Supplemental Information

A double-blind trial of decoded neurofeedback intervention for specific phobias

Cody A. Cushing, Hakwan Lau, Mitsuo Kawato, Michelle G. Craske, Vincent Taschereau-Dumouchel

Supplemental Methods

Participants

Participant Recruitment
Recruitment was accomplished through flyers, campus website announcements, and posting on online forums (e.g. Nextdoor, etc). Participants completed the modified Fear Survey Schedule (1) in order to identify healthy controls who reported no phobias and individuals who endorsed at least two specific phobias of animals from the ones included in our image dataset. Participants were excluded if they did not meet criteria for MRI scanning safety.

Diagnostic Assessment
All participants underwent a diagnostic interview, using the Anxiety Disorders Interview Schedule-5 (2), administered by trained and reliability certified study staff (Bachelors degree), with each interview reviewed for final consensus by the Principal Investigator (MGC).

Participants were excluded if they 1) did not have normal/corrected to normal vision or hearing; 2) unable to understand informed consent or could not complete the consent form correctly; 3) unable to respond adequately to screening questions; 4) unable to maintain focus/stillness during assessment; 5) had a history of neurological disease or defect; 6) were diagnosed with PTSD, OCD, SUD, current MDD, Bipolar, Psychosis, or any other neurologic diagnoses or unstable serious medical conditions (all assessed using the ADIS-5); 7) currently prescribed psychotropic medication. Participants were not screened for active behavioral treatments.

Groups
Healthy Control Group: No animal type specific phobias or fears, ascertained from administration of the ADIS-5.

Phobia Group: Rated by a diagnostic interviewer (using the ADIS-5 (Brown et al., 2014) to have at least moderate fear or avoidance of at least two animals, with each one associated with an overall rating indicative of at least mild clinical severity. Fear and avoidance were each rated by the interviewer on a 0-8 point scale (0 = no fear/never avoids, 8=very severe fear/always avoids). Clinical severity was rated by the interviewer on a 0-8 point scale that combined symptom severity, distress and impairment associated with each animal stimulus (0=no symptoms, distress and impairment, 8=very severe symptoms, distress, and impairment). Phobias were only eligible if the clinical severity rating was at least mild (a score of 2+). For the 23 participants that were enrolled and started a pre-treatment session, participants had a mean (s.d.) of 2.39 (0.65) phobias. Target phobias had a mean (s.d.) fear rating of 5.17 (1.11), avoidance rating of 5.30 (1.43), and clinical severity rating of 4.65 (1.33). Control phobias had a mean (s.d.) fear rating of 5.91 (1.00), avoidance rating of 5.78 (1.24), and clinical severity rating of 4.70 (1.26).

MRI scanning parameters

All fMRI data were acquired on a 3T Siemens Prisma scanner using a 32-channel head coil at the UCLA Ahmanson-Lovelace Brain Mapping Center.

Decoder Construction

Across 6 task runs during decoder construction, fMRI data were collected with a multi-band sequence with an acceleration factor of 8 and phase encoding in the posterior (P) to anterior (A) direction in order to minimize dropout in the ventral temporal brain area. Voxel sizes were $2.0x2.0x2.0mm^3$ with a $208x208mm^2$ Field of View. Images were collected across 72 interleaved slices with a TR of 800ms, TE of 37.00 ms, and flip angle of 52 degrees. Anatomical data were collected using a T1-weighted imaging sequence with volumetric navigators (vNAV) with prospective motion correction (TR: 2500ms/TI: 1000ms/Flip Angle: 8.0 degrees/Voxel Size: 0.8x0.8x0.8mm/Matrix Size: 256x256/Num. Slices: 208/Slice Thickness: 0.8mm).

Multi-voxel neuro-reinforcement

Prior to the cessation of data collection due to the COVID-19 pandemic, fMRI data during the fear test task and affective Stroop task were collected across 2 runs each using the same sequence described for *Decoder Construction* for 7 participants. However, during the COVID-19 shutdown, this sequence was replaced with a similar but modified sequence better tailored for capturing BOLD activity in subcortical regions such as the amygdala. This replacement sequence used for the remaining 11 participants was a multi-band sequence with an acceleration factor of 6 and phase encoding in the A-P direction. Voxel sizes were $2.0x2.0x2.0mm^3$ with a $192x192mm^2$ Field of View. Images were collected across 72 interleaved slices with a TR of 1000ms, TE of 30.00ms and flip angle of 60 degrees. Accompanying Spin Echo Field Maps were collected in opposing phase encoding directions (A-P/P-A) before functional runs in order to be used for offline distortion correction. FMRI data during online neuro-reinforcement were collected using a multi-band sequence with an acceleration factor of 6 and phase encoding in the P-A direction to minimize dropout in the ventral temporal area. Additional parameters were voxel size: $2.0x2.0x2.0mm^3$, FOV: $208x208mm^2$, num. slices: 72, TR: 1000ms, TE: 37.00 ms, and flip angle: 60 degrees.

