PNAS RESEARCH ARTICLE | NEUROSCIENCE

Distinct prelimbic cortex ensembles encode response execution and inhibition

Rajtarun Madangopal[a,1], Yuan Zhao[b], Conor Heins[a], Jingfeng Zhou[c,2], Bo Liang[a,3], Giovanni Barbera[a], Ka Chun Lam[b], Lauren E. Komer[a], Sophia J. Weber[a], Drake J. Thompson[a], Yugantar Gera[a], Diana Q. Pham[a], Katherine E. Savell[a], Brandon L. Warren[a,4], Daniele Caprioli[a,5,6], Marco Venniro[a,7], Jennifer M. Bossert[a], Leslie A. Ramsey[a], Hank P. Jedema[a], Geoffrey Schoenbaum[c], Da-Ting Lin[a], Yavin Shaham[a,8], Francisco Pereira[b,8], and Bruce T. Hope[a,1,8]

Affiliations are included on p. 11.

Edited by Huda Akil, University of Michigan, Ann Arbor, MI; received March 8, 2025; accepted August 4, 2025

Learning when to initiate or withhold actions is essential for survival, requiring the integration of past experiences with new information to adapt to changing environments. The prelimbic cortex (PL) plays a central role in this process, with a stable PL neuronal population (ensemble) recruited during operant reward learning to encode response execution. However, it is unknown how this established reward-learning ensemble adapts to changing reward contingencies, such as reward omission during extinction. Specifically, does the same ensemble adjust its activity to support behavior suppression, or is a distinct ensemble recruited for this new learning? Our data reveal that operant extinction learning recruits a distinct PL Extinction ensemble to support response inhibition, and concerted engagement of both ensembles encodes both ongoing and subsequent context-specific behavior. Using single-cell calcium imaging, we longitudinally tracked PL neurons in rats as they pressed a lever for food rewards (Training), learned to suppress responding upon reward omission (Extinction), and reinstated responding following a noncontingent "priming" pellet (Reinstatement). We trained decoders on individual rats' PL activity patterns to predict trial-wise responses and used an in silico deletion approach to identify separate PL Training and Extinction ensembles associated with response execution and inhibition, respectively. Critically, both ensembles were reengaged and maintained their distinct roles during Reinstatement. These findings highlight ensemble-based encoding of multiple, even opposing, learned associations within the same region, demonstrating how selective ensemble recruitment enables behavioral flexibility under changing contingencies.

operant learning | extinction | reinstatement | longitudinal in vivo calcium imaging | behavioral decoding

Significance

Until recently, the question of how neuronal ensembles support behavioral flexibility in the face of changing contingencies has remained unanswered. Using longitudinal single-cell calcium imaging, rat-specific decoders, and in silico ensemble deletions, we show that operant training and extinction learning recruit distinct prelimbic cortex ensembles to support opposing behavioral responses, and these same ensembles are reengaged during reinstatement. We provide ensemble-based evidence that the brain encodes extinction as a new learned association rather than through weakening of prior associations and that reinstatement relies on the interplay of preestablished training and extinction ensembles rather than recruitment of a distinct third ensemble. Overall, these results highlight how the brain encodes and engages multiple learned associations through distinct ensembles, even within a single brain region.

Learned stimulus–response–outcome sequences are central to survival and can guide response execution (or inhibition) in pursuit (or avoidance) of the original outcome. The medial prefrontal cortex (mPFC) plays a crucial role in various aspects of learned appetitive and aversive behaviors, such as attention, decision-making, and memory (1–6). Early inactivation studies using the extinction-reinstatement model led to the hypothesis that the dorsal (prelimbic cortex, PL) and ventral (infralimbic cortex, IL) mPFC subdivisions play opposite roles in behavioral expression and inhibition, respectively (7–10). However, subsequent studies have questioned this PL-go/IL-stop dichotomy and suggest instead that sparse and intermingled populations of neurons (neuronal ensembles) within these regions (11–20) control the selection of distinct task-dependent responses, rather than merely promoting or inhibiting behavior (6, 21–23).

In line with this hypothesis, in vivo electrophysiological recordings have identified diverse cue, response/inhibition, and reward-related response patterns at the single-neuron level in both subregions across multiple behavioral tasks (21, 22, 24, 25). However, as most studies used electrode arrays or bundles, they were limited to sparse sampling within the region during a given session and were not optimized to identify concerted population-level activity patterns (i.e., ensembles) during behavior. More recent studies applied single-cell resolution calcium imaging to simultaneously monitor hundreds of PL neurons during behavior (26–29) and identified PL ensembles mapping onto behavioral sequences during operant food self-administration training (29), as well as unique PL ensembles during Pavlovian reward learning (26). Interestingly, these ensembles showed stable coding patterns established during learning and could be used to infer task states (cue/trial type, reward delivery) and predict both task-related (press/lick) and non-task-related behaviors during the imaging session.

[1]To whom correspondence may be addressed. Email: rajtarun.madangopal@nih.gov or bhope@intra.nida.nih.gov.

[2]Present address: State Key Laboratory of Cognitive-Neuroscience and Learning, Beijing Normal University & Chinese Institutefor Brain Research, Beijing 100875, China.

[3]Present address: School of Electrical Engineering & Computer Science, College of Engineering & Mines, University of NorthDakota, Grand Forks, ND 58202.

[4]Present address: Department of Pharmacodynamics, University of Florida, Gainesville, FL 32610.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2505378122/-/DCSupplemental.

Published September 8, 2025.

What happens to these stable "training" ensembles when individuals adjust their behavior in response to changing reward contingencies, such as suppressing responses during extinction (reward omission) or reengaging in reward-seeking during reinstatement? Does behavioral flexibility rely on changes within the same ensemble's activity patterns, or does it involve the recruitment of distinct PL neuron populations tailored to each new behavioral context? Specifically, we asked whether the same PL training ensemble supports training, extinction, and reinstatement of responding for a reward through changing activity patterns, or if separate nonoverlapping ensembles are recruited to mediate the appropriate action during each of these phases.

