (FACS)(7) (Figure 1A).

Blendshapes are weighted vector offsets in a mesh (Figure 1A) which control movement of a set of facial features from a neutral source $face_0$ to a target expression $face_k$, that can be combined linearly to cover an extensive range of emotions using relatively few parameters(28); variation in any

one blendshape can affect the facial morph of the overall expression (Figure 1B). This is controlled via changes in blendshape weights w , typically between 0 and 1, such that any facial expression can be derived as follows, for k blendshapes:

$$Expression = face_0 + \sum_{k=1}^n w_k (face_k - face_0) \quad (1)$$

1.2 Genetic Algorithm

Genetic, or evolutionary algorithms adapt a reinforcement learning approach, relying on sampling and evaluation of parameter combinations to assess the fitness of outcomes. In the case of holistic e-fit systems and the paradigm described in the current study, this scheme can be adapted to be “user guided”, that is the sampling and evaluation of solutions is done by the participant. This process involves asking participants to select a subset of facial expressions (typically 3 or 4 from a set of 10-12) that they think most closely represents a specified emotion, with the process repeating over a number of iterations. With each iteration, a new set of facial expressions is created, based on the previously selected subset (with noise added, or random mutations). To generate the stimuli (facial expressions) for each iteration, the optimisation framework of a GA is implemented. GAs are thus termed because analogous to natural selection, expressions are evolved over several *generations of sampling, selection, breeding and mutation* until the *gene pool* converges or evolves toward the participant’s ideal expression.

To demonstrate the functionality of the GA within this paradigm using the analogy of evolutionary processes, we begin by defining these key conceptual parameters of the GA. *Genes* represent what is being evolved, in this case physical observations of facial expressions or combinations of blendshapes. A string of blendshape weights (i.e. a blendshape ‘rig’) encodes face shape to generate an expression, much as genes encode physical phenotypes. Here, the given facial expression is represented as a collection of these blendshape weights w for K blendshapes:

$$Expression = \{weight_1, weight_2, \dots, weight_k\} \quad (2)$$

1.3 Initiation & Evolution

To initialize the GA, two sets are defined. The current population or search space (i.e. *GenePool*) consists of all possible parameter combinations that sampling can select from. A *Sample* face is derived by sampling k weights from the existing gene pool at time t . Thus, the sets exhibit the following relationship:

$$weight_k^t \in Expression_k \in Expression \in Sample \subset GenePool$$

Each blendshape for a given face in the population is uniformly sampled between 0 and 1 from the model. In this paradigm, this produces 10 facial expressions per generation (Figure 1B).

$$Sample = \{weight_k | weight_k \sim U([0,1]), 1 \leq k \leq K\} \quad (3)$$

Generating faces can be generalized as sampling gene configurations from the updated population/gene pool, where a stochastic selection scheme is implemented.

$$Sample = \{weight_k | weight_k \sim Selection(Expression_k), 1 \leq k \leq K\} \quad (4)$$

Participants' selections drive the gene selection scheme, as participants choose sample faces with the most desired characteristics for their search. Participants are presented with 10 face samples and select 3 per iteration with the most favourable phenotype (Figure 1C), denoted as *Sample**. The genes of these faces are then added to an updated gene pool, where t denotes time point, i.e. generation number:

$$GenePool^{(t+1)} = \{Sample * _1, Sample * _2, Sample * _3\} \quad (5)$$

Participants are also asked to choose their *elite* face (Figure 1D), which represents the best sample they selected from the current generation (i.e. the face with the expression closest to their representation of that expression). This elite face has a much higher influence on subsequent sampling compared to the other two selected sample faces. This elite face remains intact and is

148 directly propagated through to the next generation, ensuring that the search cannot move
149 backwards. This is placed at the start of the current gene pool such that $GenePool_1 = Elite \in Sample^*$.

150

Figure 1. Schematic overview of task stimuli and procedure (A) Neutral depiction of the rendered face rig used in the Genetic Algorithm (GA) task with an overlaid mesh of vertex points. Offset in any one vertex results in distortion at that point. (B) Schematic overlay of blendshape features as they correspond to rig controllers. Change in any one controller translates to a corresponding change in the associated facial feature based on vertex offset in image (A). (C) Step one of the GA procedure. The participant is presented with 10 pseudo-randomly generated faces and a target emotion (e.g. 'fear') to identify. They are asked to reflect on how they subjectively represent the target emotion. (D) The participant relies on this representation to guide selection of 3 faces from the population which are most aligned with the target emotion, including an "elite" face (in yellow). (E) This process is repeated and the population is evolved. (F) After a fixed number of generations, a final "prototype" facial expression is generated that approximates that participant's internalized representation of the target emotion (in green). This is the expression that is then quantified and analysed.

159

1.4 Evolution

Once the participant has chosen three samples and selected one of those three as an elite face, they proceed to the next generation by updating the gene pool such that $G^{(t)} \leftarrow G^{(t+1)}$. After the first generation, a new sampling scheme is introduced for sampling new genes, which generates new sample faces through *breeding*, *mutation*, or both (Figure 1E).