Importantly, this change in scanning sequence was only for the fear test task and affective Stroop task. No changes were made that could impact the building of decoders, or multi-voxel neuro-reinforcement itself. Additionally, this scanning sequence change had no effect on overall amygdala Beta estimates in the fear test task ($t(16)=-0.391$, $p=0.70$) or the affective Stroop task ($t(15)=0.853$). Hence, it is highly unlikely any findings in this study are due to the effects of this change in sequence during the covid shutdown.

Decoder Construction

Decoder Construction: task

In place of phobic images, phobic participants viewed happy human faces using stimuli from the Chicago Face Database and NimStim Set of Facial Expressions (3,4). These stimuli have their emotional expression verified by independent raters and were used to provide a non-disturbing stimulus replacement that was sufficiently orthogonal to the task image set of animals and objects. The decoder construction task consisted of 6 runs of 600 trials each. Each trial consisted of a .98 second image presentation with no inter-trial interval. This rapid event-related design was used to maximize the number of images each participant viewed. To ensure attention, participants were given the task of pressing a button each time the image category changed (i.e. a 1-back task). Image categories were presented in chunks of 2, 3, 4, or 6 consecutive images.

Decoder construction: fMRIprocessing

Decoder construction fMRI data were processed using a combination of SPM12 (Statistical Parametric Mapping; www.fil.ion.ucl.ac.uk/spm) and custom python scripts using pyMVPA and sklearn packages (5,6). All 6 runs of the task were concatenated and preprocessed in SPM using default parameters unless otherwise explicitly specified. Data were realigned to the first image from the first run of the task and segmented into tissue classes. Anatomical and functional data were coregistered using the gray matter image from segmentation as a reference. Motion was then regressed out of the functional data using the 6 head motion parameters from realignment. Single-trial estimates were then generated with pyMVPA using the least-squares 2 (LS-2) method (7) in which a separate GLM is computed for each trial where the current trial is assigned to one regressor while the remaining trials are equally split between two “rest” regressors.

Using hyperalignment, single-trial estimates from healthy controls in the target brain region (ventral temporal cortex) were functionally transformed to the current phobic participant’s brain and used to train a machine-learning pattern classifier (decoder) using the phobic images that the participant did not see (Fig. 1). To ensure double-blind treatment target selection, the target for treatment was automatically selected by a computer program that calculated which phobic category had the highest cross-validated area under the receiver operating characteristic curve (AUC) during binary one vs. all classification.

To determine AUC metrics, a 6-fold cross-validation (CV) procedure was used. FMRI data for each participant were loaded and masked to the ventral temporal (VT) area in their own native space using an anatomical mask derived from combining the entirety of the Freesurfer parcellations of the fusiform, lingual, parahippocampal, and inferior temporal areas (8). Single-trial parameter estimates were standardized by feature within subject and within each of the 6 task runs. The data were split into 6 folds for training and testing based on the 6 runs completed by each participant. That is, for each CV split, the withheld testing set consisted of all the data from each participant for one of the six task runs. The remaining preprocessing was calculated using only the training data to avoid overfitting. As hyperalignment requires a stable number of features across participants, 1000 voxels were selected within the VT area via F-test to select which voxels accounted for the most variance elicited by all image categories across all training trials. For each phobic participant, a unique set of hyperalignment transformation

parameters into the common model space was calculated for the current phobic participant and all healthy controls. The fitting of the hyperalignment parameters was done using trials for all image categories except the current participant's phobias. For example, if a phobic participant had spider and snake phobias, all spider and snake trials were withheld from all participants when fitting the transformation parameters.

After hyperalignment transformation parameters were determined, the data from all healthy controls were moved into the native space of the current phobic participant by transforming the data into the common model space and then reverse transforming the data from the common model space into the native space of the current participant. The transformed data included the previously withheld phobic category images from the healthy controls as well as the testing dataset.