To answer this question, we used miniature epifluorescent microscopes (Miniscopes) (28, 30, 31) for longitudinal single-cell resolution calcium imaging of PL neurons in rats first as they made operant responses to receive palatable food pellet rewards in a trial-based procedure (Training), next as they learned to suppress responding when the reward was omitted (Extinction), and finally when they reinstated responding upon receiving a noncontingent "priming" pellet (Reinstatement) a manipulation known to reinstate operant responding after extinction (32, 33). We analyzed active neuron populations and activity patterns as rats modified their behavior and developed decoders operating on individual rats' PL activity and in silico deletion tests to determine whether the same PL ensemble neurons support training, extinction, and reinstatement of responding for food reward through changing activity patterns, or if separate nonoverlapping ensembles are recruited to mediate the appropriate action during each of these phases.

Results

Training, Extinction, and Reinstatement of Palatable Food Seeking Using a Trial-Based Operant Procedure. Prior to imaging, we trained rats to self-administer palatable food pellets (32, 34) in three stages consisting of increasing numbers of trials and progressively shorter lever availability periods (*SI Appendix*, Fig. S1*B*). Rats learned to make responses on the active (food-paired) lever to receive palatable food pellets during initial self-administration training (*SI Appendix*, Fig. S1*C*, left; *Lever* × *Session*: $F_{1.8,14.6} = 12.85$, $P < 0.001$), and maintained stable responding during both phases of postsurgery training (*SI Appendix*, Fig. S1*C*, center and right; *Lever* × *Session*: phase 1: $F_{3.2,25.6} = 9.8$, $P < 0.001$; phase 2: $F_{4.1,32.8} = 4.93$, $P = 0.003$). Following training, we recorded activity dynamics of PL neurons during the last two sessions of food self-administration (Training sessions; T1 and T2), four consecutive sessions of extinction (Extinction sessions; E1 to E4), and one session of palatable food-primed reinstatement (Reinstatement session; R) (Fig. 1*B*, 100 trials/session). Rats maintained stable lever pressing during the two imaged self-administration training sessions, decreased lever pressing across the four imaged extinction sessions, and reinstated lever pressing in response to noncontingent priming food pellet delivery during reinstatement testing (*Lever* × *Session*: $F_{1.3,9.02} = 18.75$, $P = 0.001$). Detailed statistical results for behavior during training and imaging sessions are provided in *SI Appendix*, Table S1.

Following imaging, we applied image processing routines to extract fluorescence time courses and spatial locations of neurons for each rat using concatenated imaging data from all seven imaging sessions (Min/Median/Max = 124/159/350 neurons, $n = 9$ rats). We used individual rats' PL neuronal activity patterns and trial-wise behavioral measures to 1) quantify the number and percentage of neurons active during specific session types and trial periods; 2) analyze the proportions of overlapping and unique neurons between session types and trial periods; 3) train rat-specific decoders to predict trial-wise behavioral responses using PL neuronal activity during Training and Extinction sections, and 4) test the contributions of specific neuron populations of interest to behavioral decoding accuracy during operant Training, Extinction, and Reinstatement.

Number and Percentage of Active Neurons by Session and Trial Period. We first broadly quantified PL activity levels during different behavioral epochs and across imaging sessions. For each rat, we estimated neuron spiking from calcium transients and identified *active* neurons during four periods of interest (Fig. 1*C*)—1) trial period: start to end of behavioral trial, 2) pre-lever period: trial start to active lever presentation, 3) lever period: active lever presentation to retraction, and 4) post-lever period: active lever retraction to trial end. For each session, a neuron was deemed *active* in a period if its firing rate within that period was significantly different from that of the baseline period 5 s prior to start of the trial (*SI Appendix*, Fig. S2*A*). For each rat, we divided the *active neuron* counts (*SI Appendix*, Fig. S2*B*) by the total number of detected neurons and compared *percent active* neurons within each trial period, across the seven imaging sessions (Fig. 1*D*), or collapsed by session-type [three session types: Training (T), Extinction (E), and Reinstatement (R); *SI Appendix*, Fig. S3*A*]. Overall, the percentage of active neurons increased during the lever availability period and decreased during the post-lever period during Extinction sessions (relative to Training and Reinstatement sessions), while the percentage of active neurons was unaltered for the pre-lever period and for the entire trial period across sessions.

The repeated-measures two-way ANOVA of *Percent active neurons* across all seven sessions, which included within-subjects factors of Imaging session (seven sessions: T1–2, E1–4, and R; Fig. 1*D*), and trial period (pre-lever, lever, or post-lever), showed a significant interaction between the two factors ($F_{2.7,21.9} = 15.63$, $P < 0.001$). Follow-up one-way ANOVAs were significant for the lever period ($F_{2.3,18.6} = 13.16$, $P < 0.001$) and post-lever period ($F_{2.8,22.7} = 11.38$, $P < 0.001$). Post hoc analysis indicated a significant increase in percent active lever period neurons (E2, E4 vs. T1) and a significant reduction in percent active post-lever period neurons (E1, E2, and E4, vs. T1) during extinction imaging sessions. In contrast, similar levels of activity were observed across sessions when considering the entire trial period ($F_{1.9,15.5} = 1.89$, $P = 0.19$), or the pre-lever period ($F_{1.9,14.9} = 1.01$, $P = 0.38$). The same pattern was maintained when we collapsed the sessions by type (three session types: T, E, R; *SI Appendix*, Fig. S3*A*). Detailed statistical results are provided in *SI Appendix*, Tables S2 (seven sessions) and S3 (three session types).

Overlapping vs. Unique Active Neurons by Session and Trial Period. Next, we asked whether the same (overlapping) or different (unique) sets of neurons are active over the seven imaged sessions. For each rat and period of interest, we identified overlapping vs. unique active neuron populations between pairs of the seven imaged sessions (Fig. 1*E*), and between pairs of the three session-types (*SI Appendix*, Fig. S3). We calculated 1) Overlapping neurons—the percent of *active* neurons shared between two sessions/session-types as %Overlap = active in both Sessions A and B/active in Sessions A or B; i.e., $|A \cap B|/|A \cup B|$ (Fig. 1 *E-1* and *SI Appendix*, Fig. S3*B*) and 2) Unique neurons—the percent of *active* neurons that were unique to a session/session-type as $\%Unique_A$ = active in Session A but not Session B/active in Session A = $|A-B|/|A|$ (Fig. 1 *E-2* and *SI Appendix*, Fig. S3*C*).

