Comparing supervised and unsupervised approaches to emotion categorization in the human brain, body, and subjective experience.

Bahar Azari^{1,*}, Christiana Westlin^{2,*}, Ajay B. Satpute², J. Benjamin Hutchinson³, Philip A. Kragel⁴, Katie Hoemann², Zulqarnain Khan¹, Jolie B. Wormwood⁵, Karen S. Quigley^{2,6}, Deniz Erdogmus¹, Jennifer Dy¹, Dana H. Brooks^{1,†}, and Lisa Feldman Barrett^{2,7,8, †}

¹Department of Electrical & Computer Engineering, College of Engineering, Northeastern University, Boston, MA, USA

²Department of Psychology, College of Science, Northeastern University, Boston, MA, USA

³Department of Psychology, University of Oregon, Eugene , OR, USA

⁴Institute of Cognitive Science, University of Colorado Boulder, Boulder, CO, USA

⁵Department of Psychology, University of New Hampshire, Durham, NH, USA

⁶Edith Nourse Rogers Veterans Hospital, Bedford, MA, USA

⁷Department of Psychiatry, Massachusetts General Hospital, Boston, MA, USA

⁸Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Boston, MA, USA

*Indicates shared first authorship.

†Indicates shared senior authorship.

Correspondence should be addressed to Dana H. Brooks (brooks@ece.neu.edu) or Lisa Feldman Barrett (l.barrett@northeastern.edu)

Machine learning methods provide powerful tools to map physical measurements to scientific categories. But are such methods suitable for discovering the ground truth about psychological categories? We use the science of emotion as a test case to explore this question. In studies of emotion, researchers use supervised classifiers, guided by emotion labels, to attempt to discover biomarkers in the brain or body for the corresponding emotion categories. This practice relies on the assumption that the labels refer to objective categories that can be discovered. Here, we critically examine this approach across three distinct datasets collected during emotional episodes – measuring the human brain, body, and subjective experience – and compare supervised classification solutions with those from unsupervised clustering in which no labels are assigned to the data. We conclude with a set of recommendations to guide researchers towards meaningful, data-driven discoveries in the science of emotion and beyond.

Introduction

Psychology became a science in the mid-19th century when scholars first used the research methods of physiology and neurology to search for the physical basis of mental categories that were inherited from ancient Western mental philosophy - categories that describe thinking, feeling, perceiving and acting. Scientific leaders of the day (e.g., [49, 94]) warned that these common-sense categories, which philosophers call “folk psychology” ([77, 23, 84, 83, 5]), would not map cleanly to physical measurements. Nonetheless, scientists have persisted for more than a century in their attempts to map measurements of the brain, body, and behavior to common-sense mental categories for cognitions, emotions, perceptions and so on. Many empirical efforts within psychological science rely on the assumption that Western folk category labels constitute the biological and psychological ‘ground truth’ of the human mind across cultures (e.g., mapping categories such as attention, emotion, or cognition onto functional brain networks;[80]), while other efforts are more neutral in their assumptions, inferring only that the category characterizes a participant’s behavior (e.g., rating of experience conditioned on experimenter-provided folk labels;[28]). Recently, however, a growing number of scientists have questioned the utility of folk category labels for organizing a science of the mind and behavior, noting the persistent difficulty of cleanly mapping these categories onto measurements of brain activity, physiological changes in the body, and behavior ([3, 4, 8, 13, 14, 24, 25, 47, 60]).

In this paper, we re-examine one aspect of the hypothesis that a natural description of the human mind – i.e., the biological organization of behavior and human experience – may require stepping back from the assumption that folk categories describe the ground truth. We do this in the context of machine learning to examine how unsupervised and supervised machine learning approaches can be used to build interpretative models from psychology measurements, across three quite different experimental settings. Unsupervised machine learning approaches can discover meaningful structure in data without assigning labels, providing a potentially valuable tool for scientific discovery in mapping biology to psychology. Supervised learning, by contrast, looks for structure in data that matches assigned labels. By comparing the results of supervised and unsupervised machine learning analyses, we can assess the extent to which psychological categories can reasonably be considered the ground truth for what exists in some objective way. Our goal is to demonstrate the importance of critically examining the labels being used, not to explain the discrepancies between supervised and unsupervised solutions or to provide evidence for a particular hypothesis. This is a proof-of-concept paper, rather than a strong inference paper, with the goal of encouraging researchers to start approaching psychology in a somewhat different way, regardless of theoretical stance. We are not suggesting that theory is unimportant in psychological science. We are instead hypothesizing that focusing empirical efforts on folk categories and building theory around those may not be the best (or only) approach. We

use emotion categories as our example because they are complex phenomena that are hypothesized to be organized as a taxonomy, and there have been numerous debates about the “ground-truth” categories. As a consequence, the domain of emotion is a clear test case in which labels that enforce strict category boundaries are traditionally imposed on the data in supervised classification methods. Whether these boundaries exist and can be discovered using unsupervised approaches is highly debated (for a discussion, see [Cicchetti 2006](#), [Adolphs 2019](#)). We selected datasets that were representative of current published research being conducted in the science of emotion, covered a range of induction techniques that have been shown to be effective ([\[59\]](#)) and common measurement modalities (self-report, psychophysiology, and fMRI), and were available to us in sufficient quantity such that variability could be adequately assessed. We present findings from supervised and unsupervised analyses of three existing datasets and show that, across all three datasets, supervised and unsupervised methods do not produce concordant results. We conclude the paper with a set of recommendations and future directions for machine learning investigations to guide researchers towards using unsupervised methods to discover biologically meaningful and reproducible discoveries about the nature of emotion and the human mind more generally

Emotion Categories: The Example

In the last decade, a number of researchers have applied supervised machine learning techniques to various types of data in an attempt to build classifiers that can identify ‘biomarkers,’ ‘signatures,’ or ‘fingerprints’ for pre-defined mental categories (for a review, see [\[56\]](#)). The science of emotion provides a particularly useful example of the problems encountered when using this approach to search for the physical basis of the human mind. Scientists begin with emotion categories for anger, sadness, fear, and so on, and then select stimuli, such as scenarios, movie clips, or music, that they expect will evoke the most frequent or common instances of each category (See [Table 1](#) for details of previous studies). These studies typically select stimuli that optimally differentiate emotion categories, and participants are exposed to these stimuli while the experimenters take measurements of their brains, bodies, and/or behavior. Scientists then aim to find evidence in the observed data for the relevant emotion categories by using supervised machine learning methods. Classifiers that are sensitive and specific to a distinct emotion category across participants and contexts ([\[56\]](#)) are then taken as evidence that the patterns or class representatives that those classifiers determine are the biomarkers of the categories. Guided by this logic, several recent studies have reported that they have identified unique patterns of brain activity, autonomic nervous system (ANS) changes, or subjective experience that discriminate emotion categories such as fear, anger, sadness, happiness, etc. ([\[52, 54, 55, 73, 74, 89, 65, 29, 66, 28, 82\]](#)). To the extent that these studies have identified markers that are

truly representative of different emotion categories, and generalize across populations and contexts, these studies may be taken by researchers as evidence that emotion categories have a reliable biological and mental basis.

However, a natural description of the human mind (i.e., the biological organization of behavior and human experience) may require stepping back from the assumption that folk categories describe the ground truth about emotions (or indeed any psychological domain). There are at least two reasons why the science of emotion is a particularly good test case to explore this question. First, the assumption that a single solution will serve as the best classifier across situations, stimuli, samples, and so on, is inherent in the language and modeling approaches used (e.g., using a single label per instance rather than multiple labels, often using deterministic rather than probabilistic models, etc; see Table 1) A close look at the published findings so far reveals that the patterns reported to be associated with a given emotion category, such as ‘fear,’ have not been consistent across studies. For example, individual studies report patterns of ANS activity that distinguish one emotion category from another, but the actual patterns vary across studies for a given emotion category, even when the studies in question use the same methods and stimuli and sample from the same population of participants (e.g., [54] and [82]). Similar cross-study, within-category variation is observed for multivoxel pattern analysis (MVPA) of blood oxygen level dependent (BOLD) signals across the brain (e.g., compare [55, 73, 89]; for a discussion, see [25]). There are several possible reasons why this has been the case. Methodological considerations might explain variation in these classification-based results ([20, 33]). For example, in brain imaging studies, small sample sizes, different affect induction methods, issues with the alignment of brains across participants, variable preprocessing workflows, and the use of variable classification algorithms might all contribute to the instability of solutions across studies (see Table 1). It is also possible that current functional brain imaging measures are insufficiently sensitive or comprehensive to identify the biomarker for a given emotion category. There may also be more than one biomarker for each category, or perhaps a single biomarker exists but multiple models capture this ground truth in different ways. But it may also be that the relationship between the presumed emotion categories, used to generate the labels for the analysis and the physiological response measured in the body or brain, are simply more complex than can be fit with a consistent relationship between the two. A purely methodological explanation may be insufficient for explaining the variation in classification-based results across studies, however, given that methodological advances for over a century have not substantially reduced the considerable variation that is observed within an emotion category across different measurement methods yielding different types of data (e.g., self-report, physiology, and brain imaging). Instead, the within-category variation that has been observed for decades (see [7, 9]) is consistent with the variation observed in classification-based results, and therefore supports the hypothesis that there is an opportunity to discover something meaningful about the

126 nature of emotion. Specifically, it suggests at least two avenues of inquiry; one would be to systematically
 127 test for one or more models across multiple datasets and modalities given a categorization. Another is to
 128 look more carefully at the relationship between the structure implied by the labels and the intrinsic structure
 129 of the acquired data across different modalities; it is an approach towards the latter inquiry that we describe
 130 here.

Study	N	# Emot. Categories	# Stimuli per Category (Unique)	Induction Method	Relevant Preprocessing	Feature Selection	Classification Algorithm
Kassam et al. (2013) [53]	10	9	2	Participant-Generated Scenario Immersion	<ul style="list-style-type: none"> • No spatial smoothing • Z-scoring • Anatomical normalization 	Voxels with the most stable activation profile	Gaussian Naive Bayes
Kragel & LaBar (2015) [56]	32	7	4	Movies, Music	<ul style="list-style-type: none"> • No spatial smoothing • Mean centering • Anatomical normalization 	All grey matter voxels	Partial least squares discriminant analysis
Saarimäki et al. (2016) [74]	48	5 (study 1) 6 (study 2)	10 (study 1) 6 (study 2)	Participant-Generated Scenario Immersion, Movies	<ul style="list-style-type: none"> • No spatial smoothing • Z-scoring • Anatomical normalization 	Voxels most sensitive to manipulation using ANOVA	Linear neural network with no hidden layers
Saarimäki et al. (2018) [75]	25	15	4	Narrative-Guided Scenario Immersion	<ul style="list-style-type: none"> • No spatial smoothing • Z-scoring • Anatomical normalization 	Voxels most sensitive to manipulation using ANOVA	Linear neural network with no hidden layers
Wager et al. (2015) [90]	2159	5	Variable	Variable	<ul style="list-style-type: none"> • Meta-analysis (peak based) • Binarized data • Anatomical normalization 	Whole brain	Bayesian spatial point process model

131

132 Table 1: **Summary of past MVPA studies of BOLD data.** These studies all claim to identify unique patterns of brain
 133 activity for specific emotion categories, yet these patterns are inconsistent across studies. Note: Other existing MVPA studies
 134 of affect (e.g., [22]), and conceptual knowledge (e.g., [79]) are less relevant and so are not listed here.