With all data in the current participant's native space, class sizes (target vs. non-target image categories) were balanced by random undersampling balanced between the 39 non-target image categories. Following previous work (9), a Sparse Multinomial Logistic Regression (SMLR) classifier was trained to perform binary (one-vs-rest) classification between the potential target category and all remaining categories (10). AUC scores for each CV split were calculated based on classifier estimates.

For the final decoder to be used in neuro-reinforcement, the same procedure was performed but trained using all 6 runs of data.

Pre/Post Test

Fear test: task

During each trial, a fixation cross was presented for 3-7 seconds, followed by a static image for 6 seconds. After the static image, a blank screen was displayed for 4-12 seconds followed by a prompt to enter how fearful they found the image on a 7-point scale. These ratings were used as the subjective fear ratings to test hypothesis H3. Images displayed either belonged to the target phobia, control phobia, neutral animal, or neutral object categories. Neutral animals and objects were randomly selected based on categories for which a given participant reported no fear during their diagnostic interview. Participants completed two runs of 15 images each with a self-paced break between runs. Within each run, they viewed 5 target phobia images, 5 control phobia images, and 2-3 neutral animal/object images, counterbalanced across runs. The first image of each run was a neutral object, always immediately followed by either a target phobia or control phobia image, counterbalanced across runs. The remaining images within a run were randomly selected from each category.

Fear test: fmri processing

FMRI task runs were distortion corrected using FSL's topup (11,12) according to spin echo field map sequences collected in opposite phase-encoding directions. Due to technical issues with spin echo field map collection, 5 participants were excluded from distortion correction. Anatomical T1 images were brain extracted using bet (13). Then, preprocessing and ICA-

decomposition were performed using FSL's melodic and FEAT (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl). During preprocessing, fMRI data were motion corrected using mcflirt (14), brain extracted using bet (13), spatially smoothed with a Gaussian kernel of FWHM 4.0mm, intensity normalized, and highpass filtered with a gaussian-weighted least-squares straight line fitting with sigma=50.0s. Images were then registered to the standard MNI space using FLIRT and then refined using nonlinear registration with FNIRT (14,15). Registration of multi-band images were improved by using a high-contrast single-band reference image collected at the start of each functional run as an initial reference image for registration.

ICA components were then manually investigated with components resulting from movement or other sources of noise removed. To further account for movement, data were processed with the Artifact Detection Tools (ART, https://www.nitrc.org/projects/artifact_detect) toolbox to generate motion regressors and identify outlier timepoints for censoring. First-level GLMs were then calculated in SPM12 with a temporal derivative to account for slice-timing differences. Regressors were fit for the onset of target phobia, control phobia, neutral animal, and neutral object images with a duration of 0 seconds to model the event-related response. Following previous work (9), only the first 2 trials within each run were analyzed for target phobia and control phobia images.

Bilateral amygdala masks were generated from the automatic Freesurfer segmentation of the T1 image and transformed into the participant's native functional space. Average parameter estimates were extracted from the Amygdala using marsbar (16). Average parameter estimates for phobic stimuli were then corrected to baseline by subtracting the average amygdala response to the neutral animal from the target phobia and control phobia, within runs. Baseline-corrected phobia responses were then averaged across runs for pre-treatment and post-treatment sessions.

Affective Stroop: task. The task started with a 1 second red fixation cross and then a brief (300 ms) image from either a phobic or neutral control category. As soon as the image appeared, participants were instructed to, as quickly and accurately as they could, make a size judgment about whether the presented animal could fit in their hand (i.e. is it the size of your hand or smaller?), by pressing one of two buttons with their index and middle finger to indicate yes or no. Response-key mappings were counterbalanced across participants. There was a 1.2 second response period (indicated by a blue fixation cross) following stimulus offset for response entry followed by a fixed 1 second inter-trial interval. Stimuli were selected from 7 animal categories: target phobia, control phobia, and 5 neutral animal categories. Similar to the fear test, neutral animal categories were selected from categories for which a given phobia participant reported no fear during their diagnostic interview. The task consisted of 210 randomly distributed trials split over 2 fMRI runs with a self-paced break between runs.

