

4 Persistent Changes in Visual Cortex Due to Aversive Conditioning

This work was done with John P. O’Doherty and Christof Koch. John O’Doherty was, as always, indispensable in the analysis design and many discussions of data interpretation. Christof Koch worked on the goal and initial design of this experiment as well as review of different analysis phases. His talents in analysis were greatly appreciated during modeling and analysis of skin conductance responses. He also maintained the thankless task of insuring that funding was in place. Imaging acquisition was completed at the Caltech Biological Imaging Center. Skin conductance recording during fMRI acquisition was possible because of extensive testing and design done in collaboration with Antoine Bruguier (see Chapter 6).

We examined changes in the representation of two face images over the course of an aversive conditioning experiment and subsequent extinction.

4.1 Introduction

Conditioning is commonly used to study learning and memory in a wide range of organisms (Baer and Fuhrer, 1982; Mackintosh, 1983; Gallistel, 1990; Thompson and Krupa, 1994; Connolly et al., 1996; Eichenbaum, 1997; Pearce et al., 1997; Tully, 1998; Squire and Kandel, 1999; Kocorowski and Helmstetter, 2001). Milner’s study of HM (Scoville and Milner, 1957) provided evidence for the separation of explicit and implicit learning, since HM was capable of non-conscious, implicit learning but not of retaining explicit memories. Most discussions of implicit learning in conditioning center on

modifications in the amygdala and thalamic structures (Morris et al., 1999). However, the related field of perceptual learning demonstrates that modifications to the cortical sensory (in this case visual) pathway can and do take place (Poggio et al., 1992; Yang et al., 1994; Goldstone, 1998; Watanabe et al., 2002). The same is also true of the auditory pathway (Fritz et al., 2003). Perceptual learning is a change in the way a subject perceives a stimulus that is believed to reflect changes in sensory cortex. Subjects in perceptual learning studies show phenomena ranging from hyperacuity (Poggio et al., 1992), an increase in effective visual resolution, to the representation of previously unknown stimulus sets (Gauthier and Tarr, 1997; Gauthier et al., 2000). While these changes may occur with explicit influence, the resulting differences are of the automatic, implicit, variety. Modification of auditory cortex due to conditioning is well described in both electrophysiology and fMRI studies (Shamma, 2004; Weinberger, 2004). In these studies, a particular auditory stimulus which is of neutral value, the conditioned stimulus (CS), is reinforced with a reward or punishment, the unconditioned stimulus (US). The representation of this particular CS shows an increase in representation in auditory cortex as a result of conditioning (Kisley and Gerstein, 2001). Correspondingly, individual cells show a shift in their response curve toward the frequency of the auditory stimulus used (Edeline et al., 1993). It seems likely that processes similar to those that take place during auditory conditioning may also cause modification of the visual pathway.

There are some fMRI studies that report differences in visual cortex during conditioning (Knight et al., 1999; Carter et al., 2006), but these studies were not directed at examining stimulus representation differences and also don't examine the persistence of such changes. Recently, a study described changes in rat visual cortex due to

conditioning (Shuler and Bear, 2006). In particular, they report changes in the timing of neurons in visual cortex that reflected reward timing. Indicating that modifications to at least the timing of neural responses in visual areas can and do take place as a result of conditioning.

These studies describe changes that result from the acquisition of conditioning. If the changes in cortical response profiles described above are lasting and not due to continued modulation by an associative area of the brain, these changes should persist after the extinction of conditioned responses.

Here we describe a study of changes in a specific cortical visual area as a result of aversive conditioning in humans. We used a simple delay protocol to condition subjects to faces and abstract images during fMRI acquisition. We posited that those stimuli that were paired with a shock would elicit hemodynamic activity that increased over time when compared to a similar unpaired stimulus. We also hypothesized that this differential response in cortical visual areas would persist beyond extinction of the conditioned response.

BOLD activity in response to a reinforced face stimulus in the fusiform face area increases over time when compared to activity elicited by an unpaired face. We find that activity in this area remains consistently elevated in spite of extinction of the conditioned response, indicating that potential changes to perception of the previously reinforced stimulus are likely to be long lasting.

