Single-Unit Responses Selective for Whole Faces in the Human Amygdala

Ueli Rutishauser,[1,2,*] Oana Tudusciuc,[3] Dirk Neumann,[3]
Adam N. Mamelak,[4] A. Christopher Heller,[5] Ian B. Ross,[5]
Linda Philpott,[5] William W. Sutherling,[5]
and Ralph Adolphs[2,3]
[1]Department of Neural Systems, Max Planck Institute for Brain Research, 60528 Frankfurt am Main, Germany
[2]Division of Biology
[3]Division of Humanities and Social Sciences
California Institute of Technology, Pasadena, CA 91125, USA
[4]Department of Neurosurgery, Cedars-Sinai Medical Center, Los Angeles, CA 90048, USA
[5]Epilepsy and Brain Mapping Program, Huntington Memorial Hospital, Pasadena, CA 91105, USA

Summary

The human amygdala is critical for social cognition from faces, as borne out by impairments in recognizing facial emotion following amygdala lesions [1] and differential activation of the amygdala by faces [2–5]. Single-unit recordings in the primate amygdala have documented responses selective for faces, their identity, or emotional expression [6, 7], yet how the amygdala represents face information remains unknown. Does it encode specific features of faces that are particularly critical for recognizing emotions (such as the eyes), or does it encode the whole face, a level of representation that might be the proximal substrate for subsequent social cognition? We investigated this question by recording from over 200 single neurons in the amygdalae of seven neurosurgical patients with implanted depth electrodes [8]. We found that approximately half of all neurons responded to faces or parts of faces. Approximately 20% of all neurons responded selectively only to the whole face. Although responding most to whole faces, these neurons paradoxically responded more when only a small part of the face was shown compared to when almost the entire face was shown. We suggest that the human amygdala plays a predominant role in representing global information about faces, possibly achieved through inhibition between individual facial features.

Results

Behavioral Performance

We recorded single-neuron activity from microwires implanted in the human amygdala while neurosurgical patients performed an emotion categorization task. All patients (12 sessions from 10 patients, 1 female) were undergoing epilepsy monitoring and had normal basic ability to discriminate faces (see Table S1 available online). Patients were asked to judge for every trial whether stimuli showing a face or parts thereof were happy or fearful (Figure 1) by pushing one of two buttons as quickly and accurately as possible. Each individual face stimulus (as well as its mirror image) was shown with both happy and fearful

expressions, thus requiring subjects to discriminate the emotions in order to perform the task (also see Figure S1D). Each stimulus was preceded by a baseline image of equal luminance and complexity ("scramble"). We showed the entire face (whole face, WF), single regions of interest (eye or mouth "cutouts," also referred to as regions of interest [ROIs]), and randomly selected parts of the face ("bubbles"; Figure 1A). The randomly sampled bubbles were used to determine which regions of the face were utilized to perform the emotion classification task using a reverse correlation technique [9]. The proportion of the face revealed in the bubble stimuli was adaptively modified to achieve an asymptotic target performance of 80% correct (Figure 1B; see Supplemental Experimental Procedures); the number of bubbles required to achieve this criterion decreased, on average, over trials (Figure 1B). Average task performance across all trial categories was 87.8 ± 4.8% (n = 12 sessions, ± standard deviation [SD]; worst performer was 78% correct; see Figure S1A for details). The behavioral classification image derived from the accuracy and reaction time (RT) of the responses showed that patients utilized information revealed by both the eyes and the mouth region to make the emotion judgment (Figure 1C; Figure S1B). Overall, the behavioral performance-related metrics confirmed that patients were alert and attentive and had largely normal ability to discriminate emotion from faces (cf. Figure S1).

Face-Responsive Neurons

We isolated a total of 210 single units (see Supplemental Experimental Procedures for isolation criteria and electrode location within the amygdala) from nine recording sessions in seven patients (three patients contributed no well-isolated neurons in the amygdala). Of these, 185 units (102 in the right amygdala, 83 in the left) had an average firing rate of at least 0.2 Hz and were chosen for further analysis. Only correct trials were considered. To analyze neuronal responses, we first aligned all trials to the onset of the scramble or face epochs and compared the mean firing rate before and after. We found that 11.4% of all units showed a significant modulation of spike rate already at the onset of the scramble (Table S2; see Figure 2A for an example), indicating visual responsiveness [10, 11], whereas 51.4% responded to the onset of the face stimuli relative to the preceding baseline (Table S2). Thus, although only about a tenth of units responded to phase-scrambled faces relative to a blank screen, half responded to the facial stimuli relative to the scramble. Some units increased their firing rate, whereas others decreased their rate in response to stimulus onset (42% and 36% of the responsive units increased their rate for scramble and face stimuli onset, respectively; Table S2; Figure S2). The large proportion of inhibitory responses may be indicative of the dense inhibitory network within the amygdala [12].