136 A second reason to question whether biology and behavior are best accounted for by a single set of emotion
 137 categories, each with a single biomarker, comes from evidence of considerable within-category variability and
 138 cross-category similarity in measurements of neural activity, ANS activity, and behavior. For example, in
 139 the domain of peripheral physiology, measurements of heart period or respiration are tied to the metabolic
 140 demands that support current or predicted actions in a given situation (e.g., cardiac output goes up when
 141 a person is about to run, but not when a person freezes and is vigilant for more information to resolve
 142 uncertainty or ambiguity; [68, 67]). People vary in their physical actions across instances of the same emotion
 143 category that occur in different situations (e.g., when experiencing fear, a person may run away, freeze, or
 144 attack); as a consequence, instances of fear will vary in their physiological features ([78]). Similar within-
 145 category variation and between-category similarities have been observed in expressive facial movements

([16]; [58]), and in neural correlates, including magnitude of the BOLD response ([93, 75, 61]), functional connectivity ([70, 86]), and single neuron recordings ([43]). Instances of an emotion category also vary in their affective features (e.g., some instances of fear can feel pleasant and some instances of happiness can feel unpleasant; [93]. This magnitude of variation, which has been replicated numerous times in the study of emotion over the past century ([37]) strengthens the possibility that imposing a single label on data may restrict or even obscure the discovery of meaningful, alternative categories in the emotion domain.

With these observations in mind, we ask whether the use of predetermined labels in building supervised classifiers might be contributing to the apparent contradiction between above-chance classification and variation across studies. Specifically, we examined whether unsupervised clustering provides a useful alternative to supervised clustering in the discovery of meaningful categories when studying the domain of emotion. Our unsupervised clustering methods consistently treat the number of clusters as a statistical parameter to be learned which, in principle, allows researchers to discover a more flexible and variable category structure should it exist, while also discovering similarities across people, should those exist. We present findings from supervised and unsupervised analyses applied to three rather diverse datasets on emotion: (1) an archival dataset during which 16 participants immersed themselves in auditory scenarios for instances labeled as inducing happiness, sadness, and fear during functional magnetic resonance imaging (fMRI) ([92]), (2) an ambulatory peripheral physiology study during which peripheral physiological signals were recorded from 46 participants throughout daily life; physiological changes (specifically, a change in interbeat interval, the time between heartbeats) triggered self-report measurement episodes in which participants freely labeled their emotional experiences ([46]), and (3) a publicly available self-report dataset in which 853 participants viewed a subset of 2,185 movie clips and reported on their experiences with categorical (yes/no) ratings of 34 emotion words ([28]) or Likert ratings of 14 affective dimensions. We subjected each dataset to a supervised classification approach with emotion category labels, as well as an unsupervised, data-driven approach which included statistically validated methods to determine the number of clusters that are objectively supported by the data. Because it is difficult to ask the same question across three different datasets that have been collected in different ways, with different constraints, we analyzed each dataset with the supervised and unsupervised approaches that were most appropriate for that particular data, while trying to the extent feasible to keep the tenets and features of the analyses consistent.

174 Results

175 Example 1: Blood Oxygen Level Dependent (BOLD) Data from [92]

176 Our first set of analyses involved an archival dataset ([92]) from a study during which sixteen participants
177 underwent fMRI scanning while immersing themselves in auditory scenarios labeled by experimenters accord-
178 ing to their intention to induce experiences of happiness, sadness, or fear. Participants heard a total of 60
179 scenarios associated with each emotion category, for a total of 180 trials across 6 runs of scanning. We ana-
180 lyzed whole-brain blood oxygen level dependent (BOLD) signals from 9s windows during which participants
181 listened to, and immersed themselves in, a single auditory scenario. Additional participant demographic
182 and task details are reported in the Methods section. Evidence suggests that participants were simulating
183 highly embodied emotional experiences when immersing themselves in the auditory scenarios. [91] analyzed
184 the same dataset used in Example 1 and observed increased activity in primary motor cortex and premotor
185 cortex and in primary somatosensory cortex while participants were lying still in the scanner, and increased
186 activity in primary visual cortex while participants’ eyes were closed. The researchers also observed increased
187 activity in brain areas that are typically considered to be important for emotion, such as the ventromedial
188 prefrontal cortex and the anterior insula, as well as primary interoceptive cortex, the thalamus, hypotha-
189 lamus, and subcortical nuclei that regulate the autonomic nervous system, immune system, and metabolic
190 systems.

191 To study whether a supervised classifier could recover the experimenter labels, we first conducted super-
192 vised classification of the BOLD data to examine classifier performance when using the folk emotion labels.
193 Specifically, we used a 3D Convolutional Neural Network ([87]) with 6-fold cross validation, iterating across
194 all combinations of training on each participant’s data from five runs and testing on a left-out run (see Sup-
195plementary Materials for classifier details). Mean within-participant accuracy was calculated by averaging
196 the accuracy across all cross-validation folds. Chance level was defined as one divided by the number of
197 categories, or $1/3$ (33.3%), and the mean accuracy was compared to chance performance using a one-sample
198 t-test. The classifier achieved a statistically significant above-chance mean within-participant accuracy of
199 46.06% ($t(15) = 9.07$, $p < 0.01$; Fig. 1a).

200 We next used an unsupervised clustering method – a Gaussian Mixture Model (GMM; [17]) – that allowed
201 us to statistically determine the number of clusters which best described the data for each participant.
202 GMMs have been used to successfully cluster BOLD data in several domains of research (e.g., [36, 88, 71]).
203 To validate the sensitivity of the GMM to detect true categories in the data, we first tested our model
204 using synthetic data generated according to a generative process described in [62]. The synthetic data were
205 designed to imitate BOLD data with clear, discoverable categories (which were classified above-chance when

206 using a supervised approach; Supplementary Fig. 5a). We included this validation step to ensure that the
 207 GMM would be able to successfully cluster the data if the relevant signal was present. Specifically, for
 208 each emotion category, we considered several randomly chosen regions in the brain with varying BOLD
 209 amplitude during each trial. Each of these regions was assigned a single radial basis function to generate a
 210 BOLD amplitude whose spatial center and width was chosen uniformly within limits of a standard human
 211 brain. The synthesized brain image for each trial was a weighted combination of these basis functions for the
 212 specific emotion category active during that trial. Zero-mean Gaussian noise was added at the last step to
 213 account for measurement noise. We then tested the sensitivity of our GMM model on the synthetic BOLD
 214 data using a concatenation of Principal Components Analysis (PCA; [50]) to reduce the dimensionality of
 215 the data followed by the Gaussian Mixture Model (GMM) to discover clusters in the lower-dimensional
 216 data (using Bayesian Information Criteria (BIC; [76]) to jointly estimate the number of PCA components
 217 and GMM clusters that best described the data). We found that the GMM was able to discover clusters
 218 in the synthetic data (Supplementary Fig. 5b). To measure our model’s ability to detect clusters in the
 219 synthetic data, we labeled the discovered clusters by the most prevalent label of the data within that cluster
 220 and then compared the proportion of samples within a given cluster whose category labels corresponded to
 221 the assigned cluster label. Even at extremely low (indeed, unlikely, given values reported in the literature
 222 ranging from 0.35 to 203.6; [90]) signal-to-noise ratio (SNR) values, the accuracy of the model was higher
 223 than chance (accuracy was 40% at an unlikely SNR of 10^{-2} , compared to chance of 33.3%). Therefore,
 224 we concluded our GMM sufficient sensitive to detect clusters that correspond to categories in real BOLD
 225 data if, in fact, clear, discoverable categories exist in the data. Additionally, the use of BIC for model order
 226 selection, which considers the likelihood of the model and a penalty term for the number of parameters,
 227 ensured that we had a sufficient sample to discover clusters. BIC is valid when sample sizes are much larger
 228 than the number of parameters in the model ([76]), and in our case, each subject had 180 trials, which is
 229 two orders of magnitude larger than the maximum number of clusters (3) in the data.

230 We then applied the same methods to estimate the GMM clusters that best described the actual BOLD
 231 data from the experiment. The results revealed variability in the number of clusters that best described each
 232 participant’s BOLD data (Fig. 1b). Furthermore, the clusters were heterogeneous, comprised of combinations
 233 of trials labeled as belonging to fear, happiness, and sadness categories, with no clear correspondence between
 234 the discovered clusters and the trial labels (Fig. 1c). We also examined the valence and arousal features of the
 235 scenarios within each cluster (available from the original dataset) but again observed no clear correspondence
 236 between these features and the clusters (Supplementary Fig. 2). Although we were unable to determine the
 237 features that the clusters were reflecting, our validation procedure with synthetic data, especially at low
 238 SNR, suggests that it is unlikely that the lack of correspondence between category labels and the clusters we

239 discovered was due to insensitivity of the unsupervised method to detect fine-grained features in the stimuli.

241 Figure 1: [BOLD Data] Results from supervised and unsupervised analyses of BOLD data from the 9s sce-
 242 nario immersion. a) Supervised clustering: Mean $\pm$ SEM classification accuracy from within-subject CNN supervised
 243 classification. The red line represents chance level accuracy (33.3%). b-c) Unsupervised clustering: b) Histogram of number of
 244 participants fit by 1, 2, or 3 GMM clusters. Specifically, eight participants had one discovered cluster, six had two, and two had
 245 three. c) Correspondence between discovered clusters and emotion category labels for two example participants (corresponding
 246 to Participant IDs 1 and 2 in part a) with two (left) and three (right) discovered clusters respectively. Bars represent the
 247 proportion of the trials from each emotion category found within each cluster. Blue bars represent trials labeled as fear, orange
 248 bars happiness, and yellow sadness. The total proportion of categories in each cluster sums to 1. The mixing proportion π_k
 249 reported below each cluster is the probability that an observation comes from that cluster, which is representative of the size
 250 of the cluster.

252 We repeated the same supervised classification and unsupervised clustering analyses using BOLD data
 253 from 3s post-stimulus intervals after each auditory scenario was finished, but before participants rated
 254 their experience of valence or arousal. During these post-stimulus intervals, participants were instructed to

continue immersing themselves in the scenario. Results from these analyses were highly similar to those that used the 9s immersion period. Specifically, we observed above-chance classification accuracy in the supervised analysis. Chance level was defined as one divided by the number of categories, or $1/3$ (33.3%), and the mean accuracy was compared to chance performance using a one-sample t-test. The classifier achieved a statistically significant above-chance mean within-participant accuracy of 47.78% ($t(15) = 25.01, p < 0.01$; Supplementary Fig. 3a). The number of estimated clusters from the unsupervised analysis varied across participants, and clusters contained a mixture of trials from each of the three labeled emotion categories (fear, happiness, sadness; Supplementary Fig. 3b,c). The valence and arousal features of trials were also mixed within each cluster (Supplementary Fig. 4). More details on this analysis are provided in the Supplementary Materials.

Example 2: Autonomic Nervous System Data from [46]

In our second comparison of supervised vs unsupervised approaches, we examined the results of a study with 46 participants who freely labeled their emotions throughout their daily life experiences while ambulatory peripheral physiological signals were recorded. It is important to note in this context that in a majority of studies, experimenters label stimuli as belonging to a specific emotion category and assume that participants would categorize the stimuli in the same way. For example, a video clip of a puppy frolicking in a park is likely to be labeled with the folk category “happiness” (i.e., is thought to induce an instance of “happiness” in all participants), yet some participants may not experience happiness when viewing this video clip (e.g., those who have just experienced a death of a pet or those who have a fear of dogs). In contrast, this study provides a useful examination of participant-, rather than experimenter-, generated labels. The participant-generated emotion labels captured variation in the categories participants used to describe affective experiences. These labels included but were not limited to the traditional folk categories of Western psychology, but they refer to categories of states that often involve intense experiences of affect (i.e., of pleasure/displeasure and arousal), two basic features that are used to characterize emotions (e.g., [2, 72]). Additionally, the labels were generated from an explicit prompt that asked participants to list an emotions they were feeling. In the words of William James, “there is no limit to the number of possible different emotions which may exist, and why the emotions of different individuals may vary indefinitely, both as to their constitution and as to objects which call them forth” ([49], pg. 454).