Overlapping vs. unique neurons within and across Training and Extinction sessions. Overall, the percentage of overlapping neurons was high within pairs of Training and Extinction sessions, but

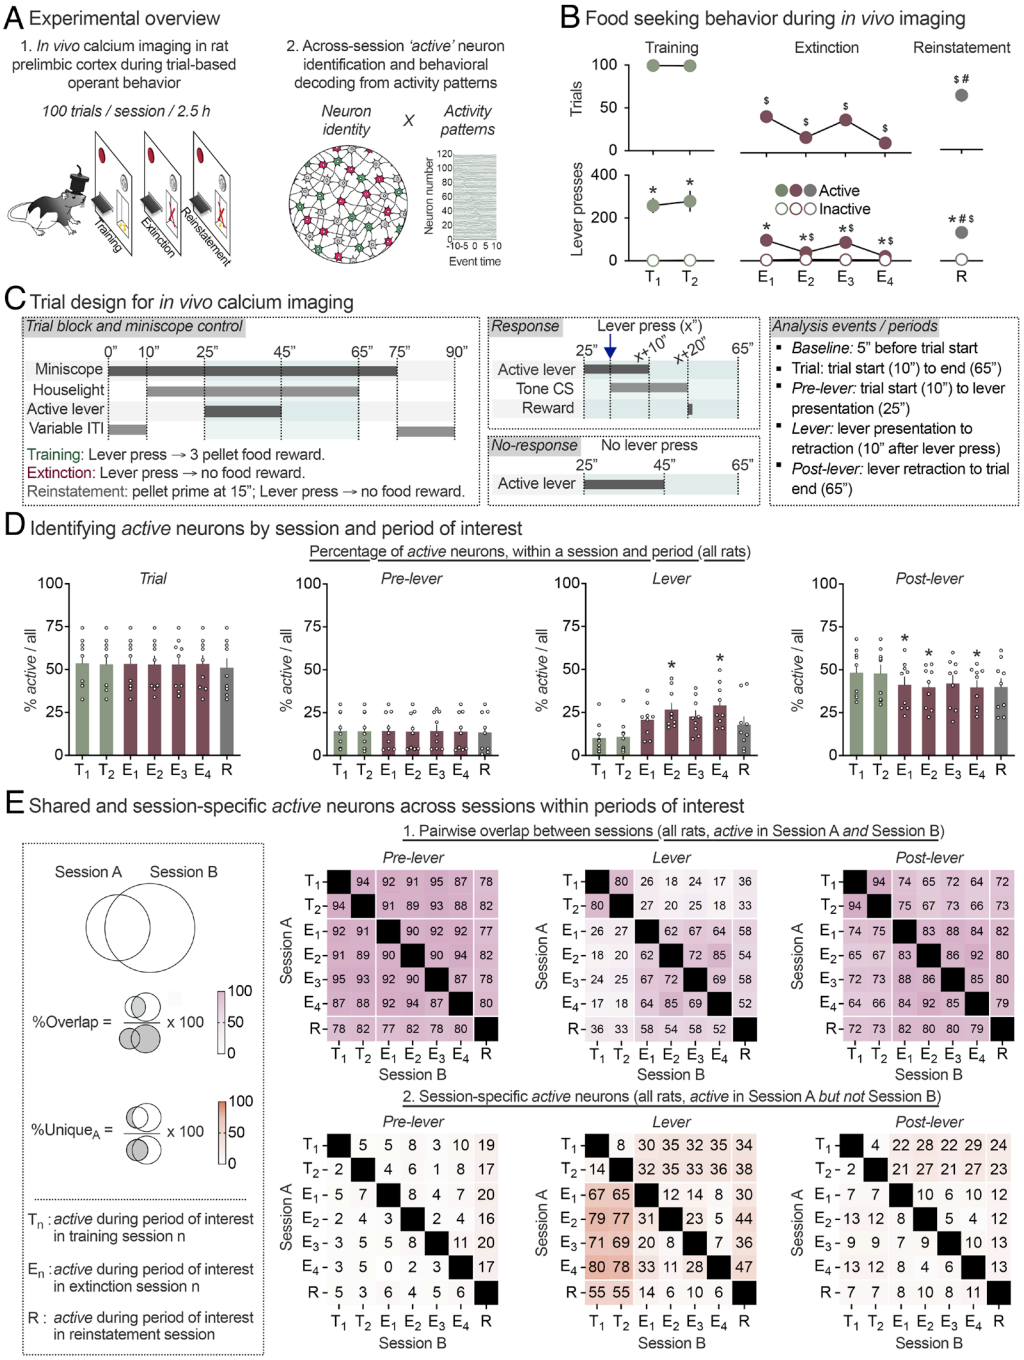

Fig. 1. In vivo calcium imaging allows longitudinal tracking of neuronal activity patterns and *active* neuron overlap in the rat PL during palatable food-seeking behavior. (*A*) *Experimental overview.* Rats were trained to self-administer palatable food pellets using a trial-based operant procedure prior to longitudinal in vivo calcium imaging during palatable food self-administration [Training (T), two sessions], extinction of palatable food seeking [Extinction (E), four sessions], and reinstatement of palatable food-seeking behavior by presentation of a priming pellet [Reinstatement (R), one session]. Neuron spatial and temporal footprints were detected using imaging data concatenated across all seven imaging sessions. (See *SI Appendix*, Fig. S1 for detailed experimental timelines, behavioral training prior to imaging sessions, GRIN lens placements, and example imaging field of view.) (*B*) *Palatable food seeking behavior during in vivo calcium imaging sessions (n = 9).* Mean (±SEM) number of response trials (top row), and active and inactive lever presses (bottom row) during training (two sessions, *Left* panel), extinction (four sessions, *Center* panel), and pellet-primed reinstatement (one session, *Right* panel). *Significant difference ($P < 0.05$) between active and inactive lever presses. $^{\$}$Significant decrease ($P < 0.05$), vs. training session 1 in number of response trials (*Top*) and active lever presses (*Bottom*). $^{\#}$Significant increase ($P < 0.05$), vs. extinction session 4 in number of response trials (*Top*) and active lever presses (*Bottom*) during reinstatement. (*C*) *Trial design and analyzed events.* Temporal sequence of events under experimenter control (*Left*), or dependent on rats' behavioral response (*Middle*, *Top*) or no-response (*Middle*, *Bottom*). Description of trial periods analyzed (*Right*). (*D*) *Percentage of active neurons by trial period and imaging session.* For each rat, spikes were estimated from calcium transients and used to calculate trial-by-trial average firing rate of individual neurons (all neurons detected across all seven sessions) within each trial period. For each session and trial period, *active* neurons were identified as those with significantly different average firing rates during that period vs. baseline period at the start of the same trial in the session. Mean (±SEM) percentage of *active* neurons identified during the entire trial (*Left*), pre-lever period (*Center Left*), lever availability period (*Center Right*), and post-lever period (*Right*), for each imaging session. Clear circles represent data from individual rats ($n = 9$). *Significant difference ($P < 0.05$) in percentage of *active* neurons vs. Training session 1. (See *SI Appendix*, Fig. S2 for *active* neuron counts for each session and period, and *SI Appendix*, Fig. S3 for *active* neuron counts by session type.) (*E*) *Shared and session-specific active neurons across sessions within trial periods.* Percentage pairwise shared 1) and session-specific 2) *active* neurons were calculated for each rat, and trial period as shown in the schematic (left). Heatmaps show mean pairwise %Overlap 1) and %Unique (2) *activ* neurons for pre-lever period (*Left*), lever period (*Middle*), and post-lever period (*Right*), for all seven imaging sessions. (See *SI Appendix*, Fig. S3 for heatmaps of %Overlap and %Unique neurons for pair-wise comparisons between three session types.)