Breeding of genes involves parameter crossover of selected faces, analogous to chromosome crossover during genetic recombination. Two parent samples P are selected at random, where each parent gene has a 50:50 chance of being chosen for the child (Equation 7). The elite face is encoded to have a much higher likelihood of being selected as a parent.

$$Breeding(P_1, P_2) = \begin{cases} gene \in P_1 \mapsto gene \in Sample & \text{if } x \sim U([0,1]) < 0.5 \\ gene \in P_2 \mapsto gene \in Sample & \text{otherwise} \end{cases} \quad (6)$$

There is also a chance of *mutation* within the genes, such that there is a small probability γ that a gene can change to a value not observed in either parent. In the context of this paradigm, mutations allow the new generation to explore portions of ‘face space’ not occupied by previous generations. This, combined with crossbreeding, offers an optimal degree of variability in the offspring that allows to fully explore the search space.

$$Mutation(Sample, \gamma) = \begin{cases} weight_k \in Sample \mapsto U([0,1]) & \text{if } x \sim U([0,1]) < \gamma \\ weight_k \in Sample \mapsto weight_k & \text{otherwise} \end{cases} \quad (7)$$

After several iterations, sample faces from the gene pool should start to converge locally toward the participant’s optimal expression. In the current study, the process stopped after 6 iterations, striking a balance between an adequate number of generations (e.g. as shown to optimise facial generation(29)) but minimising the participant ‘burn-out’ (i.e. boredom) since the faces become more and more similar with each iteration. As with most optimisation tasks, genes (blendshapes) are evaluated as the error between the evolved face and the target face. Selection (S) occurs where the participant prefers genes which attribute the lowest error. In this paradigm, target

expressions are determined by each participant's desired characteristics driven by their own internal representations of emotions.

1.5 Optimising evolution

Critical for the optimisation framework of the GA is the balance between local and global optima. Local optima are important for gradually improving parameter estimation, whereas global optima are important for obtaining the best result. To achieve the optimal balance between the two, a breeding/mutation scheme is employed, where half of the generations (i.e. 5 of 10 faces) stay faithful to the participant selected faces (local), and half offer much greater variability (global). Thus, for a select number of faces, new genes (blendshapes) are resampled, as in Eq. 6, to expand the search space. The resulting 10 new faces of the next generation are then ordered randomly so that face location is not a dependent factor in expression selection.

1.6 Coarse-to-fine search

For the current paradigm, we used 10 faces. This number was chosen to avoid information overload for participants on each trial and increase efficiency, while keeping the faces easy to view on a single screen. To compensate for this relatively small population, we employ a *coarse-to-fine* approach to more extensively cover the search space. This process involves choosing blendshapes with the greatest variability for random sampling and mutation in initial generations, with more nuanced changes in expressions chosen in later generations. As a crude form of segmentation, only the lower half of the face is sampled in the second generation, and only the upper half in the third generation. Subsequent generations sample blendshapes from the entire face.

1.7 Evolving facial expressions

To evolve facial emotional expressions, we sampled from a rig of 42 blendshapes (e.g. jaw open, brow squeeze)(28) (Figure 1B). Teeth and eyes are added separately to the faces and are static throughout the evolution process. This rig was adapted from a larger rig where blendshape

correctives were removed and left/right blendshapes were combined into a single dual-sided action. This is necessary because without such constraints, correctives and asymmetry increase the likelihood of implausible, non-realistic expressions being created. The system starts by creating a dictionary of blendshape weight values in the initial population. Rather than sampling random weights, the order of blendshapes to be sampled for the initial population is pre-selected and a coarse-to-fine search is implemented (see section 1.6 above); the most diverse shapes (e.g. frown, sneer, jaw open, smile) are presented for the first generation. The most prominent shapes for the lower half of the face are then sampled for the second generation, and the most prominent shapes for the upper half for the third. Any shapes can then be sampled for subsequent generations. Obtaining a suitable combination also necessitates incorporating an appropriate mutation rate. To achieve this, a random value was assigned between 0.25 and 0.75. Upon completion of the iterations, the final elite (“prototype”) facial expression from the last generation was parameterised (via blendshape weights associated and geometric error of vertex offsets). This was taken as a quantitative descriptor of the participant’s internal representation for the given emotion.

2. Experimental study: participants

To validate our GA paradigm for the efficient generation of participant-derived prototype emotional facial expressions, we tested 105 adults between 18 and 61 years of age (29% males, mean age 26.58 ± 8 years, 51.0% Caucasian, 42.3% Asian, 1.9% Black, 4.8% Other) recruited from the local community. A subset of 60 individuals repeated the GA task on two more occasions, once approximately 30 minutes after the first instance, and again 30-45 days later. All participants gave written informed consent, and this study was approved by the UCL Research Ethics Committee (Z6364106/2013/07/37), and data were processed in accordance with the European Union General Data Protection Regulations.

3. Behavioural task

For each generation, participants were shown 10 faces each and instructed to select 3 out of these 10 faces that best represented the emotion (happy, angry, sad, or fearful) displayed in a response window. After making their selections, participants were instructed to select one of their chosen faces which they thought best represented the target emotion (the 'elite' face). This process was repeated six times. On a given trial, participants were free to amend their selections and were not constrained by explicit time limits (though they were instructed to make their selections as quickly yet accurately as possible). Once the participant was satisfied with their choices, they selected to evolve the generation, after which point their selections could not be changed for that trial. On the subsequent generation, participants were presented with a set of 10 new faces evolved from the previous generation and they repeated the selection and evolution process. After 6 generations, blendshape weights for the final elite face were saved for further analysis. This face was termed the 'prototype face' for that participant, i.e. the face which most closely represents the individual's internal representation of the target emotion. Four emotional categories (happy, angry, sad, fear) were tested in separate blocks, in a randomised order across participants.

The GA task was run on a 33.8 cm Dell laptop (1366 x 768, 60Hz) at University College London using the Autodesk Maya interface. Participants were seated approximately 57 cms from the screen and at this viewing distance, 1 pixel subtended 1.3 arcmin. Faces were displayed in grayscale and each face subtended 1.9 x 3.0 degrees.