For 14 days over a 3-week period, participants’ electrocardiogram (ECG), impedance cardiogram (ICG), and electrodermal activity (EDA) were measured, along with their bodily movement and posture. [46] used a novel biological-triggering approach to experience sampling, where participants were prompted to

label their experience via a mobile phone app when they exhibited a pronounced and sustained change in interbeat interval (i.e., the time between successive heart beats) in the absence of physical movement. During each sampling event, participants freely labeled their current emotional experience and provided valence and arousal ratings on a 100-point sliding scale from -50 (very unpleasant/deactivated) to 50 (very pleasant/activated). Physiology data from 60s windows (30s before and after a change in interbeat interval), were analyzed. EDA data were excluded from analyses due to difficulty differentiating these signals from artifacts and because EDA changes typically occur on a slower time scale than was used in the present analyses. It was also possible to derive respiration rate (RR) from the data using the ICG signal and impedance pneumography (e.g.,[32, 34]), but we chose not to derive this measure because it is unlikely that meaningful RR changes would've occurred within the window of measurement. Six cardiovascular measures were used in the analysis at each sampling event, as detailed in Table 2. For more detail on the study design and available data, see the Methods section.

Feature	Definition	Interpretation
Interbeat interval (IBI)	Time (in ms) between heartbeats (inverse of heart rate)	IBI describes how fast the heart is beating; greater IBI values indicate slower heart rate
Respiratory sinus arrhythmia (RSA)	High frequency variability in IBI which occurs within the respiratory frequency	RSA is an estimate of parasympathetic (PNS) influence on the heart; greater RSA values indicate PNS activation
Pre-ejection period (PEP)	Time (in ms) between the beginning of electrical stimulation of the heart and the opening of the aortic valve	PEP is an inverse estimate of cardiac contractility and sympathetic (SNS) control of the heart; greater PEP values indicate reduced contractility and SNS withdrawal
Left ventricular ejection time (LVET)	Time (in ms) between the opening and closing of the aortic valve	LVET describes how long it takes the heart to pump out blood; greater LVET values are associated with greater blood volume
Stroke volume (SV)	Volume (in mL) of blood ejected by the heart with each beat	SV describes blood flow; greater SV values indicate greater blood flow per heartbeat
Cardiac output (CO)	Volume (in L) of blood circulated in the body per unit of time (m)	CO describes blood flow over time; greater CO values indicate greater blood flow

Table 2: **Cardiovascular measures derived from peripheral physiological data.** Table adapted with permission from [46].

We first conducted a supervised analysis of this dataset using each participant's three most commonly used emotion words as the labels, because the emotion labels were subject-specific and variable in number, and we only classified data from corresponding events. An average of 72.74(28.46) events were used for classification per subject across all three labels, with a minimum of 33 and a maximum of 158 events. To classify each participant's physiology data based on their top three most commonly used emotion words, we

307 trained and tuned a fully-connected neural network using 5-fold cross-validation, dividing each participant's
308 data into 5 groups and iterating across all combinations of training on 4/5ths of the data and testing on the
309 left-out 1/5th. Mean within-participant accuracy was calculated by averaging the accuracy across all cross-
310 validation folds. Chance accuracy for this dataset could not be fairly defined as the reciprocal of the number
311 of categories, because the number of events per label varied for each participant. We were unable to use
312 oversampling methods ([35]) to mitigate imbalances across categories (i.e., class imbalances) given the limited
313 number of samples (an average of 72.74 per participant). We instead evaluated statistical significance of the
314 mean accuracy across participants using a permutation test that maintained the imbalances. Specifically, we
315 re-ran the classification 1000 times for each participant using shuffled training set labels. For each iteration of
316 the permutation test we averaged the accuracy across participants, resulting in a distribution of 1000 mean
317 values for chance-level performance (chance distribution shown in Supplementary Fig. 7). Our observed
318 mean accuracy across-participants (47.1%) fell within the 5% tail of the chance distribution, indicating
319 statistically significant above-chance accuracy of our classifier at detecting whether a participant's data was
320 labeled according to one of their top three most commonly used emotion words. The per subject accuracy
321 plot is shown in Fig. 2.

322 The results of our supervised analysis were substantially different from those which resulted from an
323 unsupervised analysis reported in [46]. The unsupervised analysis used a variant of Gaussian Mixture
324 Modeling called Dirichlet Process Gaussian Mixture Modeling (DP-GMM; [17]; [19]) to discover the number
325 of clusters in the ANS data for each participant. The researchers found variable numbers of clusters for each
326 participant, with a many-to-many correspondence between participant-derived emotion labels and discovered
327 clusters in the data (Fig. 2). Participants' valence and arousal ratings also varied substantially within and
328 across the clusters they found.

329

330 Figure 2: [ANS Data] Results from supervised and unsupervised analyses of ANS data. a) Mean $\pm$ SEM classi-
 331 fication accuracy for the top three emotion words used per participant. The red line represents the mean of the null distribution
 332 for chance performance of the group. b) Distribution of participants by number of discovered physiology clusters resulting
 333 from DP-GMM. c) Correspondence between the discovered clusters and the top three participant-specific emotion categories
 334 for three representative example participants with four, five, and seven clusters respectively (corresponding to Participant IDs
 335 40, 32, and 26 in part a, with respective classification accuracy levels of 82.26%, 74.27%, and 73.80%). Bars represent the
 336 proportion of participant-generated emotion category per cluster. The total proportion of categories in each cluster sums to
 337 1. The total number of emotion category labels used by example participants 40, 32, and 26 were 26, 17, and 9, respectively.
 338 The mixing proportion π_k reported below each cluster is the probability that an observation comes from that cluster, which is
 339 representative of the size of the cluster.

341 Example 3: Self-Reports of Experience from [28]

342 In our final comparison of supervised vs. unsupervised approaches, we re-analyzed self-report ratings of
 343 experience from [28], in which 853 participants provided ratings of 2,185 evocative video clips. A subsample
 344 of participants labeled their emotional experience after viewing the film clips by choosing from among a set
 345 of 34 emotion words provided by the experimenters; each participant viewed 30 clips and made categorical
 346 (yes/no) ratings of the 34 emotion words for each clip, where a value of 0 indicated a participant was not
 347 experiencing an instance of that emotion category during the clip, and a value of 1 reflected that a participant
 348 was experiencing an instance of that emotion category while viewing the clip. A different subsample of
 349 participants rated their affective experiences; each participant in this subsample viewed 12 video clips and

350 rated them along 14 affective dimensions such as valence, arousal, effort, safety, etc., on a nine-point bipolar
 351 Likert scale (scale end-point labels were specific to each dimension, but 5 was always anchored at neutral).
 352 A traditional supervised analysis was not possible for this data set because the category labels that would
 353 be used for supervision would have to come from participants' ratings. To contrast the SH-CCA reported
 354 in [28] (see Supplementary Materials for an overview of their methods), we analyzed these data using an
 355 unsupervised clustering approach for the average yes/no emotion ratings of the 34 emotion categories across
 356 the 2,185 film clips. A GMM was not appropriate to cluster these data because the emotion category ratings
 357 were not continuous. Instead, we applied Latent Dirichlet Allocation (LDA; [18]) to discover clusters in the
 358 data. LDA is similar to GMM in that it is a mixture model, but LDA works on a collection of discrete
 359 data. It was originally developed for text corpora to discover a set of topics across a collection of documents
 360 in an unsupervised fashion. We also used LDA in this setting to perform statistical "topic modeling", in a
 361 different sense; that is, to discover the number of abstract "topics" or emotion categories (i.e., discovered
 362 clusters) occurring across all ratings of video clips. Technically, each "topic" is represented as a different
 363 probability distribution over the 34 rated emotion categories. In LDA, a perplexity plot ([18]) is used to
 364 evaluate the goodness-of-fit of the model given a fixed number of topics; its minimum indicates a suitable
 365 number of topics that optimizes the model's performance. Our perplexity plot for the ratings of these video
 366 clips showed no clear minimum value (i. e., no statistically significant minimum value; see Supplementary
 367 Materials for details) across the 34 emotion categories, indicating that the 34 emotion categories could not
 368 be reliably reduced to a smaller number of separable clusters (i.e., no unique solution was possible; Fig. 3).
 369 In other words, the LDA analysis suggested that it was not possible to discover clear 'clusters' (i.e., lower
 370 dimensional representation) in these data. As a consequence, it was not possible to examine the distribution
 371 of affective features within clusters. Our finding of no lower dimensional representation using LDA was
 372 inconsistent with the conclusion of 24 to 26 categories in [28], as well as a principal components analysis of
 373 the emotion category ratings (Supplementary Fig. 6) in which a scree plot followed by a parallel analysis
 374 [48] indicated that only five clusters best explain systematic variance, while other ways of determining the
 375 PCA truncation returned either larger values or even no truncation (consistent with the LDA analysis). In
 376 the absence of any clear consensus, one might decide that indeed there is no unambiguous structure in the
 377 data.

378

379 Figure 3: [Self-Report] Perplexity plot to discover the number of topics in LDA of emotion category self-report
 380 data. The x-axis represents the number of topics (discovered categories) and the y-axis represents the perplexity validation.
 382 There is no clear minimum value across the 34 categories, preventing us from identifying a clear clustering solution.

383 Keeping in the spirit of our first two analyses, where emotion categories were used as labels applied
 384 to BOLD and ANS data, we also approached the self-report data with a second analytic strategy. In this
 385 analysis, we used participants’ categorical emotion word ratings as the labels for the supervised analysis and
 386 the 14-dimensional vectors of mean affective ratings for the same clips as the inputs to our classification
 387 and clustering analyses. According to [28], “75% of the videos elicited significant concordance for at least
 388 one category of emotion across raters [false discovery rate (FDR) < 0.05], with concordance averaging 54%
 389 (chance level being 27%, obtained from simulated raters choosing randomly with the same base rates of
 390 category judgment observed in our data)” (pg. 3). However, when we analyzed the data we discovered that
 391 a smaller percentage of film clips could be unambiguously assigned to a single category. We assigned a label
 392 to each clip according to the highest rated emotion category for that video (e.g., if Clip 1 had a mean rating
 393 of 0.7 for sadness, 0.4 for fear, and 0.2 for anger, it was assigned the label sadness). Then, to ensure we had
 394 a sufficient number of samples for classification, we analyzed only data labeled as an emotion category that
 395 was assigned to at least 2.9% of the videos (because 34 emotion words were rated, we chose a threshold of
 396 1/34, or 2.9%). Using this criterion, we observed only 41.6% of the videos elicited concordance with a single
 397 folk emotion category. Out of the 34 emotion categories, nine categories were the highest rated category for
 398 more than 2.9% of videos (adoration, aesthetic appreciation, anxiety, awe, disgust, fear, nostalgia, romance,
 399 sexual desire). Thus we classified videos labeled with these nine categories (908 video clips out of 2185).

400 We then conducted a supervised classification of the 14 dimensional vectors of affective features for each
 401 video clip, using the nine emotion categories as labels for the corresponding videos. Specifically, we trained

402 and tuned a neural network ([17]) using 8-fold cross-validation, splitting the data into 8 groups and iterating
 403 across all combinations of training on 7/8ths of the data and testing on 1/8th of the data, to test whether a
 404 classifier could accurately detect the assigned video clip label at a rate above chance. Mean within-category
 405 accuracy was calculated by averaging the accuracy across all cross-validation folds. Chance level was defined
 406 as the reciprocal of the number of categories, or 1/9 (11.11%), and the mean across-category accuracy was
 407 compared to chance performance using a one-sample t-test. The classifier achieved a statistically significant
 408 above-chance mean across-category accuracy of 47.04% ($t(8) = 3.79$, $p < 0.01$; Fig. 4a).

409 We then examined the correspondence between participant-defined folk category analysis of affective
 410 features to an unsupervised analysis of affective features using GMM. A GMM was appropriate for this
 411 dataset because the 14 dimensional vectors of affective features are continuous variables that can be modeled
 412 with a Gaussian mean and variance. Again using the Bayesian Information Criteria (BIC) to discover the
 413 number of clusters, we found a three cluster solution of videos (see BIC curve in Fig. 4b). After estimating
 414 these three clusters in our data using a GMM, we examined the correspondence between the clustering and
 415 the emotion category labels of the videos within each cluster. The discovered clusters were comprised of
 416 a mixture of videos labeled with the nine emotion categories. For example, Cluster 1 was comprised of a
 417 mixture of videos labeled as adoration, aesthetic appreciation, awe, nostalgia, romance, and sexual desire, 6
 418 out of the 9 available labels. (Fig. 4c).