not across the two session types, only during the lever period and not the pre-lever and post-lever periods. Further, the percentage of unique neurons was significantly increased in the Extinction (vs. Training) session during the lever period. This suggests that Extinction session *active* lever period (but not pre-lever or post-lever period) neurons comprise two distinct populations 1) reactivated Training session *active* lever period neurons (i.e., neurons *active* during both Training and Extinction sessions) and 2) a new population of Extinction session-specific *active* lever period neurons (i.e., neurons *active* during Extinction sessions but not previously *active* during Training).

For the pre-lever period, we observed similar levels of high %Overlap within and across training and extinction session pairs (80 to 95%; Fig. 1 *E-1*, *Left*; *SI Appendix*, Fig. S3 *B*, *Left*), and a low proportion of $\%Unique_A$ *active* neurons during this period, irrespective of session (Fig. 1 *E-2*, *Left*) or session-type (*SI Appendix*, Fig. S3 *C*, *Left*). For the lever period, we observed higher %Overlap within Training (80% for T1-2) and Extinction (>60% for E1-4) sessions, lower %Overlap across the session types (21%; Fig. 1 *E-1*, *Middle* and *SI Appendix*, Fig. S3 *B*, *Middle*), and a higher proportions of $\%Unique_A$ *active* neurons when comparing pairs of sessions belonging to different session-types (e.g., T vs. E sessions) compared to within session-type pairs (e.g., T1 vs. T2, or E1 vs. E2–4; Fig. 1 *E-2*, *Middle* and *SI Appendix*, Fig. S3 *C*, *Middle*). Finally, for the post-lever period, we observed slightly higher %Overlap within Training and Extinction sessions compared to across the session types (Fig. 1 *E-1*, *Right* and *SI Appendix*, Fig. S3 *B*, *Right*) and a low proportion of $\%Unique_A$ *active* neurons during this period, irrespective of session (Fig. 1 *E-2*, *Right*) or session-type (*SI Appendix*, Fig. S3 *C*, *Right*).

Overlapping vs. unique neurons in Reinstatement vs. Training or Extinction sessions. For the pre-lever period, we observed similar levels of %Overlap with the Reinstatement session across all training and extinction sessions (~80%; Fig. 1 *E-1*, *Left* and *SI Appendix*, Fig. S3 *B*, *Left*). For the lever period, we found higher %Overlap between Reinstatement and Extinction (~55%), compared to Reinstatement and Training (~35%) sessions, and overall lower %Overlap of either session type with Reinstatement compared to within Training and Extinction session-types (60 to 80%; Fig. 1 *E-1*, *Middle* and *SI Appendix*, Fig. S3 *B*, *Middle*). Finally, for the post-lever period, we observed similar %Overlap when comparing Reinstatement to either Training or Extinction sessions (Fig. 1 *E-1*, *Right*) or session-types (*SI Appendix*, Fig. S3 *B*, *Right*). For the $\%Unique_A$ *active* neuron measure, we found higher $\%Unique_A$ *active* Reinstatement session neurons during the lever period vs pre- or post-lever periods, when comparing to Training (~55%) but not Extinction (~10%) sessions (Fig. 1 *E-2*, *Middle*) or session-types (51% vs. 2%; *SI Appendix*, Fig. S3 *C*, *Middle*).

Overall, the data suggest that many of the lever period neurons activated during *both* Training and Extinction sessions were reactivated during Reinstatement.

Overlapping vs. unique neurons across trial periods. Finally, we investigated the relationship between *active* neurons identified during the three trial periods. We calculated %Overlap (*SI Appendix*, Fig. S4*A*) between active neurons identified in the three periods within each session and saw low across-period overlap and for all imaging sessions. Further, the $\%Unique_A$ (*SI Appendix*, Fig. S4*B*) lever period active neurons (compared to the other two periods), stayed consistent across the seven imaging sessions, suggesting a stable ensemble representation during the lever availability period. In contrast, for the pre- and post-lever periods, $\%Unique_A$ neurons, were similar within session-types (i.e., T1–T2, or E1–E4), but differed across session types, likely reflecting different outcomes/behavioral patterns in these periods (e.g., pellet delivery, food port entry) between the session types.

Decoding Trial-Wise Behavioral Responses from Neuronal Activity. In the analyses above, we observed stable *active* neuron populations within Training and Extinction session types, and reactivation of many of these neurons during the Reinstatement session. These patterns were specific to the lever period, suggesting the possibility that activity patterns of PL neurons during this period might contain information about rats' operant response (or omission thereof) on the presented lever. To determine whether this is the case, we implemented rat-specific decoders to predict their response execution and inhibition behaviors following training and extinction learning, based on neuronal activity patterns. For each rat, we trained a binary linear decoder to predict trial-wise behavioral outcomes (response/no-response) from the vector of average firing rates of all neurons. We assessed two periods—the lever period (where we observed distinct Training and Extinction session *active* neuron ensemble patterns) and the pre-lever period (as a control where we did not observe these ensemble-like patterns). Within each period, we trained rat-specific decoders using a randomly selected subset of the data (75% of training and extinction trials) and tested the decoder on two test sets: 1) the 25% of training and extinction trials held out from decoder training (Fig. 2), and 2) 100% of trials in the reinstatement session (Fig. 3). Test set 1 evaluates whether the decoders generalize to unseen training/extinction trials. Test set 2 evaluates whether the decoders work in a new context—reinstatement session—and, therefore, whether the information they extract from the neural activity pattern is still present then. For each test set, we sampled equal numbers of response and no-response trials to ensure 50% chance level (i.e., chance accuracy = 0.5).