Of these face-responsive neurons, 36.8% responded in bubble trials, 23.8% to whole faces, and 14.1% and 20.0% to eye and mouth cutouts, respectively (all relative to scramble baseline); some units responded to several or all categories (see Figure 2A for an example). To assess relative selectivity, we next compared responses among different categories of face stimuli (see Table S2 for comprehensive summary). We
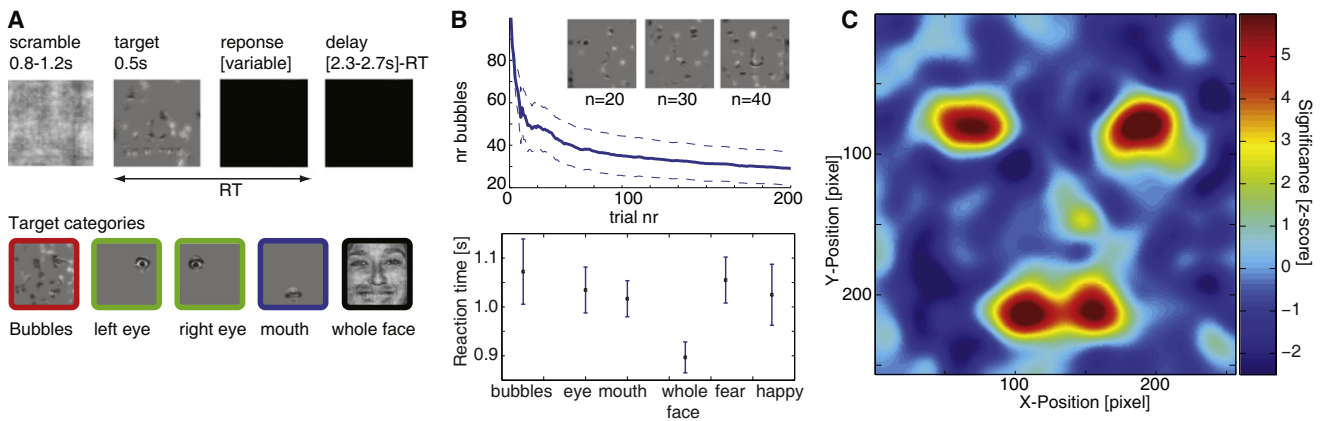
*Correspondence: urut@caltech.edu

Figure 1. Stimuli and Behavior

(A) Timeline of stimulus presentation (top). Immediately preceding the target image, a scrambled version was presented for a variable time between 0.8 and 1.2 s. Next, the target image was presented for 500 ms and showed either a fearful (50%) or happy (50%) expression. Subjects indicated whether the presented face was happy or fearful by a button press. The target presentation was followed by a variable delay. Target images and associated color code used to identify trial types in later figures are shown (bottom).

(B and C) Behavioral performances from the patients.

(B) Learning curve (top) and reaction time (bottom) (n = 11 and 12 sessions, respectively, mean ± standard error of the mean [SEM]). The inset shows example stimuli for 20, 30, and 40 bubbles revealed. Patients completed on average a total of 421 bubble trials, and the average number of bubbles required ranged from 100 at the beginning to 19.4 ± 7.9 on the last trial (n = 11 sessions, ±standard deviation [SD]; one session omitted here because the learning algorithm was disabled as a control, see results). The average reaction time was fastest for whole faces and significantly faster for whole faces than bubble trials (897 ± 32 ms versus 1,072 ± 67 ms, p < 0.05, n = 12 sessions, relative to stimulus onset).

(C) Behavioral classification image (n = 12 sessions). Color code is the z scored correlation between the presence or absence of a particular region of the face and behavioral performance: the eye and mouth regions conveyed the most information, as described previously [9]. See Figure S1 for further analyses of behavioral performance.

found that a substantial proportion of units (19.5%) responded selectively to whole faces, compared to cutouts (Figure 2; Figure S2). Only a small proportion of units distinguished between eye and mouth cutouts or between cutouts and bubbles (<10%). We found on the order of 10% of neurons whose responses differentiated between emotions, gender, or identity, similar to a prior report [6]. We thus conclude that (1) amygdala neurons responded notably more to face features than unidentifiable scrambled versions otherwise similar in low-level properties, (2) of the units responding to face stimuli, some responded regardless of which part of the face was shown, and (3) approximately 20% of all units, however, responded selectively only to whole faces and not to parts of faces, a striking selectivity to which we turn next.

Whole-Face-Selective Neurons

We next focused on the whole-face (WF)-selective units, defined in our study as those that responded differentially to WFs compared to the cutouts (n = 36). The majority of such units showed no correlation with RT of the patient's behavioral response (only 3/36 showed a significant positive correlation, and 2/36 a significant negative correlation with RT in the bubble trials; 1/36 showed a significant positive correlation with RT in the WF trials), favoring a sensory over a motor-related representation. The majority of the units (32 of 36, 89%) increased their firing rate for WFs relative to bubble trials. Focusing on these units that increase their rate (see below for the others), the first temporal epoch showing significantly differential responses to WFs and bubble trials was 250–500 ms after stimulus onset (Figure 3A). Note that this is an independent confirmation of the response selectivity, because only cutouts rather than bubble trials were used to define the WF selectivity of the neurons to begin with (see Supplemental Experimental Procedures).

How representative are the WF-selective neurons of the entire population of amygdala neurons? To quantify the differential response across all neurons to WFs compared to bubble stimuli, we calculated a whole-face index (WFI; see Supplemental Experimental Procedures) as the baseline-normalized difference in response to whole faces compared to bubbles. The average WFI of the entire population (n = 185) was 11% ± 3% (significantly different from zero, p < 0.0005), showing a mean increase in response to WFs compared to bubbled faces. The absolute values of the WFI for the previously identified class of WF-selective units (n = 36) and all other units (n = 149) were significantly different (53% ± 7% and 18% ± 2%, respectively; p < 1e-7; Figures 3C and 3D). We conclude that a subpopulation of about 20% of amygdala neurons is particularly responsive to WFs.