419

420 Figure 4: [Self-Report] Results from supervised and unsupervised analyses of self-report data. a) Mean +/-
 421 SEM classification accuracy for the most dominant six emotion category labels. The red line represents chance level accuracy
 422 (11.11%). b) Plot of the BIC criterion used to discover the number of clusters in the GMM model. The x-axis represents the
 423 number of clusters to use and the y-axis represents the BIC value; the minimum value of BIC is the one chosen by the BIC
 424 criterion (here the minimum is 3 clusters). c) Correspondence between clusters and the emotion category labels. Bars represent
 425 the proportion of trials from the dominant nine emotion categories within each discovered cluster. The total proportion of
 426 categories in each cluster sums to 1. The mixing proportion π_k reported for each cluster is the probability that an observation
 427 comes from that cluster, which is representative of the size of the cluster.

429 Discussion

430 We critically examined the use of emotion category labels in machine learning analyses across three studies,
 431 each sampling a different type of experiment and domain of data. In all three cases we contrasted the results
 432 of classification analyses supervised by labels with those from unsupervised approaches and observed that
 433 the methods did not produce concordant results. We obtained above-chance classification accuracy in the
 434 supervised analyses of all three datasets, suggesting that there is some information related to some of the

435 emotion category labels. When analyzing these same datasets using unsupervised approaches, however, the
436 clusters we discovered did not consistently correspond with the emotion category labels in any of the three
437 cases. Validation of our analytical procedures with synthetic fMRI data demonstrated that the clusters
438 discovered by unsupervised analyses in that case were reliable, and therefore explainable in some way.
439 However, we were unable to determine the simple, easily-nameable features that the clusters were reflecting
440 because the datasets tested were not designed for this purpose.

441 There are two possible explanations for this outcome that are worth emphasizing, either of which suggest
442 extreme caution when interpreting the results of machine learning analyses to test psychological hypotheses.
443 First, it is possible that the emotion category labels used in these three datasets (and most past studies of
444 emotion categories) veridically reflect biological and psychological categories that exist in nature and are
445 stable across contexts and individuals, and that latent emotion constructs do in fact, produce distinct, if
446 graded, responses in brain, body, and behavior. If this possibility is correct, then we would expect above-
447 chance accuracy for supervised classification because there is some reliable signal in the data that can be
448 classified. Unsupervised analysis of this same data could result in clusters that meaningfully correspond
449 to the labels, but could also result in clusters that do not correspond to the labels in any meaningful
450 way because unsupervised clustering is completely dependent on the measurements taken. The inability
451 of unsupervised methods to detect meaningful structure in the data could be due to an issue with the
452 measurements. Specifically, if the data are too sparse, contain too much noise from other sources, or have
453 undergone too much data reduction, the models may not be able to detect the structure in the data that
454 meaningfully corresponds to emotion categories. In this situation, the discovered clusters are still reliable
455 and therefore describing meaningful structure, but the structure is not reflective of emotion categories. For
456 example, it is possible that the unsupervised clustering did not pick up on the signal in the measurements
457 related to emotion categories if those signals are small compared to other signals that were unintended or
458 not of interest to the experimenter; in this case, the clustering would be based on characteristics that are
459 epiphenomenal to emotion (i.e., correlated features in the data that do not necessarily have anything to do
460 with the ground truth labeling, such as time of day, what a participant had for breakfast, etc).

461 There is a second possible explanation for our findings, one that was originally suggested by William James
462 and a growing number of modern scientists and philosophers: Western folk emotion categories may not be
463 equally useful for making sense of the biological signals. There may be substantial contextual differences,
464 individual differences, and cultural differences in how a human brain makes biological signals meaningful
465 when creating emotional events to guide action. Past studies have observed substantial variation in the
466 biological and psychological features of emotion categories, both within and across participants ([16, 78, 46],
467 and across cultures ([39, 38, 45]), suggesting that emotion categories are better thought of as populations

of highly variable, context-specific instances ([11, 10]). A populations view of emotion categories draws on
Darwin’s population thinking, in which a biological category, such as a species, consists of a population of
variable individuals ([31]; for a discussion, see [11, 10, 78]). In a populations view, the magnitude of variation
far exceeds the amount proposed in a classical or prototype view. A populations view treats variation among
individual instances within a category as real and meaningful, and posits that the prototype of a biological
category, as a single representation, is an abstract, statistical summary that need not exist in nature (for
an extended discussion, see [78, 12]). If this possibility is correct, then it would still be possible to achieve
above-chance classification using a supervised approach, due to other commonalities in the data that the
classifier could use to differentiate the data. Several studies have attempted to overcome this concern by
constructing models that generalize across different types of stimuli (e.g., emotion induction with movie
clips vs. music) within the same participants (e.g., [55, 52]) or the same stimuli across different individuals
(e.g., [52]). Nonetheless, these published studies were classifying a restricted range of emotional instances in
each category. They did not sample the full range of heterogeneous instances of emotion category that have
been observed in everyday life (as discussed in [89]). Instead, they cultivated a relatively small number of
“stereotypical” instances per category. This is not an unreasonable observation given that most studies on
emotion use stimulus sets that are curated using scientists’ intuitions about the nature of emotion categories
that may therefore fail to capture the full range of variation in real world emotional responses. Furthermore,
most studies do not attempt to generalize across both stimulus type and participants at the same time,
the notable exception being a classification study performed on a meta-analytic database [89]). This study,
which classified activation maps from different studies, across stimulus types and induction methods, noted
systematic differences in how experimenters induced different emotion categories across studies; and indeed,
methodological variables such as stimulus type and method of induction were classified above chance levels.
Consistent with these observations, the first data set we examined ([92] contained more stimulus variation
than is typical for most published studies using machine learning (see Table 1 for the number of unique
stimuli per category that is typical of past studies), and the second data set ([46]) used an objective,
empirical criterion for sampling stimulus events. Taken together, our results may suggest the possibility
that supervised machine learning is capable of discovering signal that is stipulated by the stimuli used to
manipulate emotion, but that signal may not be sufficiently strong to be detected with unsupervised analyses
alongside other real world variability in other features. This makes the accuracy of the labels critical, since
they impose the viewpoint from which that structure is discovered, and thus suggests the need for careful,
cross-study, skeptical interpretation of supervised results until validity of the labels can be scientifically
determined.

500 Implications for Future Research

501 The present findings do not allow us to discern which of these two interpretations is correct, but they do
502 highlight several important implications and course corrections for future research. First, there is increased
503 burden on researchers using supervised learning with category labels to think about and empirically explore
504 the validity of their labels, and potential constraints in their stimulus sets and methods, rather than assuming
505 that the labels are the ground truth and then sampling with that assumed prototype or stereotype in mind.
506 Studies that ask participants to report on their emotional experiences find wide variation in the number
507 and granularity of their emotion categories([10, 51]), and consequently, the emotion labels typically used in
508 supervised classification analyses, which correspond to Western folk emotion categories, may not accurately
509 encompass the repertoire of biological categories that describe all individuals equally well in all situations
510 and therefore fail to provide the best mapping to biological data (for a discussion, see [14]). It is also
511 important to consider the possibility that instances of emotion might be categorized in multiple ways,
512 simultaneously, which is often referred to as a 'mixed' emotion (as discussed in [44]). Future studies should
513 consider combinations of emotion words as descriptors of emotional instances to better capture the complexity
514 of emotional experiences.

515 To cultivate a robust and replicable science, researchers must design studies in a way that allows for an
516 explicit test of whether the category labels they apply are the best way to estimate the structure in their
517 data. Specifically, it would be optimal to sample many domains of features at higher dimensionality. This
518 would include sampling biological and psychological features in the same study in a temporally sensitive way.
519 An ideal study would measure a person's internal context (e.g., the person's metabolic condition, the past
520 experiences that come to mind) and external context (e.g., whether a person is at work, school, or home,
521 who else is present, broader cultural conditions). In addition to measuring internal and external context, an
522 ideal study would also measure a broad set of mental features that might describe an instance of emotion,
523 such as appraisal features (how the situation is experienced; e.g., as novel or familiar, requiring effort, or
524 conducive to one's goals, etc; [26, 27]) and functional features (the goals that a person is attempting to meet
525 during the instance; e.g., to avoid a predator, to feel closer to someone, to win a competition, to gain social
526 approval, etc; [1, 57]), all of which vary in dynamic ways over time. Combining these measures with an
527 objective, empirical way of sampling stimulus events (as in [46]) would provide an optimal way to accurately
528 capture and interpret meaningful structure in machine learning analyses of emotion.

529 Additionally, future researchers should pay special attention to variable patterns of measurements across
530 individuals or studies associated with the same labels, especially if these differing patterns result from classi-
531 fication using similar stimuli and/or methods. At the very least, researchers must start using more than one

machine learning method to analyze their data to discover the extent to which their modeling methods are responsible for any observed inconsistencies. We recommend that researchers routinely compare supervised and unsupervised approaches on their data and report both or at least use multiple supervised classification algorithms or feature selection methods on the same dataset to explore implications of methodological decisions on their solutions.

Finally, we echo recent suggestions that as a field, we must substantially increase the power of future studies ([63, 6, 21, 85]). Increasing the number of participants is not enough – the number of unique, variable stimuli per category is also important. [40] demonstrated that scanning a smaller number of subjects for a long duration reveals patterns of whole-brain BOLD responses that are not seen in traditional neuroimaging studies when a large number of subjects are scanned for only a short duration ([40]). Increased within-subject power will make it more likely that studies will either find something meaningful with unsupervised methods or to find more consistent results with a supervised approach. Most current datasets have too few instances and low signal, which does not allow researchers to rule out sparsity or noise as the reason for a lack of identifiable structure when using unsupervised clustering. Finally, if studies are designed with the possibility of discovering new categories (rather than merely confirming the existence of an experimenter’s preferred folk categories), and those studies are adequately powered to detect reliable clusters that explain more variance than the categories, then researchers will be better positioned to discover new mental categories to be studied in more depth.

Conclusion

The present study highlights the importance of questioning assumptions and the validity of using folk labels in the study of psychological categories. Our findings of inconsistency between supervised and unsupervised approaches can reflect two possibilities: (1) category labels in the study of emotion may reflect biological categories that exist in nature, and the observed inconsistency is due to measurement error or methodological considerations, or (2) these category labels are not useful as biological kinds that are stable across individuals and contexts. Instead, these words may refer to populations of highly variable, context-specific instances. The goal of the present study was not to demonstrate which of these possibilities is correct, but rather, to suggest that the variability observed in past studies of emotion classification may be reflecting a real discovery that should not be dismissed as error. Rather than claiming a particular truth, we are highlighting the fact that the current analytic strategy of using folk psychology categories to guide study development and data analysis should be considered. While we used the domain of emotion as a test case to highlight the importance of questioning category labels, the implications of our findings and suggestions for future

research apply to research on psychological categories across a variety of domains, and may ultimately provide a way forward in identifying mental categories that more cleanly map to measurements of brain activity, physiological changes in the body, and behavior.

Methods

Dataset 1: fMRI BOLD data

Data overview

We used data previously reported in [92], which includes comprehensive study design details. Briefly, 16 participants (8 female, ages 19 to 30 yrs) participated in the study. Participants completed a series of two training sessions, one 24-48 hours prior to the scan and the second immediately before the fMRI scan, in which they listened to auditory scenarios intended to induce the emotion categories of fear, sadness, and happiness. During the scan participants heard a shorter, core form of each scenario during a 9s trial and were instructed to fully immerse themselves mentally in each scenario as they listened. Following each 9s trial, there was a 3s post-stimulus interval during which participants continued immersing themselves in the scenario. Participants heard 30 unique scenarios for each of the three folk emotion categories. Each scenario was presented twice, resulting in a total of 180 trials. Given the assumption of past emotion classification studies (e.g., [52, 55]; [73, 74]) that all trials for a given emotion category should share the same pattern of neural activation, we included all repeating trials in our analyses. In other words, we assumed that stimuli would induce the same brain response on each presentation. Scenarios varied in valence and arousal, with some scenarios intended to evoke typical valence/emotion combinations (i.e., pleasant happiness) and others intended to induce atypical combinations (i.e., unpleasant happiness). All subjects provided informed consent, and all recruitment procedures and experimental protocols were approved by the Institutional Review Board of the original study’s institution (Emory University Institutional Review Board). All methods were carried out in accordance with relevant guidelines and regulations.