A linear decoder uses the average firing rate of each neuron during the lever period in each trial as input (heatmap shown in Fig. 2 *A*, *Left*). The decoder computes a weighted linear combination of the firing rates of all neurons (firing rates multiplied by the corresponding decoder weights are shown in Fig. 2 *A*, *Center Left*) to predict whether or not the rat will press the active lever: The prediction is yes if the linear combination is positive, or no if the linear combination is negative. Fig. 2 *A*, *Center Right*, shows the linear combination over every trial in the two training and four extinction sessions in the first column. The adjacent columns show, respectively, whether the lever was pressed in each trial, the decoder prediction, the true label (same as lever press), and the trials where the decoder made an error. Fig. 2 *A*, *Right*, breaks down decoder performance as a confusion matrix (% of trials where prediction matches the true response, or not, for each response), both for the example rat used in the rest of the figure (Fig. 2 *A-2*, left heatmap), or averaged across all rats (Fig. 2 *A-2*, right heatmap). Fig. 3*A* provides the analogous plots for all the trials in the Reinstatement session. Individual decoder accuracy confusion matrices (split by rat and decoder training period) are provided in *SI Appendix*, Figs. S5 and S6. Detailed results for decoder accuracy on the two test sets are provided in *SI Appendix*, Tables S4 and S6, respectively.

Trial-wise decoding during Training and Extinction. The decoder trained on lever period activity exhibited overall high accuracy (Mean: 0.96, SEM: 0.005; Cohen's d = 30.5) with slightly higher accuracy for no-response prediction (Mean: 0.99, SEM: 0.001; Cohen's d = 90.8) vs. response prediction (Mean: 0.94, SEM: 0.008; Cohen's d = 18.1). In contrast, for the decoder trained on pre-lever period activity, overall accuracy was at chance level (Mean: 0.5, SEM: 0.01; Cohen's d = 0.02) and similar for no-response prediction (Mean: 0.62, SEM: 0.11; Cohen's d = 0.37) and response prediction (Mean: 0.38, SEM: 0.11; Cohen's d = −0.37).