4.2 Methods

4.2.1 Participants

From an initial pool of seventeen subjects who participated in the experiment, our analysis includes twelve subjects who showed some indication of conditioning (greater SCRs for the CS+ than CS- at $p < 0.2$ for both faces and abstract images, see below). Five of the original seventeen were not included in the final analysis because they either did not condition (2) or because of technical failures with the SCR recording (3). Subjects included in the analysis: six male and six female, age range 19-31 (mean 24.25). All subjects gave prior informed consent. This study was approved by the Human Subject Committee at the California Institute of Technology.

4.2.2 Experimental Procedure

Prior to entering the magnet suite, subjects were told the experiment was a learning and memory experiment. They were told that they would be presented with images and some of the images would be paired with a shock while others would not. They were asked to keep their eyes on a fixation mark and pay attention to the images presented. They were not asked to perform any other task. After confirming that they understood the instructions, the skin conductance and shock electrodes were attached and their shock level was determined. The subject was then positioned in the magnet. An anatomical scan was acquired first, followed by a 5 minute retinotopic scan using a rotating wedge stimulus. At this point, instructions for the main experiment were confirmed with the subject, and the acquisition portion of the experiment was performed.

Acquisition consisted of 60 presentations of each of the four different images (240 in total) with a 75% reinforcement rate for the CS+, and no shock pairings for the CS-. Intertrial intervals ranged from 7-9 seconds. Stimuli were presented so that no more than two of identical images occurred in a row and that nonreinforced test trials were spread evenly throughout acquisition. The stimulus that served as the CS+ was chosen at random. For subjects included in the analysis, each abstract image was reinforced for six of the twelve participants. For face stimuli, the dark haired individual was reinforced for seven of the twelve subjects and the image of the individual with lighter hair for the remaining five subjects.

Post acquisition, subjects were told they would be presented with the same images but would no longer receive any shocks (that the shock electrodes had been unplugged). They were reminded to keep their eyes on the fixation mark and to pay attention to each image presented. Each CS was then presented 20 times for 1 second with an intertrial interval of 6 to 8 seconds as part of an extinction phase. Stimuli were presented so that no more than two of a particular stimulus occurred in a row.

4.2.3 Stimulus Presentation

Face images (Figure 4-1 top left) were obtained from the Ekman collection. Abstract images (Figure 4-1 top right) were adapted from a previous study (Carter et al., 2006). All images were presented at a size of approximately four degrees wide by six degrees tall for a period of one second. For reinforced stimuli, a shock was presented for one second overlapping with the last half second of the image presentation (see Figure 4-1 bottom). Images were presented on a grey background approximately 20 degrees in width

and 15 degrees in height at refresh frequency of 60Hz using LCD goggles from Resonance Technology Inc. (Los Angeles, CA).

Eight millimeter diameter Ag/AgCl electrodes containing a conductive paste (MedAssociates TD-246, 0.5M NaCl suspension) were attached to the top of the subject's right foot for stimulation. The shock stimulus was 60Hz alternating current of constant amplitude. Current levels for shock stimulation were chosen by each subject before the experiment, using an ascending rating method where the current amplitude was raised until the subject gave a rating of 9 on a scale of 1-10. A rating of 1 indicated the subject could barely feel the shock and 10 indicated the shock was painful and could not be used in the experiment. Subject shock levels ranged from 1.2 to 4.8mA with a median of 1.8mA.

4.2.4 Skin Conductance Conditioning

Skin conductance was recorded throughout the acquisition and extinction phases of learning. Eight millimeter diameter Ag/AgCl electrodes containing a conductive paste (MedAssociates TD-246, 0.5M NaCl suspension) were attached to the arch of the subject's left foot (Boucsein, 1992). These electrodes were attached to carbon fiber leads (Biopac, Goleta, CA) that ran through a wave guide into the control room, where they were filtered using the equipment described in Chapter 6. The signal was then recorded using equipment from Contact Precision Instruments (Boston, MA). Synchronization pulses were sent from the parallel port of the stimulation computer to the CPI standalone module. Special care was taken to ensure minimal radiofrequency induced contamination of the SCR signal (see Chapter 6).