The supervised classification and unsupervised clustering analyses explained below were conducted twice, in an identical manner. Analyses were first conducted on whole-brain beta maps from the 9s scenario immersion window of each trial, and then repeated on whole-brain beta maps from the 3s post-stimulus interval after the auditory scenario was presented, while participants were still immersing themselves in the scenario.

591 Preprocessing

592 fMRI preprocessing steps were conducted in AFNI ([30]) and were identical to those in [92]: slice time
593 correction, motion correction, spatial smoothing (6mm FWHM Gaussian kernel), and percent signal change
594 normalization for each run. Spatial smoothing has not typically been performed in published MVPA studies
595 of emotion. The issue of spatial smoothing in MVPA studies is complex, however, and the utility of smoothing
596 depends on the data and question at hand ([64]). In the absence of functional alignment across individuals,
597 we expect that signals useful for clustering and classification at this spatial scale will be relatively low
598 frequency and that smoothing will improve performance.

599 Data from each 9s trial was then convolved with a canonical hemodynamic response function to generate
600 a single “beta” value for each trial, resulting in 60 “beta maps” each for the nominal emotion categories of
601 fear, sadness, and happiness. Beta maps were also generated for the 3s post-stimulus interval for each trial
602 using this same method.

603 Supervised Classification

604 To test whether a supervised classification algorithm could detect distinguishable patterns of neural activity
605 for emotion categories, we carried out within-subject classification using a 3D convolutional neural network
606 (CNN; additional details of the neural network configuration are in Supplementary Materials). This analy-
607 sis was conducted using the deep learning toolbox in MATLAB ([https://www.mathworks.com/products/
608 deep-learning.html](https://www.mathworks.com/products/deep-learning.html)). The 3D CNN was used to classify data as belonging to one of three possible condi-
609 tions: happiness, sadness, or fear. Classification was conducted on whole-brain data, and the neural network
610 was used to automatically perform feature selection, as is common in many image classification reports
611 ([81]). Specifically, dropout layers were used on the input feature layer as well as subsequent hidden layers to
612 optimize the automatic feature selection procedure and regularize the neural network weights. The classifier
613 was trained using a leave-one-run-out procedure where training was performed on BOLD data from 5 runs
614 and testing was applied to the data from the remaining 6th run. Cross-validation was performed across all
615 iterations of leaving-one-run-out and accuracy was calculated as an average percentage of correct classifica-
616 tion across all cross-validation runs. Chance level performance was defined as the reciprocal of the number of
617 categories, or 33.3%. A one-sample t-test against chance level accuracy was conducted to determine whether
618 mean classifier performance across participants was significantly above chance.

Figure 5: [BOLD Data] Probabilistic graphical model for GMM analysis.

Unsupervised Clustering

Our unsupervised clustering approach used a combination of dimension reduction and Gaussian Mixture Modeling, as described next.

Principal Component Analysis (PCA). We applied PCA to the high-dimensional (total number of fMRI voxels) beta maps. Dimension reduction was necessary to enable reliable cluster estimation given the limited number of data samples available. Statistical assumptions that underlie the choice of PCA for this purpose are described in Supplementary Materials but we note here that the number of principal components used was determined jointly with the number of Gaussian Mixture Model clusters identified, using the Bayesian Information Criterion (BIC, see Supplementary Materials for details). The result of the PCA was a feature vector α_t for each trial for each participant.

Probabilistic Graphical Model. To replace the assumption that the labels are known, we applied an unsupervised generative probability model for automatically discovering brain labels from the data α_t . The probabilistic graphical model (PGM) that describes our probabilistic dependence / independence assumptions is shown below (Fig. 5). Circles indicate random variables and the statistical dependence between pairs of random variables is shown by directed arrows (graph edges). Shaded circles indicate that the random variable is observed, while clear circles indicate that it is unobserved. Rectangles show repetitions of each random variable (e.g. over emotion category or run) with the number of repetitions shown in the right bottom corner of the rectangle.

637 The input features to the PGM were the low-dimensional representation of beta maps, the random vectors
638 α_t , with dimension D equal to the number of retained principal components obtained from PCA for each trial
639 t and for each participant (we do not use additional indices for each participant to simplify the notation).
640 There are a total of T trials per participant, each with an unknown brain label Γ_t to be estimated and an
641 emotion label Λ_t . The key innovation here is this distinction; while Λ_t is known, supplied by the experiment
642 designer, the Γ_t 's are unknown and latent. We assume that each value of Γ depends on both the associated
643 Λ and a latent "mixing coefficient" γ_{bl} , indexed by both cluster label b and emotion index l . For each
644 participant, there is a mixture of B Gaussian distributions, or soft clusters, each of which is associated to
645 a brain label. The Gaussian distribution for brain label b has a mean μ_b and a covariance matrix Σ shared
646 across the brain labels. In other words, the mixture of B Gaussian distributions models all the trials for
647 each participant (i.e., the total of T trials are soft-clustered into B clusters, each of which is represented
648 by a distinct Gaussian distributions). The GMM mixing proportions for each participant's γ_{bl} is the prior
649 probability that the given beta map comes from the b th cluster given label l .

650 ***Gaussian Mixture Modeling.*** Based on the model shown in the PGM, we estimated the parameters
651 of a GMM for each participant across all trials and all emotion category labels. Given a vector α_t calculated
652 by PCA, and an emotion category label l , the GMM has the form:

$$p(\alpha_t \mid \Lambda_t = l) = \sum_{b=1}^B \gamma_{bl} g(\alpha_t \mid \mu_b, \Sigma). \quad (1)$$

653 where b represents the labels of underlying brain activity, $b = 1, 2, \dots, B$, Λ takes on the values of each a
654 priori emotion labels, γ_{bl} is the probability of observing brain label b given label l , and μ_b and Σ
655 are the GMM model parameters. Σ is taken to be diagonal (reasonable since we applied PCA first) and
656 shared across brain labels. Maximum likelihood estimates of the GMM parameters are obtained using an
657 iterative Expectation-Maximization algorithm 1. Note that B along with the number of PCA components
658 (D , dimension of α) are determined using the Bayesian Information Criteria (BIC) across a range of possible
659 values for both (see Supplementary Materials for details).

660 **Dataset 2: Autonomic Nervous System Data**

661 **Data Overview**

662 Here we report on analyses of data from a concurrent study ([46]) in which ambulatory peripheral phys-
663 iology and self-report data from each of 46 participants was collected for 14 days over a 3-week period.
664 The signals acquired included electrocardiography (ECG), impedance cardiography (ICG), and electroder-

mal activity (EDA), along with participant’s bodily movement and posture. During the two week period, participants were prompted to respond to a survey via a smartphone device whenever their interbeat interval (IBI) exhibited a pronounced increase or decrease sustained over more than 8 seconds in the absence of any physical movement or posture change. The IBI-change threshold was tailored for each individual participant throughout the study to produce an appropriate number of survey prompts per day. The survey asked what participants were doing when they received the prompt and included a rating of their overall valence and arousal on a scale from -50 to 50 (unpleasant to pleasant; low to high arousal), and free labeling of the emotions they were experiencing. Participants received an average of 9 prompts a day. All subjects provided informed consent, and all recruitment procedures and experimental protocols were approved by the Institutional Review Board of the original study’s institution (Northeastern University Institutional Review Board). All methods were carried out in accordance with relevant guidelines and regulations.

Preprocessing

Raw physiology data was preprocessed using the researchers’ in-house Python pipeline that involved signal dependent filtering, quality control checks and feature extraction. EDA data were excluded from the analyses due to issues with differentiating these signals from artifacts, as well as the slower time scale on which changes in EDA typically occur. As part of the initial processing, researchers calculated six measures: respiratory sinus arrhythmia (RSA), interbeat interval (IBI), pre-ejection period (PEPr), left ventricular ejection time (LVET), stroke volume (SV), and cardiac output (CO). For each of the six measures, at each event, they calculated a change score as a difference of means between the 30 second period before and after the change in interbeat interval that triggered the survey prompt. This resulted in a data vector of length 6 for each event, which was used as input in the analysis.

Supervised Classification

Because both the number and value of emotion category labels were specific to each subject, we only classified events corresponding to the top three emotion words for each participant. On average, for a single participant, we classified 72.74(28.46) events out of 122.87(18.05) total events. We trained and tuned a fully-connected neural network using cross-validation, employing dropout layers to prevent overfitting. We ran within-subject classification using a 5 fold cross validation approach. Because participants varied in the number of times they used each label, chance performance could not be defined as simply the reciprocal of the number of categories. Instead, we used a permutation test to evaluate statistical significance. Specifically, we ran 1000 permutations of shuffled training labels followed by classification, averaging the classification accuracies across all participants for every iteration of the permutation. This resulted in a null distribution

696 comprised of 1000 group mean chance accuracies. We compared the actual group mean performance to the
697 null distribution of chance performance obtained by this process.

698 **Unsupervised Clustering**

699 To cluster the data vector of each event for each participant, the researchers who conducted this study
700 ([46]) used Dirichlet Process Gaussian Mixture Model (DP-GMM) to discover the number and membership
701 of clusters for each participant. DP-GMM is an infinite mixture model ([17, 19]) with the Dirichlet Process
702 as a prior distribution on the number of clusters. In practice, the approximate inference algorithm uses a
703 truncated distribution, not the infinite one, with a fixed maximum number of clusters, but the number of
704 clusters nonetheless depends on the data (see [46] for more details). They used the Scikit-learn implementa-
705 tion of DP-GMM ([69]), using number of events as the initial number of clusters for each participant. Full
706 covariance type was used with Dirichlet Process weight concentration prior and mean of the data as the
707 mean prior. Since the Scikit-learn implementation of DP-GMM tends to give slightly different results based
708 on *random_state* parameter, the method was run 100 times and the solution with the highest lower bound
709 value on the likelihood was kept.

710 **Dataset 3: Self-Report Data**

711 **Data overview**

712 Self-report data used in our analysis was taken from ([28]). In that study, 853 participants provided ratings
713 for 2,185 film clips. A subset of participants provided ratings of 14 affective dimensions on a nine-point Likert
714 scale. Each video clip was rated by nine participants, and we used the mean rating for each clip for each
715 of the 14 affective dimensions as input to our analyses. A separate subset of participants made categorical
716 (yes/no) ratings of 34 emotion categories. We used mean response data, where a value of 1 indicated that
717 the given emotion category was chosen by all participants, and a value of 0 indicated that the category was
718 never chosen. Each film was rated by 9 to 17 participants. All subjects provided informed consent, and all
719 recruitment procedures and experimental protocols were approved by the Institutional Review Board of the
720 original study’s institution (the Institutional Review Board at the University of California, Berkeley). All
721 methods were carried out in accordance with relevant guidelines and regulations.