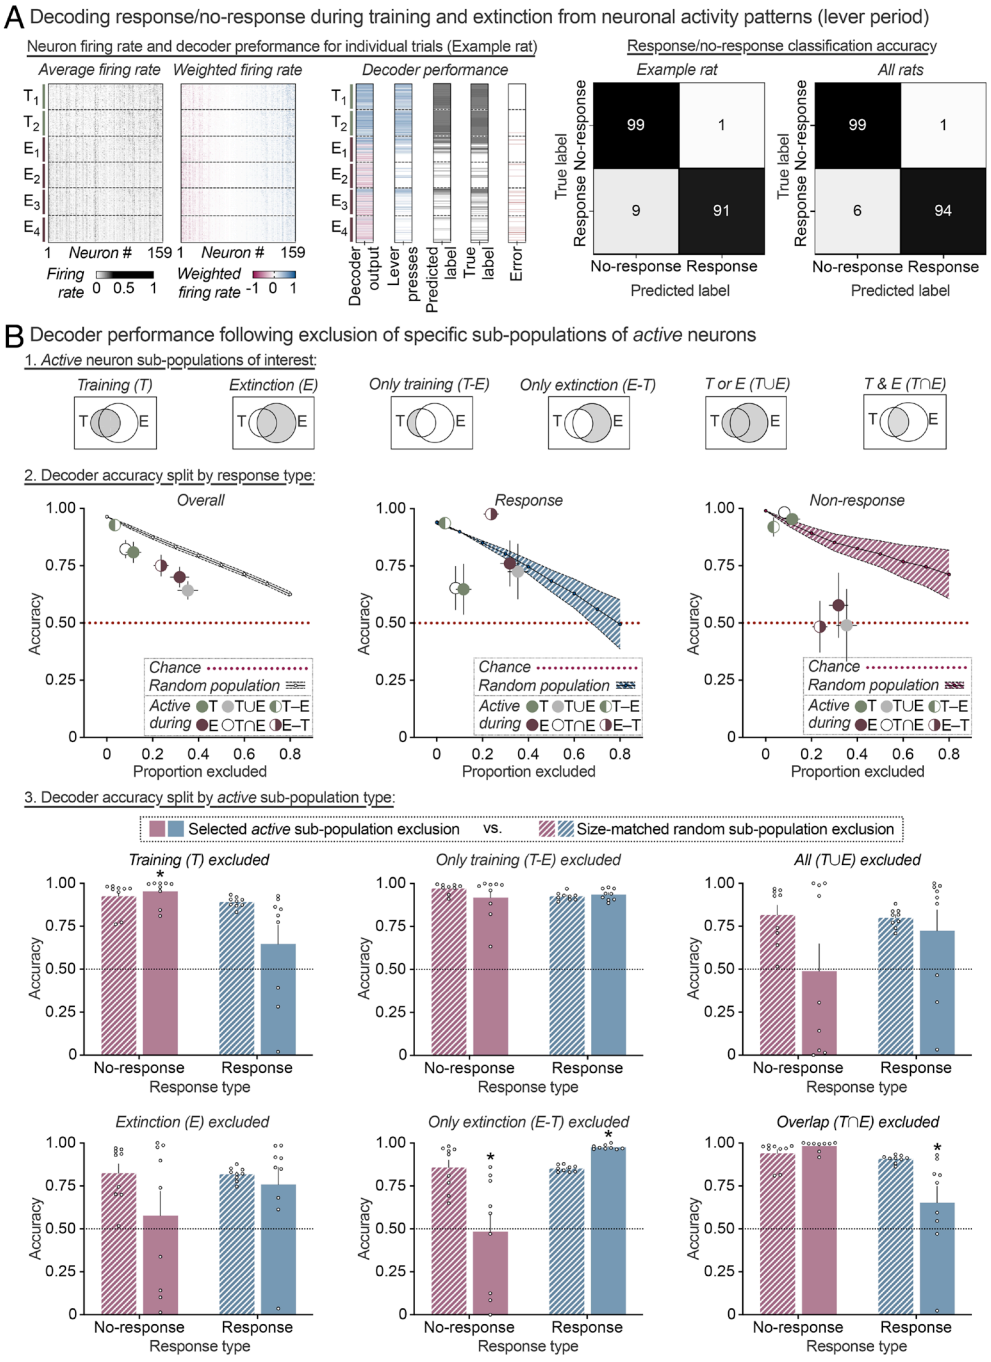

Fig. 2. Distinct prelimbic cortex activity patterns support response/no-response behavior during Training and Extinction sessions. (*A*) *Prelimbic cortex neuronal activity supports trial-by-trial behavioral decoding during Training and Extinction.* For each rat, a binary (response/no-response) linear decoder was trained using the average firing rate vector during lever availability period, for a randomly selected subset of training and extinction trials (75/25 stratified split between training and test trials). The decoder was then tested on the holdout subset of training and extinction trials and showed high decoding accuracy for all rats. Heatmaps for one example rat showing raw firing rate of each neuron (*Left*), firing rate multiplied by the corresponding decoder weight (*Center*) and decoder performance (*Right*) during the two training and four extinction sessions. Each gray pixel in the left plot represents the lever-period-average firing rate of one neuron (column) in one trial (row). Each colored pixel in the center plot represents the weighted firing rate (decoder weight × firing rate) of one neuron (column) in one trial (row). Sessions and trials are displayed in chronological order from top to bottom and neurons are sorted by weight from negative (red, *Left*) to positive (blue, *Right*). The decoder output column shows the sum of weighted firing rates across all neurons, which determines the decoder prediction (blue = response, red = no response). The lever press column shows number of responses by trial. The predicted label column shows binarized decoder output by trial (black tick = response predicted). The true label column shows binarized response pattern by trial (black tick = response). The error column shows disagreement between predicted and true labels (red tick = error). For each rat (n = 9), an equal number of response and no-response trials were sampled to set chance accuracy (red dotted line) at 0.5. Decoder accuracy (%true label) on the holdout test set (training and extinction) split by response and shown as a confusion matrix for one example rat (left heatmap) or averaged across all rats (right heatmap, n = 9). (*B*) *Distinct subsets of active Training and Extinction session neurons support response and no-response prediction accuracy.* For each rat, lever period decoder accuracy was measured following exclusion of one of six combinations of Training and Extinction session-specific *active* neuron populations (top row). Note that the decoder was trained prior to the exclusion; the weights of excluded neurons are set to 0. Summary plots (row two, all rats) of Mean (±SEM) overall accuracy (*Left*), response prediction accuracy (*Center*), and no-response prediction accuracy (*Right*) following exclusion of each *active* population (colored circles) vs. random population exclusion (line plot, 0 to 80% of all neurons per rat). Bar charts (third and fourth rows, split by population) of Mean (±SEM) response (blue) and no-response (red) prediction accuracy after exclusion of *active* population of interest (solid bars), vs. exclusion of a random population of the same size for each rat (patterned bars). Clear circles represent data from individual rats (n = 9). *Significant difference (P < 0.05) in response or no-response prediction accuracy between specific population exclusion and size-matched random population exclusion.