4. Statistical analysis

4.1 Emotion specificity and response variability

Statistical analyses were performed in Matlab (2018b, The Mathworks, Inc., Natick, MA, United States) and R version 3.5.0(30). Multivariate dispersion(31) was used to assess variability of individual prototype faces from the population norm within each of the four emotion categories. In this process, the spatial median (i.e. population norm) for each category is computed within 42-dimensional 'face-space'(22). The dissimilarity between each data point and its category's spatial

median is calculated, and Levene's F-statistic is obtained for comparisons between categories. The significance of these differences is calculated by permuting units between groups and recalculating the Levene's F-statistic.

We used t-distributed stochastic neighbour embedding (t-SNE) to visually inspect and evaluate the 2-dimensional clustering of generated faces based on their respective categories (happy, angry, sad, fear). This method is a nonlinear dimensionality reduction technique well-suited for the visualisation of high-dimensional data(32). In this case, each face consisted of a 42-dimensional face-space to comprise a single emotional expression. T-SNE converts this high-dimensional dataset into two-dimensional data that can be visualised in a 2D scatterplot while preserving as much of the significant structure of the high-dimensional data as possible in the low-dimensional map. Effectively, this method converts high-dimensional Euclidean distances between data points into conditional probabilities which represent similarities. The t-SNE algorithm consists of two stages modelling pairwise similarities, using the Euclidian distance between points as the basis of the similarity metric.

To assess the blendshapes which were most important in forming prototype expressions for each emotion, we calculated a scaled metric of influence based on the extent to which each Blendshape affects the overall mesh. This was done by scaling the blendshape weights by the geometric error (sum of squared errors from the expression minus the neutral expression vertex positions) based on the Euclidean distances of the vertex offsets at each vertex within the mesh. This allowed the ranking of blendshape influence by their scaled degree of change in the mesh.

4.2 Test-retest reliability and within-subject response variability across time

To assess whether the GA task provides a reliable measure of individual differences, we examined whether between-subject individual differences in prototype faces were greater or smaller than within-subject variance across time. This analysis also allowed us to confirm that after six generations, the prototype face was not generated as a result of either a specific mutation or

local minimum in the algorithm's iterative process. We first calculated a dissimilarity matrix for each emotion from the Euclidean distances in 42-dimensional space between all the faces produced at the three time points. We then coded an individual identity matrix for each emotion, where each row and column were different faces, and within-subject comparisons were assigned a value of 0 and between-subject comparisons were assigned a value of 1. Mantel tests were used to test for correlations between the individual identity matrices and the dissimilarity matrices using the 'vegan' package in R(33) with 9999 permutations and using a Spearman test statistic.

To further evaluate test-retest reliability of the GA task, resulting distance values from multivariate dispersion analysis within each category for each time point were used to compute two-way mixed effects intraclass correlation coefficients (ICCs).

5. External validity

To assess external validity of the prototype faces generated from the above procedure, 26 faces (25% of the total number of prototype faces generated from time point 1 of the study) were randomly selected for each emotion from the GA task, resulting in a total of 104 stimuli. We used Amazon's Mechanical Turk (<https://www.mturk.com/>) to recruit 108 individuals (66% male, 66.7% white, 30% Asian, 7.4% black or mixed) between the ages of 18-65 (average age 34.5 years). The faces were presented using an online platform designed for behavioural experiments (www.gorilla.sc), and participants were asked to describe, using one word in a free-text response box, the emotion that they perceived the face to be displaying.

Responses were coded independently and cross-validated by CC and FH as being either a) an exact match (i.e. "happy/happiness", "angry/anger", "sad/sadness", "fear/fearful"), b) correct (i.e. synonymous with "happy/angry/sad/fearful" as defined by Merriam-Webster online thesaurus (www.merriam-webster.com/thesaurus) c) incorrect (not a synonymous match for the target emotion but still a plausible description of a different emotion) or d) anomalous (indecipherable, e.g.

“d” or “2”; not consistent with a standard word used to describe basic facial emotion, e.g. “flirty”, “lying”; or ambiguous, e.g. “intense”, “indifferent”).

Results

1. Emotion specificity and response variability

Multivariate dispersion

Distributions of Euclidean distances of participants’ prototype faces (represented in 42-dimensional face-space) from the spatial median (conceptualized as the population ‘norm’ face) are displayed for each emotion in Figure 2A. Population-level deviation from the norm was lowest for happy, with the highest density of distances closest to 0, and greatest for fear, with the highest density of distances furthest from 0. This was confirmed by comparisons of distances from the spatial median between each emotion: participants’ prototype happy faces deviated less from the population norm (spatial median) compared to sad [$F(1,358)=98.79$, $p<.0001$, adjusted $R^2=.21$], angry [$F(1,358)=134.10$, $p<.0001$, adjusted $R^2=.27$], and fearful faces [$F(1,358)=14.67$, $p<.001$, adjusted $R^2=.04$]. Fearful faces deviated more from the spatial median compared to sad [$F(1,358)=29.12$, $p<.0001$, adjusted $R^2=.07$], and angry faces [$F(1,358)=48.05$, $p<.001$, adjusted $R^2=.12$]. There was no difference in the degree of deviation from the spatial median for prototype angry faces compared to sad [$F(1,358)=2.75$, $p=0.10$, adjusted $R^2=.005$]. All significant findings survive Bonferroni correction for multiple comparisons at an adjusted threshold of $p=0.05/6=0.0083$.