Nonlinear Face Responses

We next systematically analyzed responses of WF-selective neurons as a function of the proportion of eye and mouth region that was revealed in each bubble trial (number of bubbles that overlap with the eye and mouth ROI) (Figure 4A). Because mean firing rates varied between 0.2 and 6 Hz (cf. Figure 3B), we assured equal weight from each unit by normalizing (for each unit) the number of spikes relative to the number of spikes evoked by the WF. The resulting normalized response as a function of the proportion of the ROI that was revealed across the bubble trials is shown for several representative single units in Figure 4B. We found several classes of responses: some did not depend on the proportion of the face revealed (Figure 4B3), some increased as a function of the proportion revealed (Figures 4B6 and 4B7), and some decreased (Figures 4B1, 4B2, 4B4, and 4B5). Statistically, most individual units had a response function whose slope did not achieve significance (28 of 36 units), and thus most
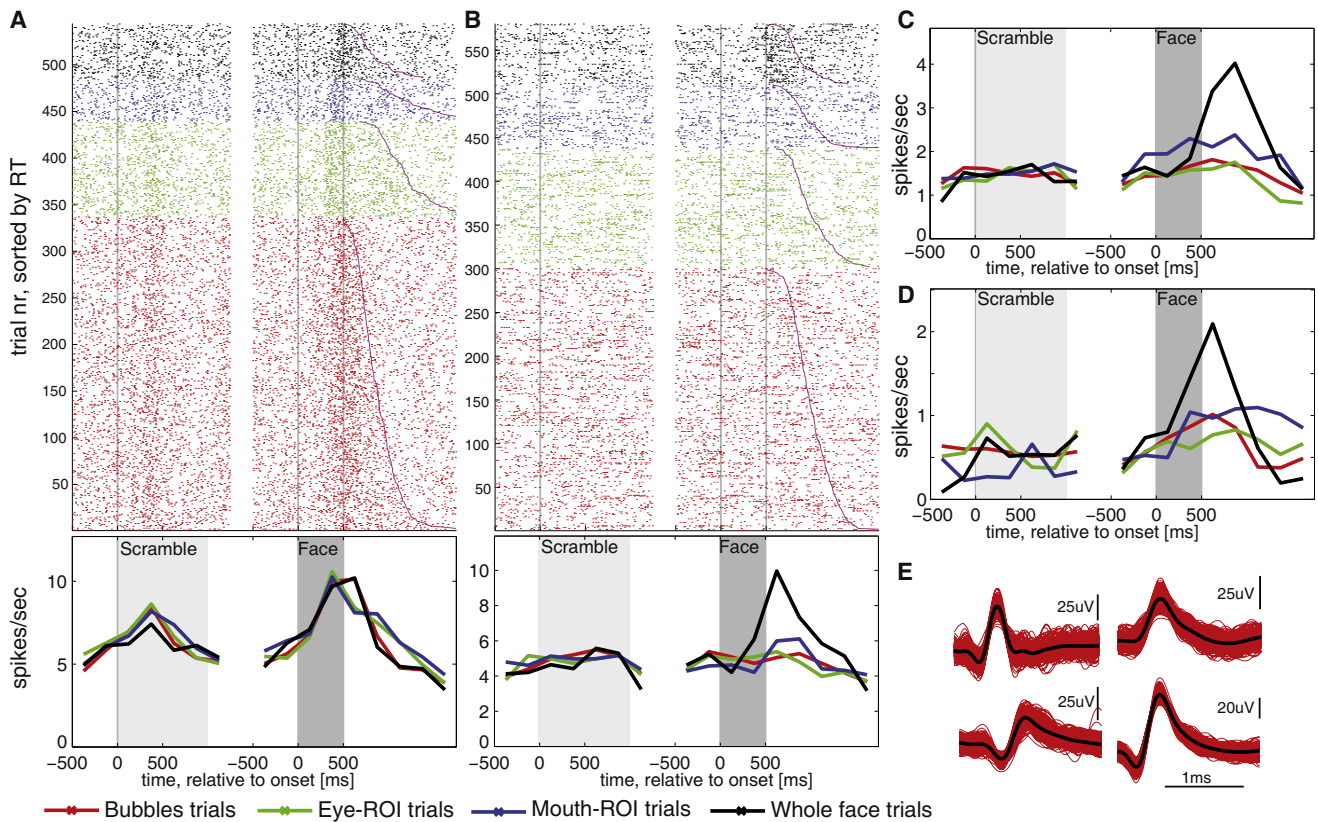
Figure 2. Single Unit Responses in the Amygdala

(A–D) Examples of responses from four different neurons. Some responded to all trials containing facial features (A), whereas others increased their firing rates only to whole faces (B–D). Each of the units is from a different recording session, and for each the raster (top) and the poststimulus time histogram (bottom) are shown with color coding as indicated. Trials are aligned to scramble onset (light gray, on average 1 s, variable duration) and face stimulus onset (dark gray, fixed 500 ms duration). Trials within each stimulus category are sorted according to reaction time (magenta line).

(E) Waveforms for each unit shown in (A)–(D). Figure S2 shows the rasters for the units shown in (C) and (D).

units did not clearly increase or decrease their firing rate as a function of the proportion of the face revealed. However, in nearly all cases, there was a striking discrepancy between responses to bubbles compared to whole faces: responses to bubble trials were not at all predictive of responses to WFs, even when substantial portions of the face or its features were revealed. We next quantified this observation further.

The population average of all single-trial responses of all units that increase their rate for WFs (32 units) showed a highly significant negative relationship with the amount of the eye and mouth revealed in the bubble trials (Figure 4C). This negative relationship was statistically robust across all trials as well as units, as assessed by a bootstrap statistic (mean slope -0.18 ± 0.05; see Figure 4C for details). Although this result was based on normalized firing rates, an even more significant negative relationship was found when considering absolute firing rates (Figure S3D) or the proportion of the whole face revealed (Figure S3E). The slope of the curve became more negative as the partial face became more similar to the WF (Figure 4C). Moreover, the same pattern, but with opposite sign, was found for the population average of all units that decreased their spike rate to WFs (n = 4): these units increased their spike rate with greater proportion of the face or ROI revealed (Figure S3A). Thus, in both cases, the population average of neurons that were WF selective (as defined by the initial contrast between WF and cutouts) showed a strong

and statistically significant relationship with the amount of the face that was shown, despite a complete failure to predict the response to whole faces. For neurons that were not WF selective to begin with, there was no systematic effect in response to the proportion of the face revealed—the slope was not significantly different from zero (Figure 4D). However, even for these non-WF-selective neurons, there was still a surprising difference between full-face and bubble trials, indicating that some of the non-WF units remain sensitive to WFs to some degree (also see Figure 3C). None of the above effects could be explained by mere differences in visibility or contrast (quantified by the contrast threshold) between the bubble and WF trials (Figure 4E; Figure S3C).