Affective Stroop: analysis. Reaction times were extracted for target phobia, control phobia, and neutral animal stimuli using custom scripts in MATLAB (Mathworks Inc., Natick, MA). Responses were coded as correct or incorrect based on unanimous agreement from 8

independent raters who judged whether each of the 30 potential animal categories was the size of their hand or smaller: unanimity was obtained for 24 animal categories; animal categories without consensus (bird, bat, fish, gecko, turtle, and guinea pig) were treated as correct as long as a response was recorded.

Multi-voxel neuro-reinforcement

Online real-time fMRI processing. Real-time fMRI processing for multi-voxel neuro-reinforcement was conducted in MATLAB with the decoded neurofeedback software developed at Advanced Telecommunications Research Institute International (https://bicr.atr.jp/decnefpro/software/). Incoming dicom images exported from the scanner were converted to nifti, realigned to a template image, the first dicom from the first run of the decoder construction task, then detrended based on all the data collected up to that TR in a given run. Proper alignment between the real-time fMRI data and the decoder construction data (on which the decoder was based) was ensured by correlating multi-voxel patterns between the real-time data and the decoder construction template. If pattern correlation fell below a threshold of 0.70 on a given trial, visual feedback was not displayed to the participant.

Monetary Reward. The size of the feedback disc determined the amount of reward the participant received at the end of each run, with their average feedback score determining the percentage of that run's total bonus received. For example, an average feedback score of 60% resulted in 60% of the potential $6.00 bonus being received (i.e. $3.60). An additional bonus was also given when participants were able to generate a feedback score of 70% or more for 3 trials in a row. Participants were given an additional $2.00 per high-score streak bonus which was visually indicated by the feedback disc turning blue with a written message alerting them to their high-score streak.

Skin Conductance Response (SCR)

Data collection

Skin Conductance Response (SCR) was recorded in Acknowledge software via Biopac MP-150 system using the EDA-100C module and Ag/AgCl electrodes placed distally on the index and middle fingers of the left hand. SCR recordings were taken during pre-treatment and post-treatment MRI scanning sessions. Of the 18 participants analyzed in our main analyses, 5 participants had technical issues during data collection and 4 participants were non-responders showing no discernable SCR. Consequently, 9 participants were analyzed for SCR.

Data analysis

SCR recordings were analyzed with custom code in python utilizing the bioread package. SCR data were filtered with a 1st-order 5 Hz low-pass Butterworth filter to account for influences of the magnetic field in the MRI environment. SCR recordings were then epoched according to stimulus onset times during the Fear Test task from 2 seconds preceding stimulus onset to 5 seconds following stimulus onset. Epoch timecourses were baseline corrected according to the average activity during the 2 seconds before stimulus onset. Peak SCR values were then

extracted for each trial epoch by taking the maximum SCR value in the time period of 1 second to 5 seconds following stimulus onset. If the peak SCR value was less than 0.02 microsiemens then it was coded as 0 following previous research (ref). Peak SCR values were then square root transformed in preparation for statistical analysis.

Self Report Questionnaires. The following self report questionnaires were administered at pre-treatment and post-treatment:

Depression, Anxiety, and Stress Scale (DASS-21) (17), Behavioral Inhibition/Behavioral Activation Scale (BIS/BAS) (18), Sheehan Disability Scale (SDS) (19), and Modified Fear Survey Schedule (1).

Supplemental Results

Hyperalignment decoding

Binary classification performance of hyperaligned decoders was estimated in a 6-fold cross-validation on decoder construction task data. Average target decoder AUC was 0.63 (0.03), which was significantly greater than the chance level of 0.50 ($t(22)=20.3$, $p<0.001$). This indicates that category-level visual representations can be significantly decoded in the brains of participants with specific phobia based on hyperaligned surrogate brain data from healthy controls.

Double-blinded placebo control

After neuro-reinforcement, the experimenter revealed to participants that neuro-reinforcement feedback had been based on the visual representation of one of their phobias. When asked to pick between two of their phobias (the target and control phobias, blinded to the experimenter), participants were unable to correctly guess the identity of their neuro-reinforcement target (43% accuracy; chance level 50%). Participants reported strategies for neuro-reinforcement that were unrelated to the target and control animal categories. Collectively, this indicates neuro-reinforcement was carried out in a double-blind fashion at an implicit level with participants unaware of the target of the intervention.