Skin conductance responses (SCRs) were defined as the maximum amplitude response initiated no earlier than 1second with a peak no later than 5seconds after the CS onset. SCR amplitudes were range corrected by the maximum response for that subject (Lykken, 1972). Raw skin conductance data was filtered using a second-order Savitzky-Golay smoothing algorithm with a window of four seconds. This filtering removed any remaining radiofrequency noise in the skin conductance trace.

4.2.5 fMRI

Thirty-two slice EPI images were acquired at a resolution of 3 cubic millimeters using a Siemens Trio scanner and 8 channel phased array head coil (www.siemens.com). Slices were acquired in ascending interleaved order. These settings resulted in a TR of 2 seconds. Functional images were then preprocessed and analyzed using SPM5 (Wellcome Department of Imaging Neuroscience, Institute of Neurology, London, UK). Preprocessing included motion realignment, slice timing acquisition correction, normalization to the MNI template and smoothing to a final value of 8mm FWHM.

The functional analysis for acquisition modeled each of the four CS types separately. Parametric modulations of each of the CSs were used to test for changes over time and any effects due to the presence of the shock. In controlling for the presence or absence of the shock, the reinforcement schedule for the CS+ images was mirrored for CS- trials. Time changes were modeled using an exponential decay (see inset Figure 4-5) with a half life of $\frac{1}{4}$ the length of the experiment (Quirk et al., 1997; Buchel et al., 1998b). When looking for events that show an increase to plateau over time, we test for areas anticorrelated with this term (see inset Figure 4-6).

The functional analysis for extinction was the same as acquisition, except for the term to model differences between reinforced and nonreinforced trials (since no trials were reinforced).

4.3 Results

The primary focus of this study was to look at changes in the representation of stimuli as a result of conditioning. We first describe areas of the brain that over all trials show a greater activation to the CS+ than CS-. We then focus on conditioning to faces and specifically two areas of the brain, the amygdala and a functionally defined cluster in the fusiform region thought to represent faces (Allison et al., 1994; Haxby et al., 1994; Kanwisher et al., 1997). We focus on faces rather than abstract images because of the wealth of information concerning face representation.

4.3.1 Skin Conductance Conditioning

On average, subjects used in our analysis demonstrate a conditioned aversion to the CS+. We compared skin conductance responses (SCRs) on nonreinforced CS+ trials (there were 15) to matched CS- trials. Individual plots of each subject's average response to the nonreinforced CS+ and matched CS- trials are shown in left and middle portions of Figure 4-2. It is important that CS+ vs CS- comparisons occur in a pair-wise fashion. While each subject has a larger average response to the CS+ than CS- it is not always true that each subject's average CS+ response is larger than all other subject's average response to the CS-. The box plots in the right portion of Figure 4-2 show the distribution of average differences between nonreinforced CS+ and CS- trials. Subjects show a

greater average response to the CS+ than CS-, indicating they were conditioned (P<0.0001 for abstract and face images; single sample t-test)

4.3.2 Brain Activity during Acquisition

We compared BOLD activity during CS+ trials with CS- trials. Regions of the brain that exhibit stronger responses to the face CS+ compared to the face CS- include the insula, operculum, caudate, globus pallidus, inferior parietal lobule, thalamic nuclei including superior colliculus, cingulate cortex and areas of the cerebellum (see Figure 4-3 and Table 4-1). These results are similar to those described in past studies (Buchel et al., 1998b; Carter et al., 2006). Similar areas were active when comparing responses for the abstract CS+ > CS-.

Differences in Face and Abstract Image Representation

BOLD responses to the images that were never reinforced, the CS-, were compared for faces and abstract images. Contrasts were masked for positive responses to the image type of interest. This comparison identifies those areas that preferentially respond to one particular class of stimuli. It does so without the shock or association formation as potential confounds. There is, however, the possibility that the CS- serves as a safety signal and could therefore be considered rewarding (Seymour et al., 2005). Similar results are obtained if the CS+ trials are used.

Areas of the brain that preferentially respond to faces include the fusiform face area, superior temporal sulcus, and medial occipital cortex (Table 4-2a; also see the left half of Figure 4-4). Areas of the brain that preferentially respond to abstract figures

include the occipital pole and parahippocampus (Table 4-2b also see the right half of Figure 4-4).