Might the above effect somehow result from the fact that the majority of bubble trials only revealed a small proportion of the face? We tested this possibility in one patient by disabling the dynamic change in bubbles and showing a relatively fixed and large number of bubbles (~100), revealing a large proportion of the face on all trials (Figure S3F shows examples; typically >90% of the eyes and mouth ROI was revealed). In this patient, we found 5 out of 21 units (24%) that were WF selective (see Figure S3H for an example), and these neurons showed a similar nonlinear response profile (Figure S3G). The average WFI for the WF units and the entire population was 130% ± 14% and 37% ± 8%, respectively. Once again, we found that neurons that selectively respond to
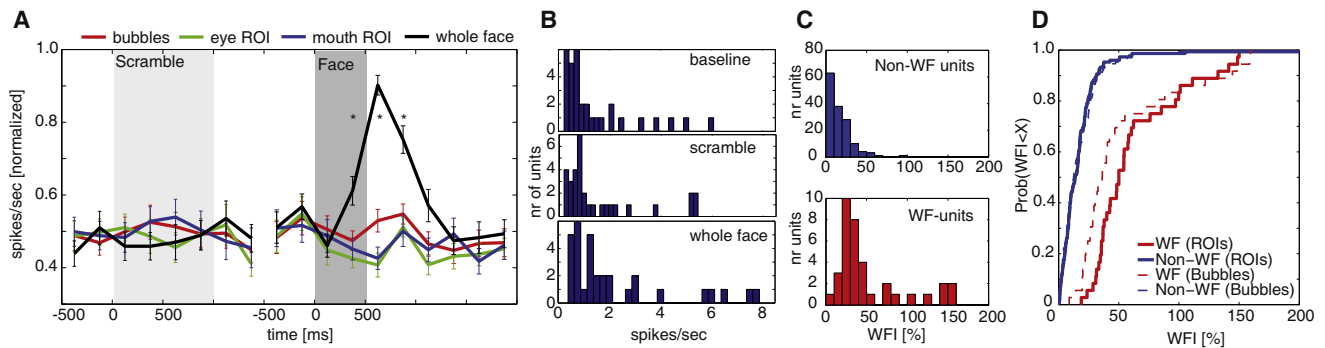
Figure 3. Whole-Face-Selective Neurons

(A) Mean response of all whole-face (WF)-selective units that increased their spike rate for WFs compared to bubble trials (n = 32 units, ±SEM, normalized to average response to WFs for each unit separately). Asterisk indicates a significant difference between the response to WFs and bubble trials (p < 0.05, two-tailed t test, Bonferroni-corrected for 14 comparisons). The response to eye and mouth regions of interest (ROIs) that was used for selecting the units is shown but not used for statistics.

(B) Histogram of firing rates before and after scramble onset as well as for WFs (n = 32 units). Mean rates were 1.4 ± 0.24 Hz, 1.7 ± 0.3 Hz, and 2.2 ± 0.4 Hz, respectively.

(C) Histogram of the whole-face index (WFI) for all recorded units (n = 185), according to whether the unit was classified as a non-WF-selective (top) or WF-selective (bottom) unit. The WFI was calculated as the baseline-normalized difference in response to whole faces compared to bubbles (which was independent of how we classified units as WF selective or non-WF selective).

(D) Distributions (plotted as cumulative distributions) of the WFI across the entire population for both WF- and non-WF-selective units (n = 36 and n = 149, respectively), calculated for both ROI trials (bold lines) and bubble trials (dashed lines). The two WFI populations for bubble trials were significantly different (p < 1e-9, two-tailed Kolmogorov-Smirnov test). Note the similarity of the distributions for cutouts (bold lines) and bubble trials (dashed), indicating that the response to both is very similar.

WFs failed to respond to parts of the face, in this case even when almost all of the face was revealed.

Could a difference in eye movements contribute to the responses we observed? This issue is pertinent, given that the human amygdala is critical for eye movements directed toward salient features of faces: lesions of the human amygdala abolish the normal fixations onto the eye region of faces [13]. Although we did not record eye movements in the present study as a result of technical constraints, we measured eye movements in a separate sample of 30 healthy participants (see Supplemental Experimental Procedures) in the same task. The mean and variance of the fixation patterns along the x and y axes did not differ between whole and bubbled faces (p > 0.20, two-tailed paired sign test). Similarly, in a previous study we found that fixation times on eyes and mouth in WFs and bubbled faces did not differ [14].

Finally, to examine the possible effects of recording from neurons that were in seizure-related tissue, we recalculated all analyses excluding any neurons within regions that were later determined to be within the epileptic focus. After excluding all units from that hemisphere in which seizures originated (see Table S1), a total of 179 units remained. Of those, 157 had the minimal required firing rate of 0.2 Hz, and 32 of those units (20%) were WF-selective units. Using only those units, all results remained qualitatively the same. It is also worth noting that, with one exception, all the patients with a temporal origin of seizures had their seizure foci in the hippocampus rather than the amygdala (Table S1), further making it unlikely that the inclusion of neurons within seizure-related tissue might have biased our findings.