We also used distances obtained from multivariate dispersion analysis to observe, within each emotion category, the number of instances where the distance of an individual data point was closer to the spatial median of a different category than it was to the spatial median of its own category (Figure 2B). These results were interpreted as an index of ‘confusability’ among the prototype faces generated for each emotion. The more instances where the distance of a prototype face was closer to the spatial median of a category other than its own, the higher degree of

‘confusability’ with other emotions among prototype faces from that category. That is, the higher number of instances where the distance was closest to the median of the same category, the higher the level of agreement among prototype faces within the population for that category. Happy was the most consistent category, with 103 of 105 data points closest to the spatial median of happy; there were only two data points for which the spatial median was closer to that of a different category (1 closer to sad, 1 closer to fear, 0 closer to angry). Angry was the next most consistent category, with 82 of 105 distances closest to the spatial median for angry (9 closer to sad, 9 closer to fear, 5 closer to happy). Sad and fear showed higher levels confusability with one another, with roughly half of the data points for one of these emotions being spatially closer to the median of the other (Sad: 41 closer to fear, 5 closer to happy, 2 closer to angry; Fear: 53 closer to sad, 2 closer to happy, 2 closer to angry).

t-SNE

Results of the t-SNE data reduction and visualisation are presented in Figure 2C. Reducing the 42-dimensional face-space data into 2-dimensional space, t-SNE represented participants’ responses as a single metric clustered by the four emotion categories. Among these clusters, happy and angry were the most separable and least variable categories, forming distinct clusters with little overlap among the other emotion categories. Sad and fear exhibited greater overlap with one another compared to happy or angry. Moreover, t-SNE visualisation showed that these two categories appeared less tightly-clustered (more diffuse) in two-dimensional space relative to happy and angry, suggesting there was greater variability in these faces among participants.

Blendshape influence

The blendshapes that were most important in forming happy, sad and fearful expressions all involved mouth-related features, primarily those determining the respective curvature of the mouth (Figure 2D). Happy was most strongly linked to changes in *smile controller* (mean weight 0.29 ± 0.26) and *smile open* (mean weight 0.39 ± 0.36), although it is important to note that the mean weights for

354 these blendshapes were less than 0.5, indicating that on average, a threshold of less than 50% of the
 355 full expression of these features was required to form the prototype happy expression. Sad was
 356 associated with *lip corner depressor* (mean weight 0.44 ± 0.37). Fear was also linked to *lip corner*
 357 *depressor* (mean weight 0.30 ± 0.30) as well as *lip down* (mean weight 0.47 ± 0.32). Only angry faces
 358 prioritised change in a feature other than the mouth (*nose wrinkler* (mean weight 0.83 ± 0.24)), with a
 359 comparatively higher weight for this feature required to generate a prototype fearful expression.

360

362 **Figure 2. (A)** Histogram plots showing the population distribution of Euclidian distances from the spatial
 363 median for each emotion category. A smaller distance represents less variation from the population norm. **(B)**
 364 For prototype faces generated within each emotion category, percent of the total number of responses

(N=105) where the distance to the spatial median was closest to each respective emotion's spatial median. **(C)** Results of the t-stochastic neighbour embedding (t-SNE) data reduction and cluster visualisation (N=105). 2D spatial location of the clusters is arbitrary except relative to one another. **(D)** For illustrative purposes, neutral rig examples displaying the fullest expression of only the blendshape which was deemed most important in forming happy, angry, sad and fearful expressions. Blendshape weights for these features were set to 1, whilst all other blendshape weights in the rig were kept at 0 (neutral).

2. Intra-individual reliability

Mantel tests (Figure 3A) showed that faces generated for all four emotions were positively correlated with the individual-identity matrix [happy: $\rho = .048$, $p < .001$; sad: $\rho = .018$, $p < .001$; angry: $\rho = .010$, $p < .01$; fear: $\rho = .016$, $p < .001$], indicating that variation between subjects was greater than within-subject variability across the three time points. Moreover, there was no significant difference of correlations in each participant between time point 1 and time point 2, time point 2 and time point 3, or time point 1 and time point 3 for any emotion, indicating that the amount of time between visits (10 minutes vs. 4-6 weeks) did not significantly impact within-subject variability.

ICCs calculated across the three time points using distances from the spatial median indicated moderate to good response reliability for all four emotion categories: happy: ICC = .585, 95% CI = .363-.739, $F(59, 118) = 2.410$, $p < .001$; sad: ICC = .576, 95% CI = .350-.733, $F(59, 118) = 2.361$, $p < .001$; angry: ICC = .419, 95% CI = .107-.634, $F(59, 118) = 1.720$, $p < .01$; fear: ICC = .574, 95% CI = .345-.731, $F(59, 118) = 2.345$, $p < .001$ (Figure 3B).

Figure 3. (A) Violin plots illustrating results of Mantel tests within each emotion category, comparing dissimilarity matrices of within- versus between-subject responses. Within each category, left plots represent within-subject differences across the three time points, and right plots represent between-subject differences. Boxplots represent the median and interquartile range. *** indicates significant at $p < 0.0001$, ** indicates significant at $p < 0.001$. Plots indicate that between-subject variability was greater than within-subject variability across time for all four emotions. **(B)** Intraclass correlation coefficients for distances from the spatial median measured across the three time points. *** indicates significant at $p < 0.001$. Error bars represent standard error. Red=poor reliability, orange=fair reliability, green=good reliability, blue=excellent reliability(34).

3. External Validity

Representations of the most commonly provided responses to prototype faces within each emotion category are displayed in Figure 4A. The two most common responses for happy, angry, and sad were exact matches for the emotion displayed. The ten most common responses for happy faces were all synonymous with happy, suggesting high population-level agreement for happy prototype faces. However, the most common response when participants were shown a fearful face was “sad”, and “disgust” was among the ten most common responses for angry, sad and fearful faces, indicating a higher degree of ‘confusability’ among these expressions, particularly fear.