BOLD interactions with time

Previous studies (Buchel et al., 1998b) find that activity in the amygdala fits an exponential decay function. Changes in BOLD activity over time were therefore modeled using an exponential decay function (see methods, Figure 4-5 inset) in an effort to reproduce this result. We tested for an exponential decay over CS+ trials when compared to CS- trials. In line with the study by Buchel et al., we find that BOLD activity in the amygdala correlates with an exponential decay function for conditioning to both faces and abstract images (Figure 4-5). For faces, the peaks of amygdala activity were located at MNI coordinates 24, -6, -21 (right) and -15, -6, -24 (left). For abstract images, the peaks of amygdala activity were located at 30, -6, -15 (right) and -21, -9, -15 (left). These amygdala peaks were all significant at $P < 0.05$ (family wise error corrected for multiple comparisons using a 10 millimeter diameter sphere around the peak voxel).

We hypothesized that reinforced stimuli would come to elicit stronger responses in areas that were capable of differentiating within that class of stimuli. We therefore tested for areas of the brain that showed a constrained increase in activity over time when compared to the CS- trials. The modeled increase is the inverted decay function of the previous section (Figure 4-6 inset). We specifically examined the areas identified in Table 4-2 (Figure 4-4) that preferentially represented each class of stimuli. We found that none of the areas preferring abstract figures correlated with an increase in activity over time (see discussion). However, the hypothesis held for faces. We found that activity in the fusiform area (peak at 42, -63, -15) showed increasing responses to CS+ trials over

time (see Figure 4-6; $P < 0.05$ corrected for small volumes using a 10 millimeter diameter sphere centered at the peak). This cluster of voxels ($P < 0.01$ uncorrected) fell within the face sensitive ROI ($P < 0.001$ uncorrected). Figure 4-8 shows the extent of overlap between the fusiform face area and those regions that exhibited larger responses to the CS+ over time. Increasing activity to the face CS+ is not apparent in the superior temporal sulcus or medial occipital cortex using the thresholds described above. Using lower thresholds, activity becomes apparent in both the superior temporal sulcus and right medial occipital lobe. However, these clusters are a sizeable distance away from the ROI peak in another anatomical region (> 20 millimeters) or do not survive correction for multiple comparisons. An increase in activity in the fusiform is consistent with an increase in representation of the CS+ image compared to the CS-.

4.3.3 Extinction

Following the acquisition of conditioning, subjects were told they would no longer receive any shocks. They were then shown 20 presentations of each CS with no shocks (extinction). Extinction of the previous association can be confirmed by the absence of differential skin conductance responses (SCRs). Testing of SCRs to both face and abstract images revealed no significant differences between CS+ and CS-, indicating an effective extinction of the previous association (see Figure 4-7a).

In contrast to the loss of differential skin conductance responses during extinction, the fusiform face area (peak at 36, -45, -24) remains more responsive to the CS+ than CS- (see Figure 4-7b; $P < 0.05$ small volume corrected for a 10 millimeter diameter sphere

centered at the peak). At a threshold of $P < 0.001$ uncorrected, it is clear from Figure 4-7b that the majority of the visual pathway shows a greater response to the face CS+.

The bottom section of Figure 4-7b shows activity in the geniculate nucleus. The left half is an SPM showing activity that is greater for the CS+ than CS- during extinction (peaks: left -21, -27, -6; right 24, -24, -3; $P < 0.05$ small volume corrected for a 10 mm diameter sphere centered at the peak). The right half displays the overlap with visual hemifield stimulation using a standard rotating wedge (Wunderlich et al., 2005), confirming the responsiveness of the area to visual activity in one hemifield in a separate session. We find it surprising that activity in an area so early in the visual stream responds to one face more strongly than another.

There is no activity in the fusiform face area correlated with a differential increase or decrease over extinction, as might be expected if activity in perceptual areas mirrored association strength.