Discussion

Recording from single neurons in the amygdalae of seven neurosurgical patients, we found that over half of all neurons responded to face stimuli (compared to only 10% of neurons responding to phase-scrambled faces), and a substantial

proportion of these showed responses selective for whole faces as compared to pieces of faces (WF-selective). Also, most neurons (31 of 36) did not show any association with reaction time, arguing that the majority of WF-selective neurons in the amygdala are driven by the sensory properties of whole faces rather than decisions or actions based on them. The earliest responses to WFs occurred within 250–500 ms after stimulus onset (Figure 3A). WF-selective neurons showed a highly nonlinear response, such that their response to WFs was inversely correlated with their response to variable amounts of the face or its features (eyes or mouth) that were revealed. Neurons that decreased their response as a function of the amount of the face revealed increased their response to WFs (Figure 4C). In contrast, neurons that increased their response as a function of the amount of the face revealed decreased their response to WFs (Figure S3A). In both cases, neurons showed the greatest difference in response between WFs and pieces of faces when facial features shown were actually the most similar between the two types of stimulus categories. Thus, the response to partially revealed faces was not predictive of how the unit would respond to WFs. These findings provide strong support for the conclusion that amygdala neurons encode holistic information about WFs, rather than about their constituent features.

We identified WF-selective neurons based on comparisons with the eye and mouth cutout trials. Because the remainder of the analysis was based on responses of these neurons in the bubble trials, the selection and subsequent analysis are statistically independent. This also allows later comparison of the response to the cutouts with the bubble trials (Figures 4C and 4D), which reveals that the cutout responses (unlike the WF responses) are consistent with what the bubble trials predict.

Our subjects performed an emotion categorization task, but amygdala responses to faces have been observed also in a variety of other tasks [2, 4, 6, 7, 11]. Also, classification images for face identification tasks are very similar to those
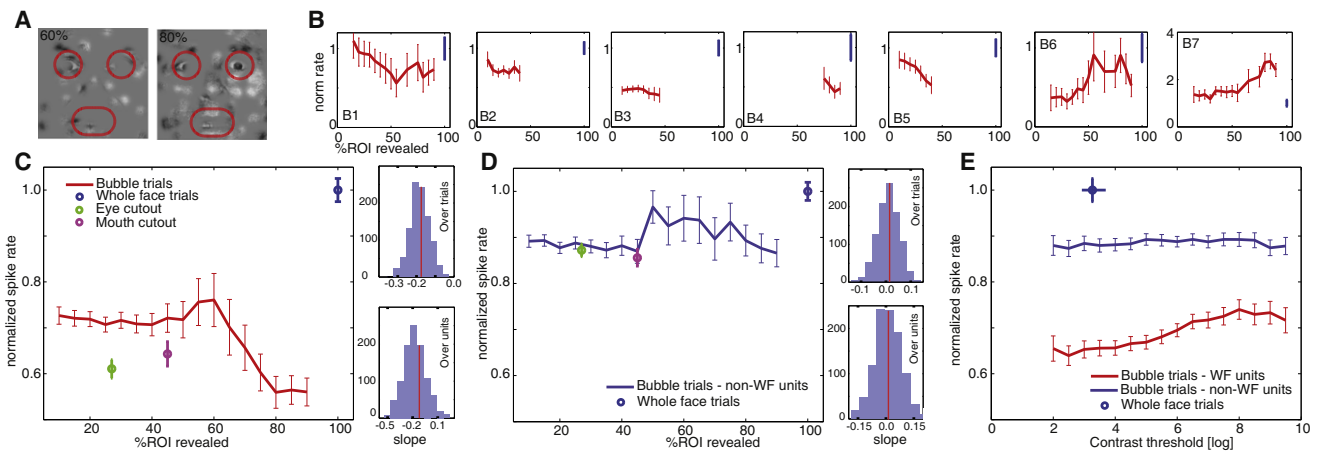
Figure 4. Response Profiles of Whole-Face-Selective Neurons

The response to partially revealed faces did not predict the response to whole faces.

(A) Example stimuli for 60% and 80% of the ROIs revealed.

(B1–B7) Example neurons. Normalized response as a function of how much of the eye and mouth ROIs were revealed (B1–B6 increase rate for WFs, B7 decreases) is shown. Only data points to which at least ten trials contributed are shown (bin width is 15%, steps of 5%). Most neurons showed a nonlinear response profile when comparing bubbles and WFs: those that decreased their response with more bubbles increased it to WFs and conversely. The slope of the regression of percentage of ROI on spike rate was significant for units B1, B5, and B6. Different portions of the x axis are plotted for different neurons, because the patients were shown different densities of bubbles contingent on their different performance accuracies (unit B4 is from the patient who was shown a very high density of bubbles for comparison; cf. Figure S3 for further details).

(C) Population responses of all WF-selective neurons that increase their response to WFs (n = 32 units, 5,686 trials). The slope was −0.18 and significantly negative (linear regression, p < 1e-11) and the last and first data points are significantly different (two-sided t test, p < 1e-5). We verified the statistical significance of the slope across the units that contributed to this population response using a bootstrap statistic over trials and units (right side, average slope −0.18 ± 0.05 and −0.17 ± 0.13, respectively; red line indicates observed value). For the first half (5%–40%) the slope was −0.07 (p = 0.27, not significant [NS]), and for the second half (45%–90%) it was −0.44 (p < 1e-8).

(D) Population response for non-WF-selective units (n = 149, 15,922 trials). The slope was not significantly different from zero for the curve shown (p = 0.62) and across the population over trials and units (bootstrap statistic; right side, average slope 0.01 ± 0.04 and 0.01 ± 0.07, respectively). The last and first data points were not significantly different.