722 **Latent Dirichlet Allocation**

723 We conducted analyses on the mean response data for each video using a topic modeling approach based on
724 Latent Dirichlet Allocation (LDA) to discover the abstract topics occurring in the collection of videos. In

725 general, Latent Dirichlet Allocation (LDA) is a generative statistical model that allows sets of observations
 726 to be explained by latent topics that explain the similarity between parts of the data ([18]). In LDA, the
 727 data from a text corpus of documents is analysed, where words in documents are used to determine both the
 728 distribution over topics and the distribution of topics over the documents. Thus in this case the assumption
 729 is that each subject’s response corresponds to an unknown number of unknown topics and the same is true for
 730 each clip In LDA, each data point is viewed as a mixture of various topics where each data point is considered
 731 to have a set of topics that are estimated by means of the LDA algorithm. The assumption typically made
 732 in LDA is that the topic distribution has a sparse Dirichlet prior, which is guided by the assumption that
 733 each data point only corresponds to a small set of topics and that each topics itself is a distributions over
 734 a small set of actual emotion categories. Or goal using LDA was to find a low-dimensional set of topics
 735 (so here discovered emotion categories) that could capture all the 34 emotion categories on which the clips
 736 were rated. Thus if low-dimensional structure was statistically justified in the ratings according to the LDA
 737 model, meaning that the emotion categories could be grouped into a smaller number of “emotion topics”,
 738 LDA is presumed to be able to discover that structure. The LDA model is shown as a probabilistic graphical
 739 model in Supplementary Figure 1. More detail on the LDA model is reported in Supplementary Materials.

740 **Supervised Classification**

741 To test whether we could successfully classify video clips as belonging to specific emotion categories based on
 742 ratings of 14 affective features, we trained and tuned a fully connected neural network using cross-validation.
 743 First we labeled the video clips based on the highest rated category across subjects (i.e., we choose the
 744 category with the maximum mean rating as the label for each video clip). Then, to ensure we had enough
 745 instances for classification of each category, we chose only those categories that were used as labels for more
 746 than 2.9% of the total video clips, which is the percentage of data samples per emotion category if data
 747 was uniformly distributed among 34 emotion categories. In other words, we only classified categories that
 748 were chosen as the highest rated category for more than 67 out of 2185 clips. Only nine emotion categories
 749 met these criteria: adoration, aesthetic appreciation, anxiety, awe, disgust, fear, nostalgia, romance, sexual
 750 desire, resulting in classification of 980 videos. The classifier was trained with a 6-fold cross validation
 751 approach; to impose a more balanced class distribution we over-sampled the minority classes (i.e, the class
 752 with the least number of samples) to be equivalent in frequency to the majority class (i.e., the class with the
 753 most number of samples). The average accuracy was calculated as the percentage of correct classifications
 754 across all validation sets. Chance level performance was defined as the reciprocal of the number of categories,
 755 or 11.11%. A one-sample t-test against chance level accuracy was conducted to determine whether mean
 756 classifier performance across emotion categories was significantly above chance.

757 **Unsupervised Clustering**

758 Similar to the fMRI analysis, here we adopted a GMM clustering technique to cluster the video clips based
759 on their 14 dimensional affective features. Each of the 14 features is a mean rating across a subset of subjects,
760 which varies from 1 to 9. The 14 dimensional affective features are, therefore, continuous bounded values in
761 each dimension and can be fit reasonably well by a GMM. The number of cluster was determined using the
762 BIC across a range of possible values for the the number of clusters, similar to what was done in Dataset 1.

763 **Author Contributions**

764 All authors were involved in conceptualizing the analyses; B.A. and C.W. performed analyses; K.H., Z.K.,
765 J.D., J.B.W., K.S.Q. analyzed data from Dataset 2; P.A.K., A.B.S., J.B.H., J.B.W., D.E, provided critical
766 feedback and edited the paper; C.W., B.A., D.H.B., and L.F.B. wrote the paper.

767 **Competing Interests**

768 No competing interests declared.

769 **Data Availability**

770 Datasets 1 and 2 are available from the corresponding authors upon request. Dataset 3 is publicly available
771 through the cited manuscript.

772 **Code Availability**

773 The code to analyze data is available from corresponding authors upon request.

774 **Note Regarding Dataset 2**

775 Our manuscript involves utilizing three existing datasets as exemplars for comparison of supervised and
776 unsupervised methods. For one of these exemplars (Dataset 2), we discuss results from an unsupervised
777 clustering analysis conducted in [46]. We present findings from this study to illustrate the main message
778 of our manuscript rather than to present a novel result. Table 2 is adapted from this manuscript with
779 permission from the authors.

References

- [1] R. Adolphs. How should neuroscience study emotions? By distinguishing emotion states, concepts, and experiences. *Social Cognitive and Affective Neuroscience*, 12(1):24–31, 2017.
- [2] R. Adolphs and D. Anderson. *The neuroscience of emotion: A new synthesis*. Princeton University Press, 2018.
- [3] M. L. Anderson. Neural reuse: A fundamental organizational principle of the brain. *Behavioral and Brain Sciences*, 33(4):245–266, 2010.
- [4] M. L. Anderson and B. L. Finlay. Allocating structure to function: the strong links between neuroplasticity and natural selection. *Frontiers in Human Neuroscience*, 7:918, 2014.
- [5] L. R. Baker. Folk psychology. In R. Wilson and F. Keil, editors, *MIT Encyclopedia of Cognitive Science*. MIT Press, 1999.
- [6] M. Bakker, A. van Dijk, and J. M. Wicherts. The rules of the game called psychological science. *Perspectives on Psychological Science*, 7(6):543–554, 2012.
- [7] L. F. Barrett. Are emotions natural kinds? *Perspectives on Psychological Science*, 1(1):28–58, 2006.
- [8] L. F. Barrett. The future of psychology: Connecting mind to brain. *Perspectives on Psychological Science*, 4(4):326–339, 2009.
- [9] L. F. Barrett. Variety is the spice of life: A psychological construction approach to understanding variability in emotion. *Cognition and Emotion*, 23:1284–1306, 2010.
- [10] L. F. Barrett. *How emotions are made: The secret life of the brain*. Houghton Mifflin Harcourt, 2017.
- [11] L. F. Barrett. The theory of constructed emotion: An active inference account of interoception and categorization. *Social Cognitive and Affective Neuroscience*, 12(1):1–23, 2017.
- [12] L. F. Barrett. John P. McGovern Award Lecture in the Behavioral Sciences: Variation is the norm: Darwin’s population thinking and the science of emotion. *American Association for the Advancement of Science 2020 Annual Meeting*, 2020. URL <https://www.youtube.com/watch?v=FUohKL5WwI8>.
- [13] L. F. Barrett and A. B. Satpute. Large-scale brain networks in affective and social neuroscience: towards an integrative functional architecture of the brain. *Current Opinion in Neurobiology*, 23(3):361–372, 2013.

- [14] L. F. Barrett and A. B. Satpute. Historical pitfalls and new directions in the neuroscience of emotion. *Neuroscience Letters*, 693:9–18, 2017.
- [15] L. F. Barrett, Z. Khan, J. Dy, and D. Brooks. Nature of emotion categories: Comment on cowen and keltner. *Trends in Cognitive Sciences*, 22(2):97–99, 2018.
- [16] L. F. Barrett, R. Adolphs, S. Marsella, A. M. Martinez, and S. D. Pollak. Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological Science in the Public Interest*, 20(1):1–68, 2019.
- [17] C. M. Bishop. *Pattern recognition and machine learning*. Springer, 2006.
- [18] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning research*, 3:993–1022, 2003.
- [19] D. M. Blei, M. I. Jordan, et al. Variational inference for dirichlet process mixtures. *Bayesian Analysis*, 1(1):121–143, 2006.
- [20] R. Botvinik-Nezer, F. Holzmeister, C. Camerer, et al. Variability in the analysis of a single neuroimaging dataset by many teams. *Nature*, 582:84–88, 2020.
- [21] K. S. Button, J. P. Ioannidis, C. Mokrysz, B. A. Nosek, J. Flint, E. S. Robinson, and M. R. Munafò. Power failure: Why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5):365, 2013.
- [22] J. Chikazoe, D. H. Lee, N. Kriegeskorte, and A. K. Anderson. Population coding of affect across stimuli, modalities and individuals. *Nature Neuroscience*, 17(8):1114, 2014.
- [23] P. M. Churchland. *A neurocomputational perspective: The nature of mind and the structure of science*. MIT press, 1989.
- [24] P. Cisek. Resynthesizing behavior through phylogenetic refinement. *Attention, Perception, & Psychophysics*, pages 1–23, 2019.
- [25] E. Clark-Polner, T. D. Johnson, and L. F. Barrett. Multivoxel pattern analysis does not provide evidence to support the existence of basic emotions. *Cerebral Cortex*, 27(3):1944–1948, 2017.
- [26] G. L. Clore and A. Ortony. Appraisal theories: How cognition shapes affect into emotion. *M. Lewis, J. M. Haviland-Jones, & L. F. Barrett (Eds.), Handbook of Emotions*, pages 628–642, 2008.

- [27] G. L. Clore and A. Ortony. Psychological construction in the OCC model of emotion. *Emotion Review*, 5(4):335–343, 2013.
- [28] A. S. Cowen and D. Keltner. Self-report captures 27 distinct categories of emotion bridged by continuous gradients. *Proceedings of the National Academy of Sciences*, 114(38):E7900–E7909, 2017.
- [29] A. S. Cowen, H. A. Elfenbein, P. Laukka, and D. Keltner. Mapping 24 emotions conveyed by brief human vocalization. *American Psychologist*, 74(6):698, 2019.
- [30] R. W. Cox. Afni: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical research*, 29(3):162–173, 1996.
- [31] C. Darwin. *On the origin of species by means of natural selection*. John Murray, 1859.
- [32] E. de Geus, G. Willemsen, C. Klaber, and L. van Doornen. Ambulatory measurement of respiratory sinus arrhythmia and respiration rate. *Biological Psychology*, 41:205–227, 1995.
- [33] M. Elliott, A. Knodt, D. Ireland, et al. What is the test-retest reliability of common task-functional mri measures? New empirical evidence and a meta-analysis. *Psychological Science*, 31:792–806, 2020.
- [34] J. Ernst, D. Litvack, D. Lozano, J. Cacioppo, and G. Berntson. Impedance pneumography: Noise as signal in impedance cardiography. *Psychophysiology*, 36:333–338, 1999.
- [35] A. Fernández, S. Garcia, F. Herrera, and N. V. Chawla. Smote for learning from imbalanced data: progress and challenges, marking the 15-year anniversary. *Journal of Artificial Intelligence Research*, 61:863–905, 2018.
- [36] G. Garg, G. Prasad, L. Garg, and D. Coyle. Gaussian mixture models for brain activation detection from fmri data. *International Journal of Bioelectromagnetism*, 13(4):255–260, 2011.
- [37] M. Gendron and L. Feldman Barrett. Reconstructing the past: A century of ideas about emotion in psychology. *Emotion Review*, 1(4):316–339, 2009.
- [38] M. Gendron, C. Crivelli, and L. F. Barrett. Universality reconsidered: Diversity in making meaning of facial expressions. *Current Directions in Psychological Science*, 27:211–219, 2018.
- [39] M. Gendron, K. Hoemann, A. N. Crittenden, S. M. Mangola, G. A. Ruark, and L. F. Barrett. Emotion perception in hadze hunter-gatherers. *Scientific Reports*, 10, 2020.
- [40] J. Gonzalez-Castillo, Z. S. Saad, D. A. Handwerker, S. J. Inati, N. Brenowitz, and P. A. Bandettini. Whole-brain, time-locked activation with simple tasks revealed using massive averaging and model-free analysis. *Proceedings of the National Academy of Sciences*, 109(14):5487–5492, 2012.