Target pattern induction

To assess the degree to which the desired pattern associated with the target phobic category was activated by patients during neuro-reinforcement, the feedback scores patients saw (representing degree of desired neural pattern activation) during neuro-reinforcement were compared to the scores patients would have seen if feedback had been based on the control phobic category pattern instead. Both these target and control scores were generated using the real-time pipeline - the only exception being the fMRI data was detrended across the entire run

(versus how much data had been collected up to a given 'trial') as this is how the data were saved at the end of the neurofeedback program once the entire run had been collected. In the 18 participants analyzed for our primary outcome, the feedback was significantly higher for the target phobic category compared to what it would have been for the control phobic category ($t(17)=12.63$, $p<0.001$) (Fig. 2B). This result indicates that the desired target pattern was successfully activated by patients during neuro-reinforcement. Results for each individual day are reported in Supplemental Fig. S1.

Amygdala response during Stroop task

Amygdala responding during the affective Stroop task did not demonstrate the same interaction we observed during the fear test task ($F(1,14)=1.075$, $p=0.317$) counter to our pre-registered hypothesis H4iii. Additionally, in the affective Stroop task, a phobia response was not observed in response to the target phobia pre-treatment as tested with a one-sample t-test on the baselined parameter estimates ($t(16)=0.19$, $p=0.85$). This lack of significant phobia response pre-treatment could be due to the increased cognitive load of this task which required rapid, reflexive judgments as soon as the stimulus appeared (compared to fear ratings in the fear test which were input many seconds after the original stimulus disappeared). Additionally or alternatively, the amygdala may have habituated during the affective Stroop task as it was always immediately preceded by the fear test.

Between-subjects analysis of dosage effects (H5)

Although circumstances outside of our control (detailed in methods) prevented us from collecting a sufficient sample size to analyze the between-subject effect of dosage with sufficient power as we initially pre-registered, we report the pre-registered analysis here. When dosage (1, 3, or 5 days of neuro-reinforcement) is treated as a between-subjects factor in a 3 (between-subjects dosage: 1, 3, or 5 days of neuro-reinforcement) x 2 (within-subjects condition: target, control phobia) x 2 (within-subjects time: pre-treatment, post-treatment) repeated-measures ANOVA, we fail to find evidence in support of H5. The 3-way interaction between dosage, condition, and time is trending but not significant ($F(2,14) = 3.236$, $p=0.07$, $\eta_p^2=0.316$). Within each group, significant differences could not be detected for target responding pre- to post-treatment. The greatest effect detected was for the 3-day group for a reduction in amygdala responding to the target phobia pre- to post-treatment ($t(5)=0.137$, $p=0.137$). This lack of evidence in support of our pre-registered hypothesis H5 is most likely due to insufficient power (6 participants in each dosage group) to detect such between-subjects effect in the current design. Future studies will be needed to address the question of how the number of neuro-reinforcement sessions an individual receives affects reduced amygdala responses to feared stimuli.

Self-Report Questionnaires

A paired sample t-test for Depression, Anxiety and Stress Anxiety Subscale was marginally significant, t(17) = 2.06, p = .055: pre-test (M = 8.9, SD = 2.7) and post-test (M = 8.2, SD = 1.6) indicating a marginal decrease in anxiety following neuro-reinforcement. There were no effects for the depression subscale or stress subscale or the total DASS score.

Assessment of additional covariate

We could not anticipate how much variance would be present in the number of phobias amongst participants with multiple phobias. While the final analyzed sample had a mean (s.d.) number of phobias of 2.39 (0.65), some participants during recruitment had as many as 5-7 phobias. This not only indicates a more widespread experience of clinical fear but also introduces a potential difference during the decoder construction process as the total number of phobias had to be withheld from data preprocessing during the hyperalignment process. For these reasons, we elected to include the number of phobias as a covariate in statistical analyses. The interpretation of the reported results remains the same when this covariate is not included in the model. More specifically, the interaction effect for H1 was still significant (($F(1,16)=5.17$, $p=0.037$, $\eta_p^2=0.244$) and the interaction effect for H4 which was near significance, remained non-significant ($F(1,15)=0.041$, $p=0.84$, $\eta_p^2=0.000069$).