(E) Response for WF-selective (red) and non-WF-selective units (blue) is not a function of contrast threshold (the contrast threshold is a model-derived index of the visibility of the stimuli). The response of the control trials for eye and mouth cutouts used for selecting the units are shown in (C) and (D) for comparison. Errors are ±SEM over trials. See also Figure S3.

we obtained using our emotion discrimination task [15] (cf. Figure 1C). This makes it plausible that WF-selective units would be observed regardless of the precise nature of the task requirements. We emphasize the distinction between responsive and selective neurons in our study—although about 50% of neurons responded to facial stimuli (compared to scrambles), this does not make them face selective because they might also respond to a variety of nonface stimuli (which were not shown in our study). Thus, the WF-selective units we found were selective for WFs compared to face parts, but their response to nonface stimuli remains unknown.

We analyzed the responses to the bubble trials by plotting neuronal responses as a function of the amount of the eye and mouth features revealed in these trials (Figure 4), as well as plotting the proportion of the entire face revealed (Figure S3E). The two measures (percentage of ROI revealed, percentage of entire face revealed) were positively correlated across trials (on average r = 0.46, p < 0.001; Figure S1C), because bubbles were independently and uniformly distributed over the entire image and the average number of bubbles (typically converging to around 20 during a session) was sufficiently high to make clustering of all bubbles on one ROI unlikely. As expected, the response as a function of the proportion of the entire face revealed (Figure S3E) thus shows a similar relationship at the population level. We used percentage of ROI for our primary analysis because it offered a metric with greater range, due to variability in the spatial location of the bubbles.

Facial Features Represented in the Amygdala

Building on theoretical models [16] as well as findings from responses to faces in temporal neocortex that provides input to the amygdala [17, 18], several studies have asked what aspects of faces might be represented in the amygdala. Various reports have demonstrated that the amygdala encodes information both about the identity of an individual's face, as well as about the social meaning of the face, such as its emotional expression or perceived trustworthiness [2, 4, 6, 7, 11]. Patients with amygdala lesions exhibit facial processing deficits for a variety of different facial expressions, including both fearful and happy [1, 13], and we found that most WF-selective units do not distinguish between fearful and happy, suggesting that the amygdala is concerned with a more general or abstract aspect of face processing than an exclusive focus on expressions of fear. At which stage of information processing does the amygdala participate? Because the amygdala receives highly processed visual information from temporal neocortex [19], one view is that it contains viewpoint-invariant [20], holistic [21] representations of faces synthesized through its inputs. Such global face representations could then be associated with the valence and social meaning of the face [22, 23] in order to modulate emotional responses and social behavior. This possibility is supported by blood oxygen level-dependent (BOLD) functional magnetic resonance imaging (fMRI) activations within the amygdala to a broad range of face stimuli (e.g., [3, 5]). An alternative possibility is motivated by the finding that the amygdala, at

least in humans, appears to be remarkably specialized for processing a single feature within faces: the region around the eyes. For instance, lesions of the amygdala selectively impair processing information from the eye region in order to judge facial emotion [13], and BOLD-fMRI studies reveal amygdala activation during attention to the eyes in faces [24] and to isolated presentation of the eye region [25, 26]. These opposing findings suggest two conflicting views of the role of the amygdala during face processing. Our results generally support the first possibility.

Face Responses in the Primate Amygdala

The amygdala receives most of its visual inputs from visually responsive temporal neocortex [19, 27], and there is direct evidence from electrical microstimulation of functional connections between face-selective patches of temporal cortex and the lateral amygdala in monkeys [28]. Although there is ongoing debate regarding a possible subcortical route of visual input to the amygdala that might bypass visual cortices [29], both the long response latencies and WF selectivity of the neurons we report suggest a predominant input via cortical processing. This then raises a core question: What is the transformation of face representations in the amygdala, relative to its cortical inputs?

The regions of temporal cortex that likely convey visual information about faces to the amygdala themselves show remarkable selectivity to faces [30–32] and to particular identities [33, 34] and emotions [17, 35] of faces. Regions providing likely input to the amygdala [28] are known to contain a high proportion (>80%) of face-selective cells and have highly viewpoint-invariant responses to specific face identities [20]. In humans, studies using BOLD-fMRI have demonstrated between 3 and 5 regions of cortex in the occipital, temporal, and frontal lobes that show selective activation to faces and that appear to range in encoding parts of faces, identities of faces, or changeable aspects of faces such as emotional expressions [18, 36]. Intracranial recordings in humans have observed electrophysiological responses selective for faces in the anterior temporal cortex [37, 38]. However, although there is thus overwhelming evidence for neurons that respond to faces rather than to other stimulus categories, many temporal regions also respond to specific parts or features of faces to some extent [32, 39, 40]. In contrast, the highly nonlinear face responses we observed in the amygdala have not been reported.

Single-unit responses in the monkey amygdala have described responses selective for faces [41], with cells showing selectivity for specific face identities and facial expressions of emotion [7, 42, 43] as well as head and gaze direction [44]. Interestingly, the proportion of face-responsive cells in the monkey amygdala has been reported to be approximately 50% [7, 42], similar to what we found in our patients. Similarly, cells recorded in the human anteromedial temporal lobe including the amygdala have been reported to exhibit highly specific and viewpoint-invariant responses to familiar faces [11, 45], as well as selectivity for both the identity and emotional expression of faces [6]. The present findings are consistent with the idea that there is a convergence of tuning to facial features toward more anterior sectors of the temporal lobe, culminating in neurons with responses highly selective to WFs as we found in the amygdala. The nonlinear face responses we describe here may indicate an architecture involving both summation and inhibition in order to synthesize highly selective face representations. The need to do so in the

amygdala likely reflects this structure's known role in social behavior and associative emotional memory: in order to track exactly which people are friend or foe, the associations between value and face identity must be extremely selective in order to avoid confusions between different people.