- [41] P. Good. *Permutation tests: a practical guide to resampling methods for testing hypotheses*. Springer Science & Business Media, 2013.
- [42] T. R. L. Griffiths and M. Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101:5228–5235, 2004.
- [43] S. A. Guillory and K. A. Bujarski. Exploring emotions using invasive methods: review of 60 years of human intracranial electrophysiology. *Social cognitive and affective neuroscience*, 9(12):1880–1889, 2014.
- [44] K. Hoemann, M. Gendron, and L. F. Barrett. Mixed emotions in the predictive brain. *Current Opinion in Psychology*, 15:51–57, 2017.
- [45] K. Hoemann, A. N. Crittenden, S. Msafiri, Q. Liu, C. Li, D. Roberson, G. A. Ruark, M. Gendron, and L. F. Barrett. Context facilitates performance on a classical cross-cultural emotion perception task. *Emotion*, 19:1292–1313, 2019.
- [46] K. Hoemann, Z. Khan, M. Feldman, C. Nielson, M. Devlin, J. Dy, L. F. Barrett, J. B. Wormwood, and K. Quigley. Context-aware experience sampling reveals the scale of variation in affective experience. *Scientific Reports*, 10:12459, 2020.
- [47] B. Hommel and L. S. Colzato. Learning from history: The need for a synthetic approach to human cognition. *Frontiers in Psychology*, 6:1435, 2015.
- [48] J. L. Horn. A rationale and test for the number of factors in factor analysis. *Psychometrika*, 30(2):179–185, 1965.
- [49] W. James. *The principles of psychology*. Henry Holt and Company, 1890.
- [50] I. T. Jolliffe and J. Cadima. Principal component analysis: A review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2065):20150202, 2016.
- [51] T. B. Kashdan, L. F. Barrett, and P. E. McKnight. Unpacking emotion differentiation: Transforming unpleasant experience by perceiving distinctions in negativity. *Current Directions in Psychological Science*, 24(1):10–16, 2015.
- [52] K. S. Kassam, A. R. Markey, V. L. Cherkassky, G. Loewenstein, and M. A. Just. Identifying emotions on the basis of neural activation. *PLoS One*, 8(6):e66032, 2013.

891 [53] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*,
892 2014.

893 [54] P. A. Kragel and K. S. Labar. Multivariate pattern classification reveals autonomic and experiential
894 representations of discrete emotions. *Emotion*, 13(4):681–690, 2013.

895 [55] P. A. Kragel and K. S. LaBar. Multivariate neural biomarkers of emotional states are categorically
896 distinct. *Social Cognitive and Affective Neuroscience*, 10(11):1437–1448, 2015.

897 [56] P. A. Kragel, L. Koban, L. F. Barrett, and T. D. Wager. Representation, pattern information, and
898 brain signatures: From neurons to neuroimaging. *Neuron*, 99(2):257–273, 2018.

899 [57] R. S. Lazarus and R. S. Lazarus. *Emotion and adaptation*. Oxford University Press on Demand, 1991.

900 [58] T. Le Mau, K. Hoemann, S. Lyons, J. Fugate, E. Brown, M. Gendron*, and L. Barrett. Uncovering
901 emotional structure in 604 expert portrayals of experience. (Under review).

902 [59] H. C. Lench, S. A. Flores, and S. W. Bench. Discrete emotions predict changes in cognition, judgment,
903 experience, behavior, and physiology: A meta-analysis of experimental emotion elicitations. *Psycholog-
904 ical Bulletin*, 37:834–855, 2011.

905 [60] K. A. Lindquist and L. F. Barrett. A functional architecture of the human brain: Emerging insights
906 from the science of emotion. *Trends in Cognitive Sciences*, 16(11):533–540, 2012.

907 [61] K. A. Lindquist, T. D. Wager, H. Kober, E. Bliss-Moreau, and L. F. Barrett. The brain basis of emotion:
908 A meta-analytic review. *Behavioral and Brain Sciences*, 35(3):121, 2012.

909 [62] J. R. Manning, R. Ranganath, K. A. Norman, and D. M. Blei. Topographic factor analysis: A bayesian
910 model for inferring brain networks from neural data. *PLoS One*, 9(5):e94914, 2014.

911 [63] S. E. Maxwell. The persistence of underpowered studies in psychological research: Causes, consequences,
912 and remedies. *Psychological Methods*, 9(2):147, 2004.

913 [64] M. Misaki, W. Luh, and P. A. Bandettini. The effect of spatial smoothing on fmri decoding of columnar-
914 level organization with linear support vector machine. *Journal of Neuroscience Methods*, 212:355–361,
915 2013.

916 [65] L. Nummenmaa and H. Saarimäki. Emotions as discrete patterns of systemic activity. *Neuroscience
917 Letters*, 693:3–8, 2019.

918 [66] L. Nummenmaa, R. Hari, J. K. Hietanen, and E. Glerean. Maps of subjective feelings. *Proceedings of*
919 *the National Academy of Sciences*, 115(37):9198–9203, 2018.

920 [67] P. A. Obrist. *Cardiovascular psychophysiology : A perspective*. Plenum Press New York, 1981.

921 [68] P. A. Obrist, R. A. Webb, J. R. Sutterer, and J. L. Howard. The cardiac-somatic relationship: Some
922 reformulations. *Psychophysiology*, 6(5):569–587, 1970.

923 [69] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Pretten-
924 hofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and
925 E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:
926 2825–2830, 2011.

927 [70] G. Raz, A. Touroutoglou, C. Wilson-Mendenhall, G. Gilam, T. Lin, T. Gonen, Y. Jacob, S. Atzil,
928 R. Admon, M. Bleich-Cohen, et al. Functional connectivity dynamics during film viewing reveal common
929 networks for different emotional experiences. *Cognitive, Affective, & Behavioral Neuroscience*, 16(4):
930 709–723, 2016.

931 [71] R. E. Røge, K. H. Madsen, M. Schmidt, and M. Mørup. Unsupervised segmentation of task activated
932 regions in fMRI. In *2015 IEEE 25th International Workshop on Machine Learning for Signal Processing*
933 *(MLSP)*, pages 1–6. IEEE, 2015.

934 [72] J. Russell. Core affect and the psychological construction of emotion. *Psychological Review*, 110:145–172,
935 2003.

936 [73] H. Saarimäki, A. Gotsopoulos, I. P. Jääskeläinen, J. Lampinen, P. Vuilleumier, R. Hari, M. Sams, and
937 L. Nummenmaa. Discrete neural signatures of basic emotions. *Cerebral Cortex*, 26(6):2563–2573, 2016.

938 [74] H. Saarimäki, L. F. Ejlhøj, E. Glerean, I. P. Jääskeläinen, P. Vuilleumier, M. Sams, and L. Num-
939 menmaa. Distributed affective space represents multiple emotion categories across the human brain.
940 *Social Cognitive and Affective Neuroscience*, 13(5):471–482, 2018.

941 [75] A. B. Satpute and K. A. Lindquist. The default mode network’s role in discrete emotion. *Trends in*
942 *Cognitive Sciences*, 23(10):851–864, 2019.

943 [76] G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, 1978.

944 [77] W. Sellars. *Empiricism and the Philosophy of Mind*. Harvard University Press, 1956.

[78] E. H. Siegel, M. K. Sands, W. Van den Noortgate, P. Condon, Y. Chang, J. Dy, K. S. Quigley, and L. F. Barrett. Emotion fingerprints or emotion populations? A meta-analytic investigation of autonomic nervous system features of emotion categories. *Psychological Bulletin*, 144(4):343–393, 2018.

[79] A. E. Skerry and R. Saxe. A common neural code for perceived and inferred emotion. *Journal of Neuroscience*, 34(48):15997–16008, 2014.

[80] S. Smith, P. Fox, K. Miller, D. Glahn, P. Fox, C. Mackay, N. Filippini, K. Watkins, R. Toro, A. Laird, and C. Beckmann. Correspondence of the brain’s functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, 106:13040–13045, 2009.

[81] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.

[82] C. L. Stephens, I. C. Christie, and B. H. Friedman. Autonomic specificity of basic emotions: Evidence from pattern classification and cluster analysis. *Biological Psychology*, 84(3):463–473, 2010.

[83] S. Stich and I. Ravenscroft. What is folk psychology? *Cognition*, 50(1-3):447–468, 1994.

[84] S. P. Stich. *From folk psychology to cognitive science: The case against belief*. MIT press, 1983.

[85] D. Szucs and J. P. Ioannidis. Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLoS Biology*, 15(3):e2000797, 2017.

[86] A. Touroutoglou, K. A. Lindquist, B. C. Dickerson, and L. F. Barrett. Intrinsic connectivity in the human brain does not reveal networks for ‘basic’ emotions. *Social Cognitive and Affective Neuroscience*, 10(9):1257–1265, 2015.

[87] H. Vu, H.-C. Kim, and J.-H. Lee. 3d convolutional neural network for feature extraction and classification of fMRI volumes. In *2018 International Workshop on Pattern Recognition in Neuroimaging (PRNI)*, pages 1–4. IEEE, 2018.

[88] E. Vul, D. Lashkari, P.-J. Hsieh, P. Golland, and N. Kanwisher. Data-driven functional clustering reveals dominance of face, place, and body selectivity in the ventral visual pathway. *Journal of Neurophysiology*, 108(8):2306–2322, 2012.

[89] T. D. Wager, J. Kang, T. D. Johnson, T. E. Nichols, A. B. Satpute, and L. F. Barrett. A bayesian model of category-specific emotional brain responses. *PLoS Computational Biology*, 11(4):e1004066, 2015.

973 [90] M. Welvaert and Y. Rosseel. On the definition of signal-to-noise ratio and contrast-to-noise ratio for
974 fmri data. *PLoS One*, 8:e77089, 2013.

975 [91] C. Wilson-Mendenhall, A. Henriques, L. W. Barsalou, and L. F. Barrett. Primary interoceptive cortex
976 activity during simulated experiences of the body. *Journal of Cognitive Neuroscience*, 31:221–235, 2019.

977 [92] C. D. Wilson-Mendenhall, L. F. Barrett, and L. W. Barsalou. Neural evidence that human emotions
978 share core affective properties. *Psychological Science*, 24(6):947–956, 2013.

979 [93] C. D. Wilson-Mendenhall, L. F. Barrett, and L. W. Barsalou. Variety in emotional life: within-category
980 typicality of emotional experiences is associated with neural activity in large-scale brain networks. *Social*
981 *Cognitive and Affective Neuroscience*, 10(1):62–71, 2014.

982 [94] W. M. Wundt and C. H. Judd. *Outlines of psychology*. Scholarly Press, 1897.

Supplementary Material

Supplementary Methods

All analyses were conducted in MATLAB version R2019a using custom code, unless otherwise mentioned.

Dataset 1: fMRI BOLD data

Convolutional Neural Network Details

The details of the 3d Convolutional Neural Network configuration are as follows:

- Three 3D convolutional layers with kernel size of 5 were used.
- After each convolutional layer, we used a batch normalization layer to speed up training and reduce sensitivity to network initialization.
- Between each (feature/hidden) layer we also use dropout layers to regularize the neural network weights and optimize the selected features, [81].
- A hyperbolic tangent activation layer was applied after batch normalization.
- Downsampling by a factor of 2 with stride 2 was performed using a max pooling layer on the activation outputs.
- Finally, a fully connected layer followed by a soft max produced the classification probabilities.
- In training, the batch size was 20, the maximum epoch was 10, and the learning rate was chosen as a hyperparameter to be equal 0.001.
- ADAM optimization algorithm, [53], was used to iteratively update network weights in the training procedure.

Joint Estimation of Principal Component and Cluster Number using BIC

We performed a grid search over both the value of D , the number of principle components to retain, and the values of B , number of clusters/mixture components. We evaluated the model for every combination of (B, D) in a reasonable range using the BIC. Specifically, the value for D was varied from 1 to 179 (one less than the number of trials $T = 180$) and is the maximum dimension of variance that can be captured with $T < V$ trials, where V is the number of voxels. The value for B was varied from 1 to 10. We chose the (B, D) pair from this range that yielded the minimum BIC.