Supplemental Tables

Supplementary Table 1. Participant demographics by dosage group

Race	**1 day**	**3 day**	**5 day**
White	2	3	4
Black	0	1	1
Asian/Pacific Islander	4	4	2
Other	1	0	0
Not Reported	1	1	1
Ethnicity			
Hispanic	1	1	3
Non-Hispanic	7	6	5

Gender			
Male	2	0	3
Female	6	7	5
Non-binary	0	0	0
Age: mean(sd)	25.5 (6.6)	26.3 (12.9)	27.6 (9.4)
Education Level			
High School	1	1	2
Some College	1	1	1
Associates/2-year degree or higher	6	5	5

Supplemental Figures

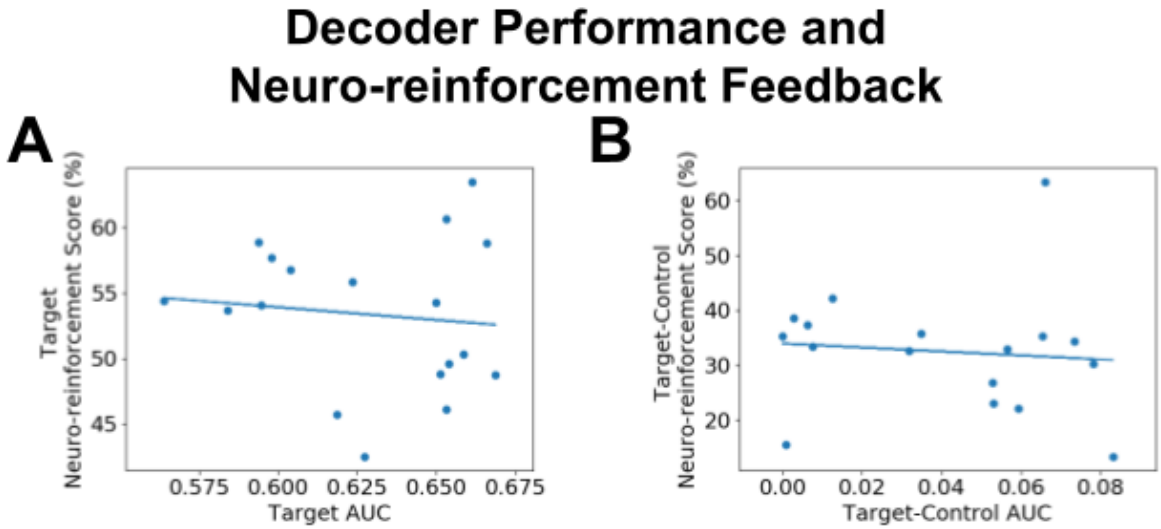

Supplemental Figure S1. No relation between decoder construction performance and neuro-reinforcement scores. Scatter plots of average cross-validated AUC scores during decoder construction and neuro-reinforcement scores from neuro-reinforcement sessions. Solid lines indicate line of best fit. (A) Association for the target phobic category (B) Association for the difference between the target and control phobic categories.

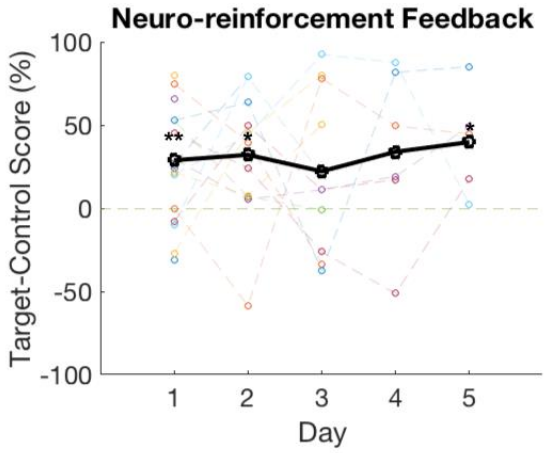

Supplemental Figure S2. Neuro-reinforcement feedback for target versus control phobias by day. Daily averages of target minus control scores for each participant are plotted in colored data points. Each color codes for an individual participant across days, connected by dashed lines. Black data points represent average from all participants on a given day, connected via solid black line. Individual days with significantly greater target feedback compared to control feedback according to paired t-tests are marked with an asterisk. * p<0.05, ** p<0.01