Of the 11,686 responses collected from external validators across all four emotion categories, 7,752 (66.3%) were classed as “correct”, of which 5,375 (69.3% of correct responses, 46%

407 of total responses) were deemed an exact match. 3,934 responses (33.7%) were classed as incorrect,
408 of which 116 responses (2.9% of incorrect responses, 1% of total responses) were classed as
409 anomalous. Accuracy rates per emotion are displayed in Figure 4B. Happy faces had the highest
410 accuracy rate, where 91.8% of happy faces were correctly identified (68.8% of which were an exact
411 match). 81.3% of angry faces were correctly identified (75.2% of which were an exact match) and
412 60.8% of sad faces (80.4% of which were an exact match). Fearful faces had the lowest accuracy
413 rate, with only 31.0% of fearful faces correctly identified (35.6% of which were an exact match). Of
414 total incorrect responses for each emotion category, 12.3% were deemed anomalous for happy
415 faces, 4.5% for angry faces, 2.5% for sad faces, and 1.7% for fearful faces.

Figure 4. (A) Word cloud and histogram representations of the most common responses among online external validator participants within each emotion category when asked to respond to randomly selected prototype faces created from the GA task using one word in a free-text response box. Histograms depict the ten most commonly provided responses, ranked from most to least common. **(B)** Bar chart depicting proportion of correct (Exact and Synonymous) and incorrect (Incorrect and Anomalous) responses made by online external validator participants.

Discussion

Combining methods from computer animation and psychophysics allowed us to uncover individual differences in the way in which people perceive and internally represent facial emotional expressions, addressing a gap in the field of emotion research(10). This study shows that GAs are an efficient method to derive a quantifiable measure of an individual's internal representation of facial emotional expressions. We first described a paradigm which implements a GA to evolve a set of facial emotional stimuli and showed that it can efficiently generate reliable individualised prototype facial expressions for four distinct emotional categories (happy, sad, angry, and fear), that are consistent across time. We next showed that an independent sample of individuals was able to identify emotions displayed by the prototype expressions generated from the GA, but that this ability was more accurate and less variable for positive (i.e. happy) relative to negative emotions (particularly fear). These findings suggest that it is possible to generate personalised stimuli that can be a) quantified to index individual differences in emotion processing and b) used in conventional, well-established facial emotion processing tasks as a way to further improve our understanding of individual differences in emotion perception. Moreover, this study provides preliminary evidence that many of the commonly used facial emotion processing paradigms using standardised sad and fearful stimuli may not be adequately capturing individual differences in how people perceive and represent these emotions.

The reliability and generalisability of many traditional paradigms used to measure emotion processing has been called into question(10, 35, 36). Recent advances in data-driven approaches to the study of facial emotion(6, 20, 25–27) have started to overcome these limitations and work towards individualised measures of emotion processing. However, these methods are often limited by their requirement of many trials, making them unsuitable for use in large or specialist study groups such as children or clinical populations. Our study represents progress in the field of emotion research by showing that GAs are a reliable and efficient alternative which can generate

individualised, quantifiable representations of facial emotion in a matter of minutes. We not only showed that there was less variability within individuals across time points compared to between participants (i.e. participants look more like themselves over time than they do each other), but we also found a high degree of reliability of task measures. This confirms the ability of this task to distinguish within versus between-subject performance, providing evidence that the measure is sensitive enough to detect nuanced individual differences. Moreover, the external validation study showed that participants were able to correctly identify two-thirds of overall responses, in line with previous work showing average accuracy rates for emotion recognition across seven categories to be between 64-67%(37).

Clustering results showed that prototype angry and happy faces were visually distinct from the other emotion categories, and these two emotions had the highest identification accuracy rates in external validation. This is perhaps unsurprising for anger, as it has been hypothesized that humans are evolutionarily “hard-wired” to detect angry or threatening stimuli(38). Previous work has shown, for example, that the detection of angry faces is faster and more efficient compared to other emotions(39). Similarly, results for happy support the hypothesized “*happy face superiority effect*”, which suggests that happiness is recognised faster and more accurately compared to other emotions as an evolutionary mechanism to promote group cohesion and social bonds(37, 40, 41). Lastly, the fact that angry and happy prototype faces required expression of blendshapes in facial movements distinct from those needed for fear and sad expressions may help clarify the distinct mechanisms driving differences in representations across emotions.

On the other hand, there was a significant degree of overlap for the representation and perception of fearful and sad prototype faces. This suggests that individuals’ internal representations of fear and sadness had more features in common (and that they were more commonly confused with one another) compared to happy or angry expressions. Moreover, external validation in a separate sample of participants showed that these two emotions were least consistently correctly

identified, with prototype fearful expressions most often perceived as sad. This finding is in line with and can help explain findings of poorer classification accuracy for expressions of fear and sadness relative to happiness and anger, and that fearful expressions are frequently confused with sad expressions(13). Additionally, people tend to respond with lower accuracy when asked to identify fearful expressions compared to other basic emotions(42). Importantly, our study extends this evidence to show that this ‘confusability’ effect is also present when perceiving personalised, rather than standardised, representations of emotion.