1009 Details of Model Performance Evaluation using Synthetic Data

1010 As described in the main text, we generated synthetic data to test the ability of our clustering model to
1011 capture the signal induced by stimuli embedded in noise. We generated synthetic data using the neurosim
1012 MATLAB package (<https://github.com/ContextLab/neurosim>) with the generative process proposed in
1013 ([62]). The synthetic data was designed to imitate BOLD data with clear, discoverable categories. For
1014 a given signal-to-noise ratio (SNR), an experimental design matrix, and brain voxel location, the package
1015 generates a set of synthesized voxel activations. For each emotion category, we used 20 randomly chosen
1016 spherical regions in the brain with varying BOLD amplitude during each trial. Each of these regions was
1017 assigned a single radial basis function whose spatial center and width was chosen uniformly but restricted
1018 to remain within the limits of a standard human brain. The synthesized “brain image” for each trial t was a
1019 weighted combination of these 20 basis functions for the specific emotion category active during that trial.
1020 Zero-mean Gaussian noise was added as the last step to simulate the measurement noise. The standard
1021 deviation of the Gaussian noise was calculated using the SNR supplied. We varied the SNR in our study
1022 from 0.01 to 100 in the range of [0.01, 0.05, 0.1, 0.5, 1, 5, 10, 100] to assess the performance of our model
1023 under varying levels of noise. In a Monte Carlo simulation, for each SNR value, we performed 500 random
1024 realizations of synthesized brain images and calculated the accuracy of our model at clustering the emotion
1025 categories. We reported the average accuracy across all random realizations. To evaluate the supervised
1026 performance on the same synthetic data we used the same structure for the CNN as was used for the actual
1027 fMRI BOLD data, described above.

1028 Dataset 2: Autonomic Nervous System Data

1029 Neural Network Details

1030 The Details of the neural network configuration are as follows:

- 1031 • Three fully-connected layers of size 5, 4, 3 were used.
- 1032 • After each fully-connected layer we used a batch normalization layer to speed up training of the network
1033 and reduce the sensitivity to network initialization.
- 1034 • A rectified linear unit (ReLU) activation layer was applied after batch normalization.
- 1035 • Finally a soft max layer generated the classification probabilities.
- 1036 • The batch size was 10, the maximum epoch was set to 50, and the learning rate was chosen as a
1037 hyperparameter to be equal 0.001.

- ADAM optimization algorithm, Kingma and Ba [53], was used to iteratively update network weights in the training procedure.

Dataset 3: Self-Report Data

Overview of Cowen & Keltner 2017 split-half canonical correlation analysis

The experimenters did not provide emotion labels for the film clips, but the clips were chosen with those specific 34 labels in mind (from Cowen & Keltner, “The videos were gathered by querying search engines and content aggregation websites with contextual phrases targeting 34 emotion categories”; pg. 2 [28]). However, it can be debated whether the initial analyses reported in [28] were free from strong assumptions about the existence of certain emotion categories. The researchers devised a split-half canonical correlation analysis (SH-CCA) in which they correlated the mean emotion category ratings for each clip from half the subsample of participants with the mean ratings from the other half. They reported that the results of the SH-CCA indicated that the data contained between 24 and 26 categories of emotional experience. Because the categories in the SH-CCA were pre-specified rather than discovered (i.e., one half of the categorical ratings constrained the analysis of data the other half), this method is more constrained than is optimal for a fully unsupervised analysis. Indeed, it might be argued that the categorical nature of the ratings and using the same labels for video selection and for participant ratings makes this SH-CCA more likely to reveal a solution consistent with a supervised classification than data-driven clustering ([15]).

Latent Dirichlet Allocation

The LDA model is shown as a probabilistic graphical model below. The interpretation of the symbols in the graph is the same as described in the Methods section for Dataset 1.

Supplementary Figure 1: [Self-Report] Probabilistic graphical model for LDA analysis

Specifically, we have a collection of D video clips, each of which is a mixture over K topics (i.e., video

clip $d \in \{1, \dots, D\}$ has a associated K dimensional vector θ_d that shows its distribution over topics). Each topic k is characterized by 34-dimensional vectors ψ_k , which is a distribution over the predefined emotion categories. For each video clip, N is the total number of 'yes' answers for all of the pre-defined emotion categories, W is a categorical variable showing the 'yes' answer for a specific predefined emotion category, and Z is the corresponding topic index. Finally, α and β are the parameters of the Dirichlet priors for the topic mixtures θ and topics ϕ , respectively. The corresponding probabilistic generative model for this LDA model is as follows:

$$\begin{aligned}\theta &\sim \text{Dirichlet}(\alpha), \\ \phi &\sim \text{Dirichlet}(\beta), \\ Z &\sim \text{Categorical}(\theta), \\ W &\sim \text{Categorical}(\phi),\end{aligned}$$

where the Dirichlet distribution is given by

$$p(\theta \mid \alpha) = \frac{1}{\mathbf{B}(\alpha)} \prod_K \theta_i^{\alpha_i - 1}$$

1061 and $\mathbf{B}(\alpha)$ denotes the multivariate Beta function.

1062 We adapted a collapsed Gibbs sampling method for LDA model based on [42] in MATLAB's Text Ana-
1063 lytics Toolbox to solve for the unknown parameters of the model, including the topic distributions and the
1064 distribution over topic ϕ for each video clip, θ .

1065 To decide on a suitable number of topics for LDA, we compared the goodness-of-fit of LDA models across
1066 varying numbers of topics. We evaluated the goodness-of-fit of an LDA model after by training on all videos
1067 except for a held-out subset by calculating the perplexity on the held-out subset:

$$\text{perplexity}(D_{test}) = \exp \left\{ -\frac{\sum_{d=1}^M \log p(W \mid \alpha, \psi)}{MN} \right\} \quad (2)$$

1068 where D_{test} is the held-out collection of M videos. The perplexity indicates how well the model described
1069 the data, with a lower perplexity suggesting a better fit. We used 10-fold cross-validation by randomly
1070 partitioning the original dataset into ten equal size subsets. Nine of the subsets were used for training, and
1071 the remaining subset for validation. The cross-validation process was then repeated 10 times for all folds
1072 and each of the 10 subsets was used exactly once as the validation set. The average perplexity across 10

1073 folds was used for choosing the number of topics (see Fig. 3 in the main document).

1074 **Neural Network Details**

1075 The Details of the neural network configuration are as follows:

- 1076 • Three fully-connected layers of size 10, 8, 6 were used.
- 1077 • After each fully-connected layer we used a batch normalization layer to speed up training of the network
1078 and reduce the sensitivity to network initialization.
- 1079 • A rectified linear unit (ReLU) activation layer was applied after batch normalization.
- 1080 • Finally a soft max layer generated the classification probabilities.
- 1081 • The batch size was 20, the maximum epoch was set to 10, and the learning rate was chosen as a
1082 hyperparameter to be equal 0.001.
- 1083 • ADAM optimization algorithm, Kingma and Ba [53], was used to iteratively update network weights
1084 in the training procedure.

1085 **Permutation Test**

1086 To evaluate the statistical significance of our classification findings, we conducted a permutation test
1087 based on the definition from [41]. The details of the procedure are as follows. Given the original data set
1088 $D = \{(\mathbf{X}_i, y_i)\}_i^n$ and a permutation function for n elements, π , one can permute the labels y of data set D
1089 with the aim to produce a new data set $D' = \{(\mathbf{X}_i, \pi(y)_i)\}_i^n$ whose marginal distributions of features $p(\mathbf{X})$
1090 and labels $p(y)$ are the same as those of the original data set while the dependence between features and
1091 labels are broken. The new data set D' is intuitively known to come from chance. Based on the definition
1092 by Good (2000), we defined $\mathcal{D}$ as a set of k randomized permutations D' of the original data set D sampled
1093 from a null distribution. We then calculated the p -value based on the following equation, where $a(f, D')$ is
1094 the training accuracy computed using the permuted labels in D' :

$$p = \frac{|D' \in \mathcal{D} : a(f, D') \geq a(f, D)| + 1}{k + 1}, \quad (3)$$

1095 The $|\cdot|$ represent the cardinality (size) of the set. The calculated p -value represents the fraction of permuted
1096 samples where the classifier’s accuracy was better in the random, permuted data than the original data,
1097 reflecting the probability that the classification accuracy for our data was obtained by chance.

1098 **No Significant Minimum in Perplexity Plot (Fig. 3)**

1099 In Fig. 3, we cannot decide in favor of a specific value for the number of clusters because there is not a clear
1100 minimum. To confirm that there is no minimum attained that is statistically significant, we ran a paired
1101 t-test for each pair consisting of the smallest value at 31 and one of the adjacent values at 30 or 32. We used
1102 the perplexity value across 10 folds of validation as inputs in the t-test. According to the t-test result, the
1103 perplexity value at 31 was not significantly different from 30 or 32 ($p > 0.1$). Therefore, we cannot choose
1104 that as the minimum perplexity.

1105

1107

1108 Supplementary Figure 2: [BOLD Data] Correspondence between clusters and valence/arousal features for
1109 BOLD data from the 9s scenario immersion. Left/right columns show data for two example participants with two and
1110 three clusters, respectively. Top row shows results for arousal and bottom row for valence. Specifically in the top row the heights
1111 of the bars represent the proportion of low arousal (blue) and high arousal (red) trials in each cluster while in the bottom row
1112 the bars represent the proportion of unpleasant (blue) and pleasant (red) trials in each cluster. The total proportion in each
1113 cluster sums to 1.

1115

1116 **Supplementary Figure 3: [BOLD Data] Results from supervised and unsupervised analyses of BOLD data from**
 1117 **the 3s post-stimulus interval.** a) Supervised clustering: Mean +/- SEM classification accuracy from within-subject CNN
 1118 supervised classification. The red line represents chance level accuracy (33.3%). b)-c) Unsupervised clustering: b) Histogram
 1119 of number of participants fit by 1 through 6 GMM clusters. Specifically, eight participants had one discovered cluster, three
 1120 participants had two clusters, four had three clusters, and one had six clusters. c) Correspondence between discovered clusters
 1121 and emotion category labels for two example participants (corresponding to Participant IDs 1 and 2 in part a) with two (left)
 1122 and three (right) discovered clusters respectively. Bars represent the proportion of the trials from each emotion category found
 1123 within each cluster. Blue bars represent trials labeled as fear, orange bars happiness, and yellow sadness. The total proportion
 1124 of categories in each cluster sums to 1. The mixing proportion π_k reported below each cluster is the probability that an
 1126 observation comes from that cluster, which is representative of the size of the cluster.

1127

1128 Supplementary Figure 4: [BOLD Data] Correspondence between clusters and valence/arousal features for
 1129 BOLD data from the 3s post-stimulus interval. Left/right columns show data for two example participants corresponding
 1130 to Participant IDs 1 and 2 with two and three clusters, respectively. Top row shows results for arousal and bottom row for
 1131 valence. Specifically in the top row the heights of the bars represent the proportion of low arousal (blue) and high arousal (red)
 1132 trials in each cluster while in the bottom row the bars represent the proportion of unpleasant (blue) and pleasant (red) trials
 1133 in each cluster. The total proportion in each cluster sums to 1.

1135

1136 Supplementary Figure 5: [Synthetic BOLD Data] Evaluating modeling performance using synthetic data.
 1137 Accuracy for supervised classification (left, (a)) and unsupervised clustering (right, (b)) for increasing SNR for the synthetic
 1138 dataset generated based on the assumption that unique patterns of brain activity exists for different emotion categories.
 1139 Unsupervised clustering accuracy represents the correspondence of the unsupervised clusters with the category labels, i.e., each
 1140 category is associated to the most probable cluster according to the GMM. Shaded bands represent standard deviation.

1142

1143 Supplementary Figure 6: [Self-Report] Results from a principal component analysis of self-report data. The
 1144 plot depicts principle component variance (i.e. eigenvalues of the covariance matrix accounted for by each successive value) as
 1145 a function of principle component index. Actual values are shown in blue squares; the line is drawn for better visualization.
 1146 Different methods indicate different numbers of components to retain. The dashed line indicates a the retention criteria of
 1147 eigenvalues greater than one, which indicates eight components. Alternately, a parallel analysis [48] to determine the number
 1148 of components suggested five components, indicated by the 95% and 5% quantile of the distribution of shuffled data.

1150

1151 Supplementary Figure 7: [ANS Data] Distribution of chance level accuracy obtained by the permutation
 1152 test. We ran 1000 permutations of classification on shuffled training labels and averaged classification accuracies across all
 1153 participants for every permutation. The average classifier accuracy across subjects was 47.1%, which fell within the 5% tail of
 1154 the chance distribution, indicating statistical significance.