It remains an important question to understand how the face representations in the amygdala are used by other brain regions receiving amygdala input. It is possible that aspects of temporal cortical face responses depend on recurrent inputs from the amygdala, because the face selectivity of neurons in temporal regions that are functionally connected with the amygdala (such as the anterior medial face patch) evolves over time and peaks with a long latency of >300 ms [20], and because temporal cortex can signal information about emotional expression at later points in time than face categorization as such [35]. Similarly, visually responsive human amygdala neurons respond with a long latency of on average around 400 ms [46]. Such a role for amygdala modulation of temporal visual cortex is also supported by BOLD-fMRI studies in humans that have compared signals to faces in patients with lesions to the amygdala [47].

In conclusion, our findings demonstrate that the human amygdala contains a high proportion of face-responsive neurons. Most of those that show some kind of selectivity are selective for presentations of the entire face and show surprising sensitivity to the deletion of even small components of the face. Responses selective for whole faces are more prevalent than responses selective for face features, and responses to whole faces cannot be predicted from parametric variations in the features. Taken together, these observations argue that the face representations in the human amygdala encode socially relevant information, such as identity of a person based on the entire face, rather than information about specific features such as the eyes.

References

1. Adolphs, R., Tranel, D., Damasio, H., and Damasio, A. (1994). Impaired recognition of emotion in facial expressions following bilateral damage to the human amygdala. Nature *372*, 669–672.
2. Winston, J.S., Strange, B.A., O'Doherty, J., and Dolan, R.J. (2002). Automatic and intentional brain responses during evaluation of trustworthiness of faces. Nat. Neurosci. *5*, 277–283.
3. Fitzgerald, D.A., Angstadt, M., Jelsone, L.M., Nathan, P.J., and Phan, K.L. (2006). Beyond threat: amygdala reactivity across multiple expressions of facial affect. Neuroimage *30*, 1441–1448.

4. Hoffman, K.L., Gothard, K.M., Schmid, M.C., and Logothetis, N.K. (2007). Facial-expression and gaze-selective responses in the monkey amygdala. Curr. Biol. *17*, 766–772.

5. Hariri, A.R., Tessitore, A., Mattay, V.S., Fera, F., and Weinberger, D.R. (2002). The amygdala response to emotional stimuli: a comparison of faces and scenes. Neuroimage *17*, 317–323.

6. Fried, I., MacDonald, K.A., and Wilson, C.L. (1997). Single neuron activity in human hippocampus and amygdala during recognition of faces and objects. Neuron *18*, 753–765.

7. Gothard, K.M., Battaglia, F.P., Erickson, C.A., Spitler, K.M., and Amaral, D.G. (2007). Neural responses to facial expression and face identity in the monkey amygdala. J. Neurophysiol. *97*, 1671–1683.

8. Rutishauser, U., Ross, I.B., Mamelak, A.N., and Schuman, E.M. (2010). Human memory strength is predicted by theta-frequency phase-locking of single neurons. Nature *464*, 903–907.

9. Gosselin, F., and Schyns, P.G. (2001). Bubbles: a technique to reveal the use of information in recognition tasks. Vision Res. *41*, 2261–2271.

10. Steinmetz, P.N., Cabrales, E., Wilson, M.S., Baker, C.P., Thorp, C.K., Smith, K.A., and Treiman, D.M. (2011). Neurons in the human hippocampus and amygdala respond to both low and high level image properties. J. Neurophysiol. *6*, 2874–2884.

11. Kreiman, G., Koch, C., and Fried, I. (2000). Category-specific visual responses of single neurons in the human medial temporal lobe. Nat. Neurosci. *3*, 946–953.

12. Mosher, C.P., Zimmerman, P.E., and Gothard, K.M. (2010). Response characteristics of basolateral and centromedial neurons in the primate amygdala. J. Neurosci. *30*, 16197–16207.

13. Adolphs, R., Gosselin, F., Buchanan, T.W., Tranel, D., Schyns, P., and Damasio, A.R. (2005). A mechanism for impaired fear recognition after amygdala damage. Nature *433*, 68–72.

14. Neumann, D., Spezio, M.L., Piven, J., and Adolphs, R. (2006). Looking you in the mouth: abnormal gaze in autism resulting from impaired top-down modulation of visual attention. Soc. Cogn. Affect. Neurosci. *1*, 194–202.

15. Schyns, P.G., Bonnar, L., and Gosselin, F. (2002). Show me the features! Understanding recognition from the use of visual information. Psychol. Sci. *13*, 402–409.

16. Bruce, V., and Young, A. (1986). Understanding face recognition. Br. J. Psychol. *77*, 305–327.

17. Hasselmo, M.E., Rolls, E.T., and Baylis, G.C. (1989). The role of expression and identity in the face-selective responses of neurons in the temporal visual cortex of the monkey. Behav. Brain Res. *32*, 203–218.

18. Haxby, J.V., Hoffman, E.A., and Gobbini, M.I. (2000). The distributed human neural system for face perception. Trends Cogn. Sci. (Regul. Ed.) *4*, 223–233.

19. Amaral, D.G., Price, J.L., Pitkanen, A., and Carmichael, S.T. (1992). Anatomical organization of the primate amygdaloid complex. In The Amygdala: Neurobiological Aspects of Emotion, Memory, and Mental Dysfunction, J.P. Aggleton, ed. (New York: Wiley-Liss), pp. 1–66.