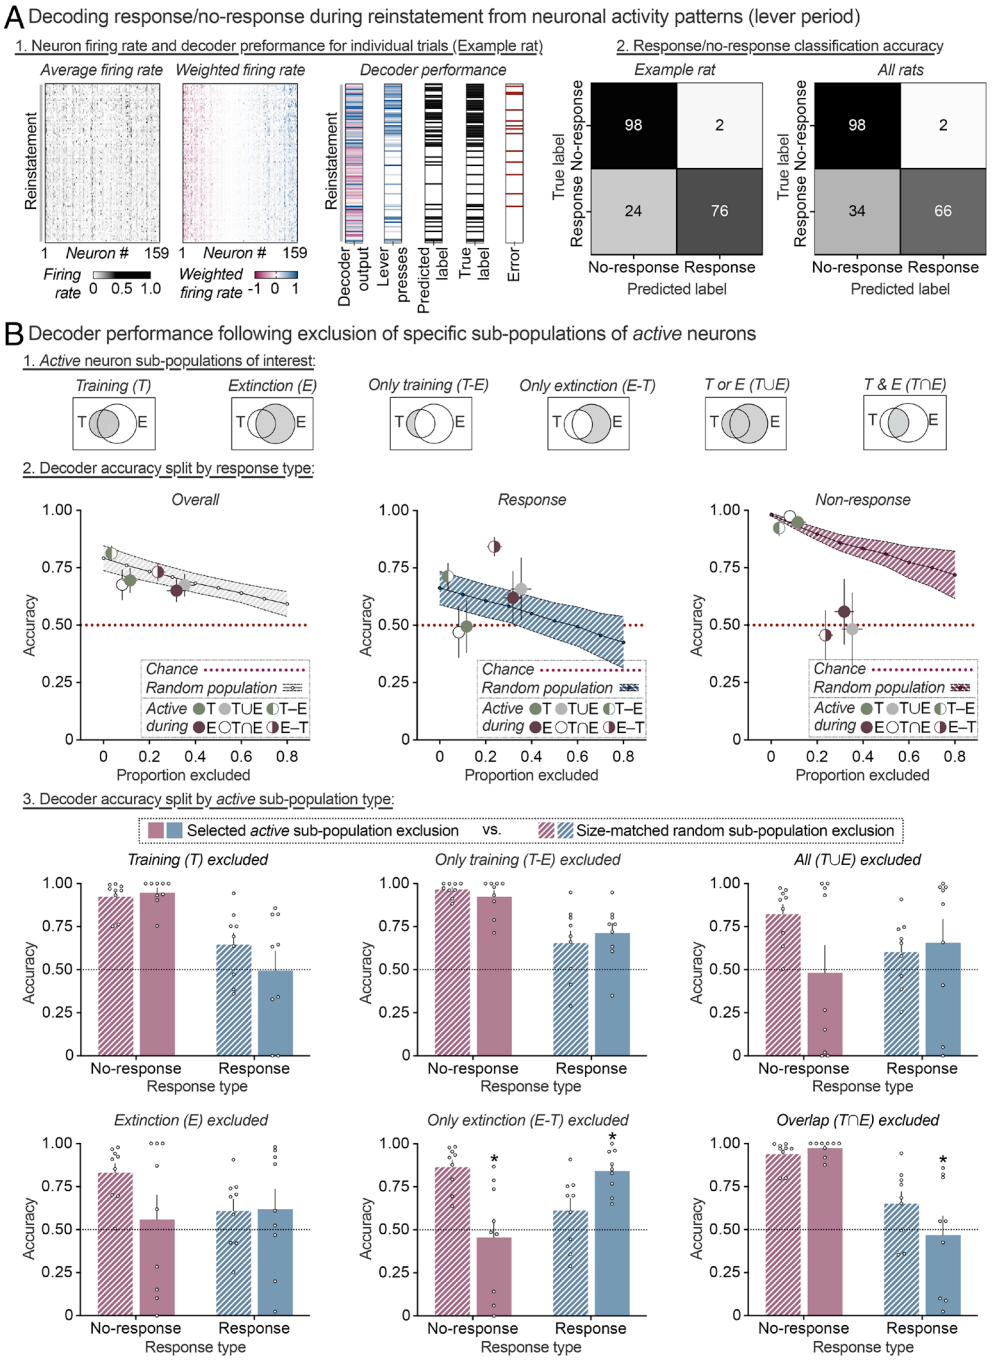

Fig. 3. Prelimbic cortex Training and Extinction neuron activity patterns support response/no-response behavior during pellet-primed reinstatement. (*A*) *Prelimbic cortex neuronal activity supports trial-by-trial behavioral decoding during Reinstatement.* For each rat, the response/no-response decoder trained using training and extinction session trials was tested on the reinstatement session trials. Heatmaps for one example rat showing raw firing rate of each neuron (*Left*), firing rate multiplied by the corresponding decoder weight (*Center*) and decoder performance (*Right*) during the reinstatement session. Each gray pixel in the left plot represents the lever-period-average firing rate of one neuron (column) in one trial (row). Each colored pixel in the center plot represents the weighted firing rate (decoder weight × firing rate) of one neuron (column) in one trial (row). Trials are displayed in chronological order from top to bottom, and neurons are sorted by weight from negative (red, *Left*) to positive (blue, *Right*). The decoder output column shows the sum of weighted firing rates across all neurons, which determines the decoder prediction (blue = response, red = no response). The lever press column shows number of responses by trial. The predicted label column shows binarized decoder output by trial (black tick = response predicted). The true label column shows binarized response pattern by trial (black tick = response). The error column shows disagreement between predicted and true labels (red tick = error). For each rat (n = 9), equal number of response and no-response trials were sampled from the reinstatement session to set chance accuracy at 0.5. Decoder accuracy (%true label) on reinstatement trials split by response and shown as a confusion matrix for one example rat (left heatmap) or averaged across all rats (right heatmap, n = 9). (*B*) *Distinct subsets of active Training and Extinction session neurons support response and no-response prediction during Reinstatement.* For each rat, lever period decoder accuracy was measured following exclusion of one of six combinations of Training and Extinction session-specific *active* neuron populations (top row). Summary plots (row two, all rats) of Mean (±SEM) overall accuracy (*Left*), response prediction accuracy (*Center*), and no-response prediction accuracy (*Right*) following exclusion of each *active* population (colored circles) vs. random population exclusion (line plot, 0 to 80% of all neurons per rat). Bar charts (third and fourth rows, split by population) of Mean (±SEM) response (blue) and no-response (red) prediction accuracy after exclusion of *active* population of interest (solid bars), vs. exclusion of a random population of the same size for each rat (patterned bars). Clear circles represent data from individual rats (n = 9). *Significant difference ($P < 0.05$) in response or no-response prediction accuracy between specific population exclusion and size-matched random population exclusion.

Trial-wise decoding during Reinstatement. The decoder trained on lever period activity also had overall high accuracy during reinstatement testing (Mean: 0.79, SEM: 0.0.055; Cohen's d = 1.8) with slightly higher accuracy for no-response prediction (Mean: 0.9, SEM: 0.08; Cohen's d = 1.65) vs. response prediction (Mean: 0.74, SEM: 0.059; Cohen's d = 1.37). In contrast, for the decoder trained on pre-lever period activity, overall accuracy was at chance level (Mean: 0.54, SEM: 0.05; Cohen's d = –0.38) and similar for no-response prediction (Mean: 0.6, SEM: 0.12; Cohen's d = 0.37) and response prediction (Mean: 0.38, SEM: 0.1; Cohen's d = 0.28).

In agreement with the previous overlapping and unique neurons matrices, these decoding analyses indicate that stable PL activity patterns are established during the lever period as the rats learn to make, and suppress, responses over the course of operant Training and Extinction learning. These patterns can be used to predict response execution or inhibition during individual trials during learning (in the Training and Extinction sessions), and the decodability is maintained when rats reengage with the lever during the Reinstatement session. In contrast, it appears these stable patterns are not present during the pre-lever period.

Contribution of Active Neuron Subpopulations to Response/No-Response Prediction. In the previous section, we established the existence of distinct, stable patterns of activity during the lever period that are associated with response execution or inhibition, respectively. Building on these results, the next question we posed is which specific subpopulations of *active* Training and Extinction session neurons contain the information that the decoder uses in predicting response behavior with high accuracy. For each rat, we assessed the contribution of six subpopulations of Training and Extinction session-specific *active* neurons to lever period decoder performance during training and extinction (Fig. 2*B*) or reinstatement (Fig. 3*B*). The six subpopulations of *active* neurons considered were 1) Training (T)—active in any of the two training sessions; 2) Extinction (E)—active in any of the four extinction sessions; 3) Only Training (T-E)—active in training but not extinction sessions; 4) Only Extinction (E-T)—active in extinction but not training sessions (i.e., the putative *Extinction* ensemble); 5) All T or E (T∪E)—active in either training, or extinction sessions; and 6) Overlap between T & E (T∩E)—active in both training and extinction sessions (i.e., the putative *Training* ensemble). In this analysis, we assumed neurons active in any one session within a given session type (Training or Extinction) contributed similarly to behavior and were interchangeable with those active in other sessions of the same type. This assumption was based on our population-level analysis, which revealed 1) high overlap between the two training sessions and between the four extinction sessions and 2) similar %Overlap between each training session and any of the extinction sessions, and similar % overlap between each extinction session and any other session (Training or Extinction).