For many years, studies have used fearful faces as stimuli to examine the association between the processing of acute fear (and associated physiological and neurocognitive responses) the presence of behavioural and psychiatric traits(43–45). These studies have assumed that the recognition and perception of pre-determined standardised stimuli is comparably represented by a diverse group of individuals. Concurrently, there have been calls to refine the measurement of emotion perception which we know vary at the individual level(10, 46, 47). The results from the present study suggest that in paradigms where participants must recognize fear relative to other emotions, the sensitivity of the fearful stimuli should be carefully considered, as there may be a large degree of population-level variability specific to fear in the way that individuals represent and respond to fearful faces. A key benefit of our new approach is the ability to directly quantify this variability and investigate its impact on wide-ranging aspects of emotion processing, from how we perceive the expressions of others to how we express our own emotions. For example, it has been argued that there is a social need to express certain emotions (e.g. happiness) more explicitly to foster cohesion among individuals, while other emotions (e.g. fear, sadness) may be masked or camouflaged by the expresser who may not want to index their vulnerability or distress(48). This innate social pressure may in part contribute to the “confusability” observed in this study between fear and sadness, along with the fact that it is unlikely one will experience concurrent feelings of happiness and e.g. fear, whereas it is possible that one might concurrently experience fear and sadness. An index of the variance of expressions generated by different participants within emotion

categories could also prove particularly useful in psychopathology research where a) less variation within a category could make it easier to identify departures from the population norm, or b) greater variation may indicate that we need to be particularly mindful of how we sensitively and reliably assess individual differences in perceiving negative emotions such as fear. Either of these cases would indicate that improvements in how we measure variation could be of great value when seeking to understand mechanistic underpinnings of problematic traits and behaviours (e.g. anxiety, aggression).

Limitations and future directions

To our knowledge, this is the first application of GAs to the study of individual differences in facial emotion representation and expression generation. However, this initial study is not without limitations and has paved the way for refinement of this approach in future studies. The faces used in the current interface are computer-generated rather than photo-realistic. It is possible that while these computer-generated faces approximate portrayals of real human faces, their representation still differs from “real-world” stimuli(49). Future iterations of this task will leverage higher fidelity photogrammetric and photographic facial data together with recent advances in neural rendering(50). It has also been suggested that morphologically similar expressions such as sadness and fear may differ primarily in their dynamics (i.e. facial movements)(51). Comparing dynamic stimuli generated by our GA could prove a fruitful future step in isolating the source of individual differences. The high test-retest reliability within subjects across time provides evidence that the GA used an appropriate number of iterations for adequate convergence, and mitigates the likelihood that the GA fell into a local minimum. However, future studies should examine the test-retest reliability of this task using a different initial population of test faces at each time point. This will further confirm that if the participant still creates quantifiably similar faces at each test session, the number of iterations was appropriate and the local minima problem was not an issue.

This task has promising implications for developmental and psychopathology research. For example, there is evidence of age-related differences in emotion recognition capability from early childhood through to older adulthood(52, 53), evidence of sex differences in emotion processing(54, 55) and evidence that brain response to emotional expressions changes across development(56, 57). Moreover, both children and adults with a wide range of psychopathology and sub-clinical behavioural traits have been shown to have abnormalities in facial affective processing(58)(59)(60). An investigation of the relationship between behavioural traits and individual differences in emotion representation can shed light on, for example, whether individuals with elevated levels of certain traits may more heavily value specific features in certain locations on the face (e.g. eyes, mouth), whether there are biases related to certain emotions in specific populations, or whether the threshold needed to produce an emotionally-salient expression is higher or lower in some individuals. This approach also enables the investigation of the stability of emotion representations across time in different individuals, presenting promise for individualising stimuli and better targeting commonly employed emotion processing paradigms which use facial stimuli. Ultimately, such refinement can help us understand why some people fail to respond appropriately to emotional expressions.

Conclusions

We showed that the GA can efficiently and reliably generate categorically distinct prototype faces for each individual participant that are quantifiable by their component features. The task described presents a progress in the measure of individual differences in emotion processing, with applications for both developmental and clinical populations in the future. Findings are largely in line with existing evidence for separable effects of these four emotion categories while highlighting the greater degree of variability within and overlap between fear and sadness, an effect which has important implications for many commonly used paradigms investigating threat processing in psychological research. Evolutionary algorithms and the GA approach present a valid and promising

foundation upon which to conduct future research into emotion representation at the individual level.

Declaration of Interests

The authors declare no competing interests.

Acknowledgements

COC received support as a Wellcome Trust Sir Henry Wellcome Postdoctoral Fellow under Grant No. 206459/Z/17/Z. KR was funded as an EPSRC doctoral student under Grant No. EP/L016540/1. DPC, EV and IM received support from the MRC as co-investigators on Grant No. MR/S011307/1. DPC is director of the Centre for the Analysis of Motion, Entertainment Research and Applications (CAMERA), supported by the EPSRC under Grant No. EP/M023281/1. EV was supported by a British Academy Mid-Career Fellowship.

Author contributions

All authors fulfil criteria for authorship. CC, DC, EV and IM conceived of and designed the study. KR developed the study paradigm. CC and FH collected the data. CC, FH, KR and RL analysed and interpreted the data. CC wrote the paper, and all authors critically edited and revised the work and approved the final version to be submitted for publication. All authors agree to be accountable for all aspects of this work.

Data and code availability

This paper has been made available online as a preprint on PsyArXiv, DOI 10.31234/osf.io/9fta2. Data and code that support the findings of this study have been deposited on figshare (DOIs 10.6084/m9.figshare.11743308; 10.6084/m9.figshare.11743281; 10.6084/m9.figshare.11743326; 10.6084/m9.figshare.11743716).