20. Freiwald, W.A., and Tsao, D.Y. (2010). Functional compartmentalization and viewpoint generalization within the macaque face-processing system. Science *330*, 845–851.

21. Andrews, T.J., Davies-Thompson, J., Kingstone, A., and Young, A.W. (2010). Internal and external features of the face are represented holistically in face-selective regions of visual cortex. J. Neurosci. *30*, 3544–3552.

22. Adolphs, R. (2010). What does the amygdala contribute to social cognition? Ann. N Y Acad. Sci. *1191*, 42–61.

23. Todorov, A., and Engell, A.D. (2008). The role of the amygdala in implicit evaluation of emotionally neutral faces. Soc. Cogn. Affect. Neurosci. *3*, 303–312.

24. Gamer, M., and Büchel, C. (2009). Amygdala activation predicts gaze toward fearful eyes. J. Neurosci. *29*, 9123–9126.

25. Morris, J.S., deBonis, M., and Dolan, R.J. (2002). Human amygdala responses to fearful eyes. Neuroimage *17*, 214–222.

26. Whalen, P.J., Kagan, J., Cook, R.G., Davis, F.C., Kim, H., Polis, S., McLaren, D.G., Somerville, L.H., McLean, A.A., Maxwell, J.S., and Johnstone, T. (2004). Human amygdala responsivity to masked fearful eye whites. Science *306*, 2061.

27. Iwai, E., and Yukie, M. (1987). Amygdalofugal and amygdalopetal connections with modality-specific visual cortical areas in macaques (Macaca fuscata, M. mulatta, and M. fascicularis). J. Comp. Neurol. *261*, 362–387.

28. Moeller, S., Freiwald, W.A., and Tsao, D.Y. (2008). Patches with links: a unified system for processing faces in the macaque temporal lobe. Science *320*, 1355–1359.

29. Pessoa, L., and Adolphs, R. (2010). Emotion processing and the amygdala: from a 'low road' to 'many roads' of evaluating biological significance. Nat. Rev. Neurosci. *11*, 773–783.

30. Bruce, C., Desimone, R., and Gross, C.G. (1981). Visual properties of neurons in a polysensory area in superior temporal sulcus of the macaque. J. Neurophysiol. *46*, 369–384.

31. Perrett, D.I., Rolls, E.T., and Caan, W. (1982). Visual neurones responsive to faces in the monkey temporal cortex. Exp. Brain Res. *47*, 329–342.

32. Tanaka, K., Saito, H., Fukada, Y., and Moriya, M. (1991). Coding visual images of objects in the inferotemporal cortex of the macaque monkey. J. Neurophysiol. *66*, 170–189.

33. Rolls, E.T. (2000). Functions of the primate temporal lobe cortical visual areas in invariant visual object and face recognition. Neuron *27*, 205–218.

34. Eifuku, S., De Souza, W.C., Tamura, R., Nishijo, H., and Ono, T. (2004). Neuronal correlates of face identification in the monkey anterior temporal cortical areas. J. Neurophysiol. *91*, 358–371.

35. Sugase, Y., Yamane, S., Ueno, S., and Kawano, K. (1999). Global and fine information coded by single neurons in the temporal visual cortex. Nature *400*, 869–873.

36. Tsao, D.Y., Moeller, S., and Freiwald, W.A. (2008). Comparing face patch systems in macaques and humans. Proc. Natl. Acad. Sci. USA *105*, 19514–19519.

37. Allison, T., McCarthy, G., Nobre, A., Puce, A., and Belger, A. (1994). Human extrastriate visual cortex and the perception of faces, words, numbers, and colors. Cereb. Cortex *4*, 544–554.

38. Allison, T., Puce, A., Spencer, D.D., and McCarthy, G. (1999). Electrophysiological studies of human face perception. I: Potentials generated in occipitotemporal cortex by face and non-face stimuli. Cereb. Cortex *9*, 415–430.

39. Fujita, I., Tanaka, K., Ito, M., and Cheng, K. (1992). Columns for visual features of objects in monkey inferotemporal cortex. Nature *360*, 343–346.

40. Freiwald, W.A., Tsao, D.Y., and Livingstone, M.S. (2009). A face feature space in the macaque temporal lobe. Nat. Neurosci. *12*, 1187–1196.

41. Leonard, C.M., Rolls, E.T., Wilson, F.A., and Baylis, G.C. (1985). Neurons in the amygdala of the monkey with responses selective for faces. Behav. Brain Res. *15*, 159–176.

42. Nakamura, K., Mikami, A., and Kubota, K. (1992). Activity of single neurons in the monkey amygdala during performance of a visual discrimination task. J. Neurophysiol. *67*, 1447–1463.

43. Kuraoka, K., and Nakamura, K. (2006). Impacts of facial identity and type of emotion on responses of amygdala neurons. Neuroreport *17*, 9–12.

44. Tazumi, T., Hori, E., Maior, R.S., Ono, T., and Nishijo, H. (2010). Neural correlates to seen gaze-direction and head orientation in the macaque monkey amygdala. Neuroscience *169*, 287–301.

45. Quiroga, R.Q., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005). Invariant visual representation by single neurons in the human brain. Nature *435*, 1102–1107.

46. Mormann, F., Kornblith, S., Quiroga, R.Q., Kraskov, A., Cerf, M., Fried, I., and Koch, C. (2008). Latency and selectivity of single neurons indicate hierarchical processing in the human medial temporal lobe. J. Neurosci. *28*, 8865–8872.

47. Vuilleumier, P., Richardson, M.P., Armony, J.L., Driver, J., and Dolan, R.J. (2004). Distant influences of amygdala lesion on visual cortical activation during emotional face processing. Nat. Neurosci. *7*, 1271–1278.