Activity in the fusiform face area increasingly responds to the CS+ face during acquisition and retains that differential response after extinction of the conditioned association. Figure 4-8 displays the overlap of regions that preferentially respond to faces (red), increasingly respond to the CS+ during conditioning (yellow), and consistently respond more strongly to the CS+ face in spite of extinction of conditioning.

4.4 Discussion

We sought to provide evidence of modification of perceptual areas representing visual stimuli similar to established work in auditory conditioning (Bakin and Weinberger, 1990; Morris et al., 1998b; Bao et al., 2001). We used an aversive conditioning paradigm

to condition subjects to faces and abstract images. Examination of activity in the fusiform face area during conditioning to faces indicates that responses to the CS+ increase over time. Tests in the other face responsive areas listed in Table 4-2 do not provide strong evidence for increasing responses to the CS+ face (see Results). We therefore restrict our discussion to the FFA. In addition to acquisition, responses in the FFA were also examined during extinction phase, where conditioned stimuli are presented without reinforcement. Extinction is meant to form a new association that the previously reinforced stimuli are now safe. Larger responses to the CS+ trials persist through this separate extinction phase, indicating that changes in representation are present after the elimination of conditioned responses. We find no evidence of differential SCRs during extinction and also see little, if any, activity in areas associated with the conditioned response such as the amygdala or insula. This leads us to believe that there is little remaining affective component to the FFA activity in extinction.

With some differences noted due to task, familiarity, and emotion, repeated presentations of visual stimuli generally result in decreasing activity in visual areas (Henson et al., 2002; Ishai et al., 2004). A reduction in responses to known stimuli is consistent with the need to maintain a sparse representation, thereby allowing a greater number of stimuli to be represented. Results in auditory conditioning show both a shift in peak neuronal responses (Edeline et al., 1993) and a greater number of neurons responding to the conditioned stimulus even without a peak shift (Kisley and Gerstein, 2001). We know of no studies demonstrating persistent BOLD response changes for visual stimuli, although recent work in rats has shown a shift in timing (Shuler and Bear, 2006).

We find that as significance is attached to an image of a particular face, BOLD responses in the fusiform face area increase. The fusiform face area is posited to be an area of the brain with enhanced face representation (Puce et al., 1996; Kanwisher et al., 1997). The innateness and degree of this specificity are often questioned (Gauthier et al., 2000; Tarr and Gauthier, 2000), but, in general, this area of the brain does seem to be involved in the representation and identification of faces. It is plausible that an increase in activity in this portion of the brain to a specific face reflects a change in the likelihood that this face will be identified (an increase in saliency). Future experiments will test for behavioral correlates that could reflect saliency increases such as reduced search times when looking for CS+ faces. Another potential behavioral manifestation of increased representation by the FFA is that any given face could be more likely to be identified as the CS+ face. When shown an image of two faces mixed by morphing we would expect that previously conditioned subjects would have an equilibrium point, the point at which subjects are equally likely to identify the image as either individual, further from the CS+ face. That is, they would require a smaller CS+ identity percentage in order to describe the image as belonging to the CS+ individual.

One possible cause for the increased responsiveness in the FFA is changes in either top-down or bottom-up attentional processes. If the result is due to top-down attention, we would expect to see that activity in brain areas known to be involved in attention networks reflect the same changes we see in the FFA. In fact, during acquisition, bilateral BOLD activity in the intraparietal sulcus (IPS, right 21,-60, 36; left -27, -57, 45) also increases over time for the CS+ ($P < 0.05$ small volume corrected for a

10mm diameter sphere centered on the peak). The IPS is thought to be involved in the control of attention; for review, see Pessoa 2003 (Pessoa et al., 2003). However, there is no evidence that this activity in the IPS is persistent throughout extinction as is the activity in the FFA. This leads us to hypothesize that connectivity analysis would reveal interaction between the FFA and IPS during conditioning in order to cause perceptual changes. The lack of persistent activity in attentional or emotional networks during extinction leads us to believe post-extinction changes in the FFA represent a bottom-up or saliency process contained in the FFA.