References

1. Chen C, Jack RE. Discovering cultural differences (and similarities) in facial expressions of emotion. *Current Opinion in Psychology*. 2017;17:61-66.
2. Ekman P, Friesen WV. Constants across cultures in the face and emotion. *J Pers Soc Psychol*. 1971;17. doi:10.1037/h0030377
3. Gendron M, Roberson D, van der Vyver JM, Barrett LF. Perceptions of Emotion from Facial Expressions are Not Culturally Universal: Evidence from a Remote Culture. *Emotion*. 2014;14(2):251-262. doi:10.1037/a0036052
4. Izard CE. Innate and universal facial expressions: Evidence from developmental and cross-cultural research. *Psychological Bulletin*. 1994;115(2):288-299. doi:10.1037/0033-2909.115.2.288
5. Blair R. Facial expressions, their communicatory functions and neuro-cognitive substrates. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences*. 2003;358(1431):561-572. doi:10.1098/rstb.2002.1220
6. Jack RE, Schyns PG. The Human Face as a Dynamic Tool for Social Communication. *Current Biology*. 2015;25(14):R621-R634. doi:10.1016/j.cub.2015.05.052
7. Calvo MG, Nummenmaa L. Perceptual and affective mechanisms in facial expression recognition: An integrative review. *Cognition and Emotion*. 2016;30(6):1081-1106. doi:10.1080/02699931.2015.1049124
8. Ekman P. *Facial Action Coding System: Investigator's Guide*. Consulting Psychologists Press; 1978.
9. Hamann S, Canli T. Individual differences in emotion processing. *Current Opinion in Neurobiology*. 2004;14(2):233-238. doi:10.1016/j.conb.2004.03.010
10. Feldman Barrett L, Adolphs R, Marsella S, Martinez AM, Pollak SD. Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. *Psychological Science in the Public Interest*. 2019;20(1). <https://journals.sagepub.com/eprint/SAUES8UM69EN8TSMUGF9/full>.
11. Eugène F, Lévesque J, Mensour B, et al. The impact of individual differences on the neural circuitry underlying sadness. *Neuroimage*. 2003;19(2 Pt 1):354-364.
12. Phan KL, Wager T, Taylor SF, Liberzon I. Functional neuroanatomy of emotion: A meta-analysis of emotion activation studies in PET and fMRI. *NeuroImage*. 2002;16:331-348.
13. Capitão LP, Underdown SJV, Vile S, Yang E, Harmer CJ, Murphy SE. Anxiety increases breakthrough of threat stimuli in continuous flash suppression. *Emotion*. 2014;14(6):1027-1036. doi:10.1037/a0037801
14. Viding E, McCrory E, Seara-Cardoso A. Psychopathy. *Current Biology*. 2014;24:R871-R874. doi:http://dx.doi.org/10.1016/j.cub.2014.06.055

15. Recio G, Schacht A, Sommer W. Classification of dynamic facial expressions of emotion presented briefly. *Cognition and Emotion*. 2013;27(8):1486-1494. doi:10.1080/02699931.2013.794128
16. Ekman P. *Emotion in the Human Face*. Cambridge Cambridgeshire; New York: Cambridge University Press; Editions de la Maison des Sciences de l'Homme; 1982.
17. Tottenham N, Tanaka JW, Leon AC, et al. The NimStim set of facial expressions: Judgments from untrained research participants. *Psychiatry Research*. 2009;168(3):242-249.
18. Gendron M, Roberson D, van der Vyver JM, Barrett LF. Perceptions of emotion from facial expressions are not culturally universal: Evidence from a remote culture. *Emotion*. 2014;14(2):251-262. doi:10.1037/a0036052
19. Ersotelos N, Dong F. Building highly realistic facial modeling and animation: A survey. *Visual Comput*. 2008;24(1):13-30. doi:10.1007/s00371-007-0175-y
20. Williams L. Performance-driven Facial Animation. In: *Proceedings of the 17th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '90. New York, NY, USA: ACM; 1990:235–242. doi:10.1145/97879.97906
21. Shin HJ, Lee Y. Expression Synthesis and Transfer in Parameter Spaces. *Computer Graphics Forum*. 2009;28(7):1829-1835. doi:10.1111/j.1467-8659.2009.01560.x
22. Yu H, Garrod OGB, Schyns PG. Perception-driven facial expression synthesis. *Computers & Graphics*. 2012;36(3):152-162. doi:10.1016/j.cag.2011.12.002
23. Lewis JP, Anjyo K, Rhee T, Zhang M, Pighin F, Deng Z. Practice and Theory of Blendshape Facial Models. 2014. doi:http://dx.doi.org/10.2312/egst.20141042
24. Frowd CD, Hancock PJ, Carson D. EvoFIT: A holistic, evolutionary facial imaging technique for creating composites. *ACM Transactions on applied perception (TAP)*. 2004;1:19-39.
25. Gibson SJ, Solomon CJ, Bejarano AP. Synthesis of Photographic Quality Facial Composites using Evolutionary Algorithms. In: *Procedings of the British Machine Vision Conference 2003*. Norwich: British Machine Vision Association; 2003:23.1-23.10. doi:10.5244/C.17.23
26. Valentine T. A Unified Account of the Effects of Distinctiveness, Inversion, and Race in Face Recognition. *The Quarterly Journal of Experimental Psychology Section A*. 1991;43(2):161-204. doi:10.1080/14640749108400966
27. Valentine T, Lewis MB, Hills PJ. Face-space: A unifying concept in face recognition research. *The Quarterly Journal of Experimental Psychology*. 2015:1-24.
28. Frowd CD, Hancock PJB. Evolving Human Faces. In: Romero J, Machado P, eds. *The Art of Artificial Evolution*. Berlin, Heidelberg: Springer Berlin Heidelberg; 2008:189-210. doi:10.1007/978-3-540-72877-1_9
29. R Core Team. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing; 2013. http://www.R-project.org/.
30. van der Maaten L, Hinton G. Visualizing Data using t-SNE. *Journal of Machine Learning Research*. 2008;9:2579-2605.