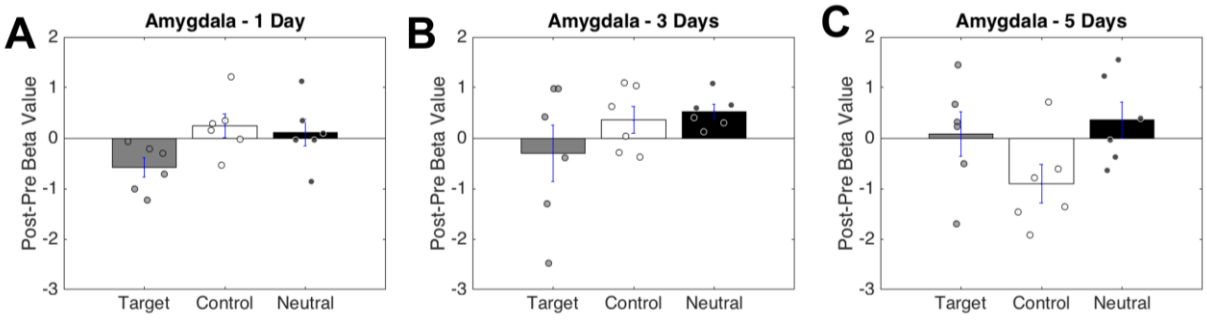

Supplemental Figure S3. Amygdala response following neuro-reinforcement by number of days of neuro-reinforcement received. Panels show changes in responses to target phobia, control phobia, and neutral animal images from pre-neuro-reinforcement to post-neuro-reinforcement, quantified as post minus pre difference. Results for participants that received 1 day (A), 3 days (B), or 5 days (C) of neuro-reinforcement.

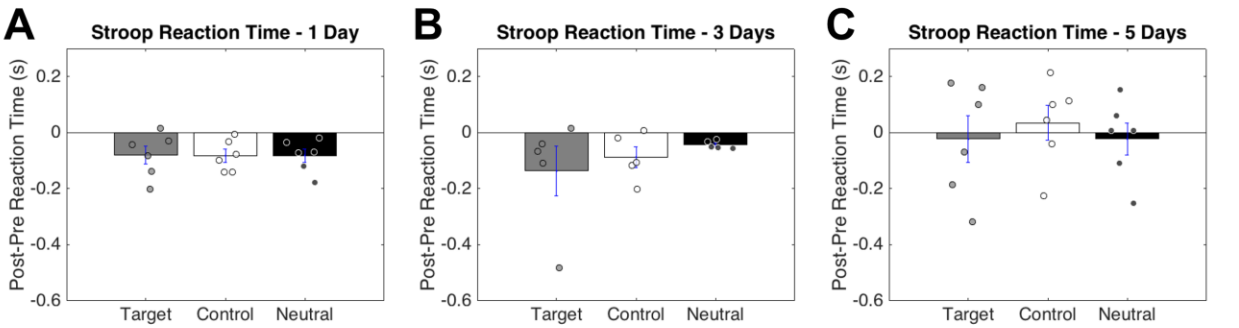

Supplemental Figure S4. Reaction times in affective Stroop task following neuro-reinforcement by number of days of neuro-reinforcement received. Panels show changes in reaction times to target phobia, control phobia, and neutral animal images from pre-neuro-reinforcement to post-neuro-reinforcement, quantified as post minus pre difference. Results for participants that received 1 day (A), 3 days (B), or 5 days (C) of neuro-reinforcement.