Also of interest is that increased responses to the CS+ face during extinction are not limited to the FFA. Areas as early as the LGN (see Figure 4-7) show similar differences. Face recognition can occur in small groups of neurons such as the brain of the honey bee (Dyer et al., 2005). However, it is more likely that LGN activity in extinction is a result of top-down influences from the FFA, since it both preferentially responds to faces and shows a pattern of learning during acquisition that was not present in the LGN. Connectivity analysis in future experiments may also be useful in disentangling driving effects in this situation.

In contrast to the results for faces, we find no evidence of representation changes in visual areas that preferentially respond to abstract images. There are many potential causes for this. First, the images chosen may not fit into an image class represented in a specific brain area. This explanation does not seem highly likely given that there is significant overlap between areas shown to prefer abstract images when comparing CS- trials and areas shown to prefer abstract images when comparing CS+ trials. The abstract images do appear to show some overlap in representation. Second, since the abstract

images are novel stimuli, there may not be adequate representation of the stimuli in the brain for differentiation to be localized. Third, as in all imaging studies, it may be that this study lacks the sensitivity to pick out a change in representation that differs between CS+ and CS- for the abstract images. The attempt to represent the novel images may overwhelm any difference between the two.

We have shown that learned changes in face representation due to conditioning are persistent throughout extinction. This persistent activity may reflect changes in saliency. Future experiments should include an assessment of the relative saliency of the conditioned face as well as connectivity analysis that could disambiguate the top-down and bottom-up effects taking place.

4.5 Tables

Table 4-1

Cluster Region	MNI Coord.	Cluster size	P value
a) Faces CS+ > CS-			
Left Posterior Insula	(-36,-12,15)	25	P<0.001
Left Insula	(-39,0,6)	111	P<0.001
Left Medial Thalamus	(-6,-27,-9)	79	P=0.001
Right Medial Thalamus	(15,-12,12)	30	P<0.01
Right Posterior Insula	(39,-18,-3)	11	P<0.01
Right Insula	(39,-3,-12)	16	P<0.01
Left Inferior Parietal Lobule	(-66,-27,27)	10	P<0.01
Left Globus Pallidus	(-12,3,-6)	8	P<0.05
b) Abstract CS+ > CS-			
Left Thalamus	(-12,-12,0)	30	P<0.01
Left Posterior Insula	(-33,-21,12)	10	P<0.01
Left Inferior Parietal Lobule	(-60,-18,18)	11	P<0.01
Left Insula	(-33,-9,6)	22	P<0.01
Right Medial Thalamus	(3,-15,6)	25	P<0.01
Left Geniculate Nucleus	(-21,-21,-6)	5	P<0.05
Left Inferior Insula	(-33,-12,-12)	6	P<0.05
Right Inferior Insula	(-33,-21,12)	7	P<0.05

Table 4-1 Areas where brain activity is greater for CS+ trials than CS- trials for faces (a) and abstract figures (b). This table lists the region of activation followed by the MNI coordinates of its statistical peak, the number of voxels (3 mm^3) in the cluster and the peak whole brain corrected P value. Cluster size is those voxels with $P<0.05$ family wise error (FWE) corrected. Clusters containing 5 or less voxels are not listed (see also Figure 4-3).

Table 4-2

Cluster Region	MNI Coord.	Learning	Sustained
a) Faces CS- > Abs. CS-			
Right Superior Temporal Sulcus	(57,-45,0)	No	Yes
Right Superior Temporal Sulcus	(45,-42,15)	No	Yes
Left Medial Occipital Cortex	(-12,-54,-3)	No	No
Right Fusiform	(39,-48,-24)	No	Yes
Right Medial Occipital Cortex	(6,-78,0)	Yes	Yes
b) Abs. CS- > Face CS-			
Right Occipital Pole	(30,-90,9)	No	No
Left Parahippocampus	(-27,-60,-12)	No	No

Table 4-2 Areas of the brain preferentially responding to faces (a) or abstract (b) images (FWE corrected $P < 0.05$, comparing CS- trials). This table lists the region of activation followed by the MNI coordinates of its statistical peak, whether or not it was active in learning during acquisition (“Learning”) and if the region remained more responsive to the CS+ than CS- during extinction (“Sustained”). An ROI was considered active if statistically significant voxels ($P < 0.01$ uncorrected) were found within the ROI mask ($P < 0.001$ uncorrected) for the given CS type (faces > abstract or abstract > faces).