31. Anderson MJ. Distance-based tests for homogeneity of multivariate dispersions. *Biometrics*. 2006;62(1):245-253. doi:10.1111/j.1541-0420.2005.00440.x
32. Oksanen J, Blanchet FG, Friendly M, et al. *Vegan: Community Ecology Package.*; 2019. <https://CRAN.R-project.org/package=vegan>.
33. Svegar D, Kardum I. Happy Face Superiority Effect in Change Detection Paradigm. *Psychological Topics*. 2013;22(2):249-269.
34. Ohman A. Fear and anxiety as emotional phenomenon: Clinical phenomenology, evolutionary perspectives, and information-processing mechanisms. In: *Handbook of Emotions*. New York: Guildford Press; 1993:511-536.
35. Fox E, Lester V, Russo R, Bowles RJ, Pichler A, Dutton K. Facial Expressions of Emotion: Are Angry Faces Detected More Efficiently? *Cogn Emot*. 2000;14(1):61-92.
36. Calvo MG, Lundqvist D. Facial expressions of emotion (KDEF): Identification under different display-duration conditions. *Behav Res*. 2008;40(1):109-115. doi:10.3758/BRM.40.1.109
37. Goren D, Wilson HR. Quantifying facial expression recognition across viewing conditions. *Vision Research*. 2006;46(8):1253-1262. doi:10.1016/j.visres.2005.10.028
38. Rapcsak SZ, Galper SR, Comer JF, et al. Fear recognition deficits after focal brain damage: A cautionary note. *Neurology*. 2000;54(3):575-575. doi:10.1212/WNL.54.3.575
39. Ekman P, Friesen WV. *Unmasking the Face: A Guide to Recognizing Emotions from Facial Expressions*. Reprint edition. Cambridge, MA: Malor Books; 2003.
40. Colonnello V, Russo PM, Mattarozzi K. First Impression Misleads Emotion Recognition. *Front Psychol*. 2019;10. doi:10.3389/fpsyg.2019.00527
41. Kätsyri J, Sams M. The effect of dynamics on identifying basic emotions from synthetic and natural faces. *International Journal of Human-Computer Studies*. 2008;66(4):233-242. doi:10.1016/j.ijhcs.2007.10.001
42. Nord CL, Gray A, Charpentier CJ, Robinson OJ, Roiser JP. Unreliability of putative fMRI biomarkers during emotional face processing. *NeuroImage*. 2017;156:119-127. doi:10.1016/j.neuroimage.2017.05.024
43. Sauder CL, Hajcak G, Angstadt M, Phan KL. Test-retest reliability of amygdala response to emotional faces. *Psychophysiology*. 2013;50(11):1147-1156. doi:10.1111/psyp.12129
44. Mill A, Allik J, Realo A, Valk R. Age-related differences in emotion recognition ability: A cross-sectional study. *Emotion*. 2009;9(5):619-630. doi:10.1037/a0016562
45. Thomas LA, Bellis MDD, Graham R, LaBar KS. Development of emotional facial recognition in late childhood and adolescence. *Developmental Science*. 2007;10(5):547-558. doi:10.1111/j.1467-7687.2007.00614.x
46. Bouhuys AL, Bloem GM, Groothuis TGG. Induction of depressed and elated mood by music influences the perception of facial emotional expressions in healthy subjects. *Journal of Affective Disorders*. 1995;33(4):215-226. doi:10.1016/0165-0327(94)00092-N

47. Killgore WDS, Cupp DW. Mood and Sex of Participant in Perception of Happy Faces. *Perceptual and Motor Skills*. 2002;95:279-288.
48. Casey BJ, Giedd JN, Thomas KM. Structural and functional brain development and its relation to cognitive development. *Biological Psychology*. 2000;54(1):241-257. doi:10.1016/S0301-0511(00)00058-2
49. Guyer AE, Monk CS, McClure-Tone EB, et al. A Developmental Examination of Amygdala Response to Facial Expressions. *Journal of Cognitive Neuroscience*. 2013;20:1565-1582. doi:10.1162/jocn.2008.20114
50. Demenescu LR, Kortekaas R, Boer JA den, Aleman A. Impaired Attribution of Emotion to Facial Expressions in Anxiety and Major Depression. *PLOS ONE*. 2010;5(12):e15058. doi:10.1371/journal.pone.0015058
51. Kang W, Kim G, Kim H, Lee S-H. The Influence of Anxiety on the Recognition of Facial Emotion Depends on the Emotion Category and Race of the Target Faces. *Experimental Neurobiology*. 2019;28(2):261-269. doi:10.5607/en.2019.28.2.261
52. Sethi A, McCrory E, Puetz V, et al. Primary and Secondary Variants of Psychopathy in a Volunteer Sample Are Associated With Different Neurocognitive Mechanisms. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*. 2018;3(12):1013-1021. doi:10.1016/j.bpsc.2018.04.002
53. Coll M-P, Viding E, Rütgen M, et al. Are we really measuring empathy? Proposal for a new measurement framework. *Neuroscience & Biobehavioral Reviews*. 2017;83:132-139. doi:10.1016/j.neubiorev.2017.10.009
54. Kozak MJ, Cuthbert BN. The NIMH Research Domain Criteria Initiative: Background, Issues, and Pragmatics. *Psychophysiology*. 2016;53(3):286-297. doi:10.1111/psyp.12518