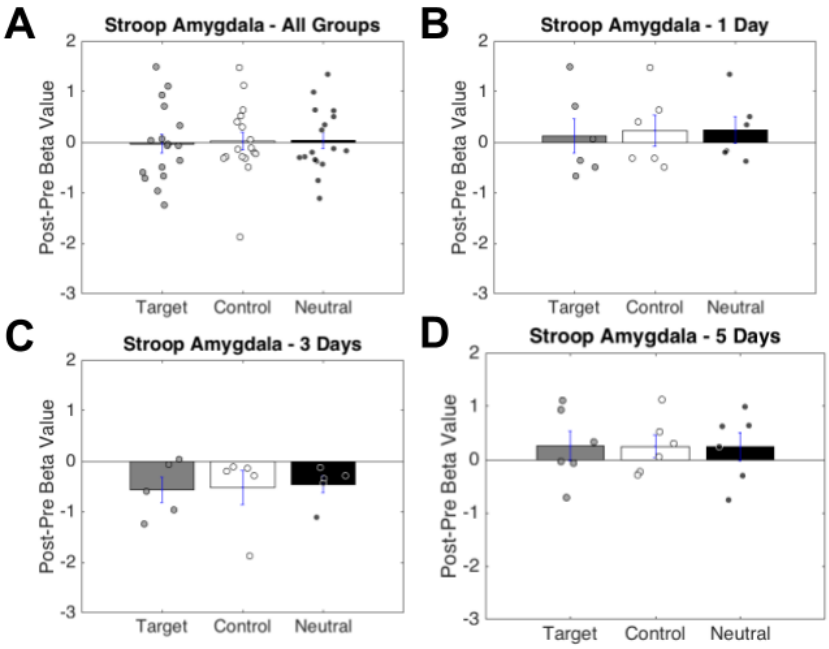

Supplemental Figure S5. Amygdala response following neuro-reinforcement in affective Stroop task. Results for all dosage groups combined (A) showing post-treatment minus pre-treatment amygdala responses to target phobia, control phobia, and neutral animals in the affective Stroop task. Also, the 1 day (B), 3 days (C), and 5 days (D) of neuro-reinforcement are also shown for illustrative purposes.

References

1. Wolpe J, Lang PJ (1964): A FEAR SURVEY SCHEDULE FOR USE IN BEHAVIOUR THERAPY. *Behav Res Ther* 2: 27–30.

2. Brown T, Barlow D (2014): *Anxiety and Related Disorders Interview Schedule for DSM-5 (ADIS-5)® - Adult Version: Client Interview Schedule 5-Copy Set*. Oxford, New York: Oxford University Press.

3. Ma DS, Correll J, Wittenbrink B (2015): The Chicago face database: A free stimulus set of faces and norming data. *Behav Res Methods* 47: 1122–1135.

4. Tottenham N, Tanaka JW, Leon AC, McCarry T, Nurse M, Hare T a., *et al.* (2009): The NimStim set of facial expressions: Judgements from untrained research participants. *Psychiatry Res* 168: 242–249.

5. Hanke M, Halchenko YO, Sederberg PB, Hanson SJ, Haxby JV, Pollmann S (2009): PyMVPA: a Python Toolbox for Multivariate Pattern Analysis of fMRI Data. *Neuroinformatics* 7: 37–53.

6. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, *et al.* (2011): Scikit-learn: Machine Learning in Python. *J Mach Learn Res* 12: 2825–2830.

7. Turner BO, Mumford JA, Poldrack RA, Ashby FG (2012): Spatiotemporal activity estimation for multivoxel pattern analysis with rapid event-related designs. *NeuroImage* 62: 1429–1438.

8. Fischl B, Van Der Kouwe A, Destrieux C, Halgren E, Ségonne F, Salat DH, *et al.* (2004): Automatically Parcellating the Human Cerebral Cortex. *Cereb Cortex* 14: 11–22.

9. Taschereau-Dumouchel V, Cortese A, Chiba T, Knotts JD, Kawato M, Lau H (2018): Towards an unconscious neural reinforcement intervention for common fears. *Proc Natl Acad Sci U S A* 115: 3470–3475.

10. Krishnapuram B, Carin L, Figueiredo MAT, Hartemink AJ (2005): Sparse multinomial logistic regression: fast algorithms and generalization bounds. *IEEE Trans Pattern Anal Mach Intell* 27: 957–968.

11. Andersson JLR, Skare S, Ashburner J (2003): How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *NeuroImage* 20: 870–888.

12. Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TEJ, Johansen-Berg H, *et al.* (2004): Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage* 23 Suppl 1: S208-219.

13. Smith SM (2002): Fast robust automated brain extraction. *Hum Brain Mapp* 17: 143–155.

14. Jenkinson M, Bannister P, Brady M, Smith S (2002): Improved optimization for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* 17: 825–841.

15. Jenkinson M, Smith S (2001): A global optimisation method for robust affine registration of brain images. *Med Image Anal* 5: 143–156.

16. Brett M, Anton J-L, Valabregue R, Poline J-B (2002): Region of interest analysis using an SPM toolbox. *NeuroImage* 16: 497.

17. Lovibond PF, Lovibond SH (1995): The structure of negative emotional states: Comparison of the Depression Anxiety Stress Scales (DASS) with the Beck Depression and Anxiety Inventories. *Behav Res Ther* 33: 335–343.

18. Carver CS, White TL (1994): Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: The BIS/BAS Scales. *J Pers Soc Psychol* 67: 319–333.

19. Sheehan DV, Harnett-Sheehan K, Raj BA (1996): The measurement of disability. *Int Clin Psychopharmacol* 11: 89–95.