4.6 Figures

Figure 4-1

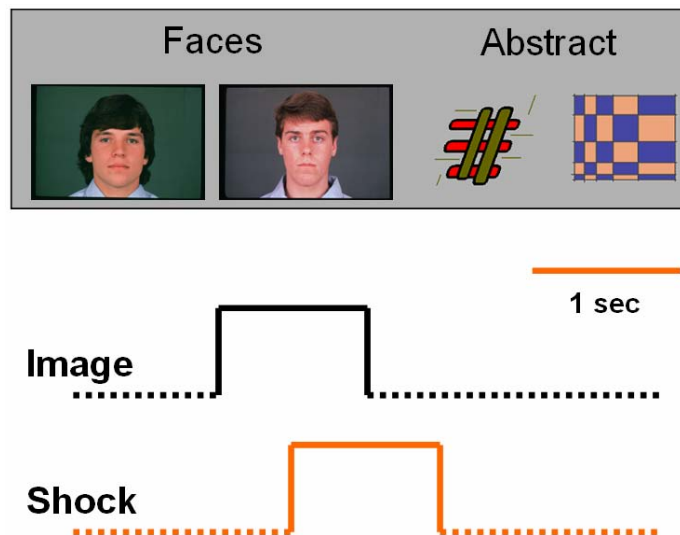


Figure 4-1 Experimental Setup – The top half of the figure shows the four images used as conditioned stimuli (CS). Each colored image was presented 60 times during acquisition and 20 times during extinction. One face and one abstract image were chosen randomly to be reinforced for each subject (the CS+). During acquisition, CS+ trials were reinforced on 75% of trials. The bottom half of the image shows the timing of stimulus presentation. Each image was shown for one second. On trials that were reinforced, a one second shock, the unconditioned stimulus (US) was given overlapping with the last half second of the image presented.


Figure 4-2

Figure 4-2 Skin conductance responses during acquisition show subjects learned to associate the CS+ stimulus with the shock for both faces and abstract images. The left two plots show average responses to nonreinforced CS+ trials and matched CS- trials for each subject. The leftmost plot is for face stimuli and the middle plot is for abstract stimuli. Each line represents a subject. Each subject shows that, on average, they had stronger responses to the CS+ than CS- and were therefore conditioned. The rightmost plot is a box and whisker plot of the average difference between CS+ and CS- stimuli for each subject, indicating that on average the group showed conditioning ($P < 0.0001$ for both abstract and face images). The median of each group is marked by a red line. The bottom and top edges of the box represent the edge of the lower and upper quartiles, respectively. The whiskers indicate the full range of data. Any outliers are indicated by a '+'.


Figure 4-3

Figure 4-3 Statistical Parametric Maps (SPMs) of those brain regions that have greater activity on CS+ trials than on CS- trials during conditioning to faces. Prominent clusters include the insula (Ins), operculum (Operc), globus pallidus (GP), caudate, inferior parietal lobule (IPL), thalamus, cingulate cortex, superior colliculus (SC) and cerebellum. These results are consistent with previous studies. SPMs are shown at a threshold of $P < 0.001$ uncorrected, overlaid on an average T1 structural image that has been normalized to the MNI template. Table 4-1a lists clusters of significant activity after whole brain correction for multiple comparisons.


Figure 4-4

Figure 4-4 Differences in representation of the CS- during acquisition (faces > abstract left; abstract > faces right; $P < 0.001$ uncorrected). Areas of the brain that respond more strongly to faces than abstract figures include the fusiform face area shown in the left image above. Abstract images elicit stronger responses than faces in extra-striate cortex as well as the para-hippocampal region. SPMs are shown overlaid on an average structural image that has been normalized to the MNI template. Glass brain images (inset) provide a means of visualizing any activity not in the slice displayed.


Figure 4-5

Figure 4-5 Differential responses ($CS+ > CS-$) in the amygdala correlate with an exponential decay function (inset) for faces (left) and abstract images (right). The SPM is shown at a threshold of $P < 0.001$ uncorrected.


Figure 4-6

Figure 4-6 Differential responses to face presentations (CS+ > CS-) in the fusiform gyrus increase over time (inset). Coordinates for the cluster of activity marked by the crosshairs are given in the upper left corner. This SPM is shown at a threshold of $P < 0.01$ uncorrected to illustrate the extent of activity in the fusiform gyrus.

Figure 4-7

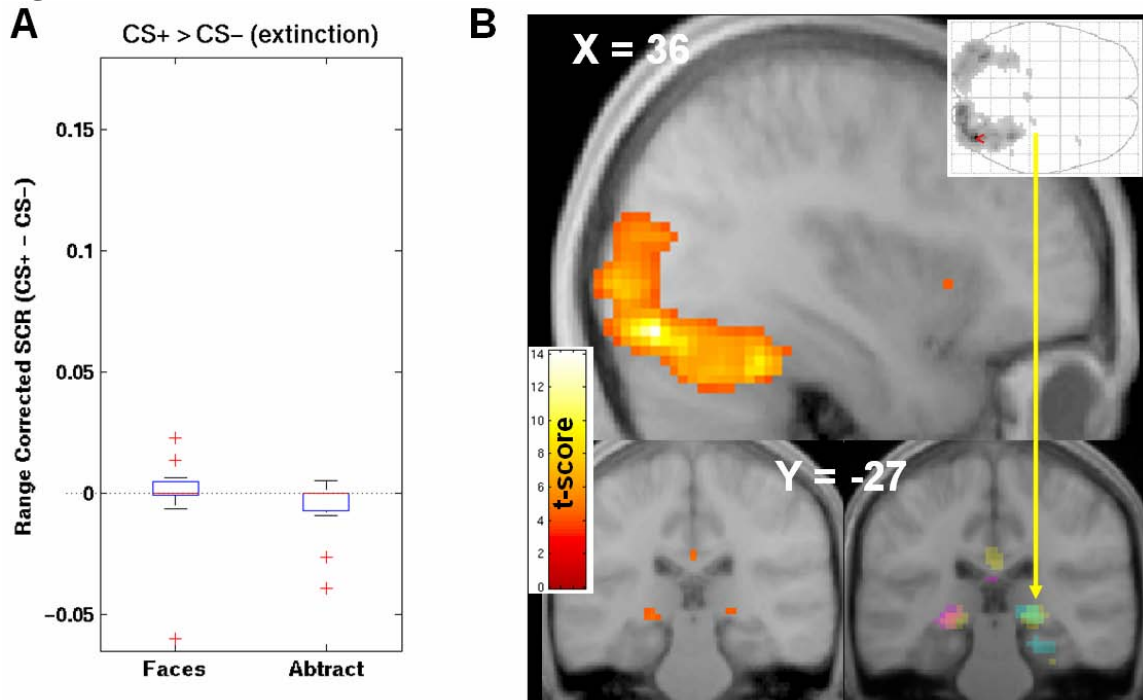


Figure 4-7 Skin conductance and brain activity during extinction. **A**) a box plot (see Figure 4-2) indicating subjects no longer show differential responses to the CS. The conditioned response has been extinguished. **B**) SPMs indicating areas that respond more strongly to the presentation of the face CS+ than face CS- during extinction. These areas include the majority of the visual stream for objects up to and including the fusiform face area (upper SPM, $P < 0.001$ uncorrected). The bottom SPM ($P < 0.001$ uncorrected) on the left is a slice centered on activity in the lateral geniculate nucleus (LGN). Localization of activity to the LGN is confirmed in the bottom right image ($P < 0.01$ uncorrected) that maps the overlap between learned responses (yellow) and activity due to presentation of a wedge stimulus in each visual hemifield (magenta – right; cyan – left).


Figure 4-8

Figure 4-8 Summary of changes in face representation over conditioning and extinction. BOLD activity in the fusiform face area (displayed above in red; see Figure 4-4 and Table 4-2a) increases during acquisition of conditioning (displayed above in yellow; see Figure 4-6). This activity is persistent throughout extinction of the conditioned association (displayed above in magenta; see Figure 4-7). Activity is displayed using a lower threshold than previous figures ($P < 0.01$ uncorrected) to indicate the extent of overlap.