We used an in silico *lesion* approach to test the contribution of each of the six subpopulations to trial-wise response/no-response prediction accuracy of the previously trained lever period decoder. To this effect, we removed specific subpopulations from decoder prediction by setting their corresponding decoder weights to 0 when computing the linear combination of neuron firing rates and calculated the accuracy of the resulting decoder (*specific-population-exclusion accuracy*). For a comparison controlling for accuracy changes based on subpopulation size, we also calculated the accuracy of a decoder obtained after exclusion of a randomly chosen size-matched sets of neurons (*random-population-exclusion accuracy*). Summary plots of overall accuracy (*Left*), response prediction accuracy (*Center*), and no-response prediction accuracy (*Right*) following exclusion of each active subpopulation (colored circles) vs. random population exclusion (0 to 80% of all neurons per rat, line with error band) are shown in Figs. 2 *B-2* and 3 *B*-2.

We first used two-way repeated measures ANOVA with within-subjects' factors of *Population* (1–6) and *Drop type* (specific, random) to test whether the subpopulations have enhanced contributions to overall decoder prediction compared to randomly selected PL neurons. Then for each subpopulation, we used two-way repeated measures ANOVA with within-subjects' factors of *Drop type* (specific, random) and *Response type* (response, no-response) to assess whether subpopulation exclusion resulted in biased effects on response vs no-response prediction accuracy (Figs. 2 *B-3* and 3 *B-3*). Detailed results following exclusion of specific *active* neuron subpopulations are provided in *SI Appendix*, Tables S5 and S7, respectively.

Prediction accuracy during Training and Extinction. Overall decoder accuracy (Response and No-response prediction combined) after excluding each of the six *active* neuron subpopulations was lower than size-matched random subpopulation exclusion (Fig. 2 *B-2*, *Left*). The overall two-way ANOVA showed a significant *Population* × *Drop type* interaction ($F_{1.96,15.7}$ = 2.55, P = 0.043) and post hoc analysis revealed significant accuracy drop for four of the six *active* subpopulations tested (E, T∪E, T∩E, and E-T; see *SI Appendix*, Tables S4 for detailed statistical outputs). When assessing decoder accuracy separately for Response vs. No-response types (Fig. 2 *B-2*, *Middle* and *Right*), two-way ANOVA showed significant *Response type* × *Drop* type interaction for three of the six subpopulations tested (T, T∩E, and E-T; see *SI Appendix*, Tables S5 for detailed statistical outputs). Exclusion of neurons belonging to the putative *Training* ensemble [Overlap (T∩E) subpopulation; Fig. 2 *B-3*, *Bottom Right*] decreased decoder accuracy to chance levels, specifically for response trials (P = 0.029). In contrast, exclusion of neurons belonging to putative Extinction ensemble [Only Extinction (E-T) subpopulation; Fig. 2 *B-3*, *Bottom Middle*] decreased decoder accuracy to chance levels for no-response trials (P < 0.001) and showed a moderate increase in decoder accuracy for response trials (P < 0.001).

Prediction accuracy during Reinstatement. Overall decoder accuracy (Response and No-response prediction combined) after excluding each of the six *active* neuron subpopulations was lower than size-matched random subpopulation exclusion (Fig. 3 *B-2*, *Left*). The overall two-way ANOVA showed a main effect of *Population* ($F_{1.64,13.1}$ = 3.85, P = 0.006) but not *Drop type* ($F_{1,8}$ = 2.7, P = 0.14), or interaction ($F_{2.04,\ 16.3}$ = 2.87, P = 0.085) between the factors (*SI Appendix*, Tables S6). When assessing decoder accuracy separately for Response vs. No-response types (Fig. 3 *B-2*, *Middle* and *Right*), the two-way ANOVA showed significant *Response type* × *Drop* type interaction for two of the six subpopulations tested (T∩E, and E-T; see *SI Appendix*, Tables S7). Again, while exclusion of neurons belonging to the putative *Training* ensemble [Overlap (T∩E) subpopulation; Fig. 3 *B-3*, *Bottom Right*] decreased decoder accuracy to chance levels, specifically for response trials (P = 0.012), exclusion of neurons belonging to putative Extinction ensemble [Only Extinction (E-T) subpopulation; Fig. 3 *B-3*, *Bottom Middle*] decreased decoder accuracy to chance levels for no-response trials (P < 0.001) and showed a moderate increase in accuracy for response trials (P < 0.001).

Overall, the decoding analyses indicate that *active* lever period neurons were significantly associated with Response/No-response (compared to similar-sized randomly selected neuron populations) both during Training and Extinction sessions, and during Reinstatement. Further, the activity patterns of PL *Training* and *Extinction* ensembles supported opposite response predictions—*Training* ensemble activity

supported response execution and Extinction ensemble activity supported response inhibition on individual trials, across Training, Extinction, and Reinstatement.

Discussion

We used single-cell resolution calcium imaging to longitudinally monitor the activity patterns of hundreds of PL neurons in rats during trial-based operant food reward training, extinction, and reinstatement. We tested whether the same PL ensemble neurons adapt their activity patterns to support training, extinction, and reinstatement of reward-seeking behavior, or if distinct, nonoverlapping ensembles are recruited to mediate the appropriate action during each phase. Population-level analyses of overlapping vs. unique *active* neurons across the Training and Extinction sessions revealed higher overlap for lever period *active* neurons within the same session type than between them; this was not the case for the pre-lever or post-lever periods. Additionally, many *active* lever period neurons from Training were reengaged during Extinction, and a new population of *active* neurons emerged specifically during Extinction.

We trained rat-specific decoders to predict trial-wise response/no-response behavior based on PL *activity patterns* during training and extinction. Decoders trained on lever period activity patterns accurately predicted individual rats' trial-wise response (execution) vs. no-response (inhibition) behavior; this was not the case for the pre-lever period.

Next, we tested the contribution of specific subpopulations of *active* Training and Extinction session neurons by omitting them from the decoder (i.e., an in silico targeted deletion). Omission of lever-period neurons *active* during both Training and Extinction sessions (T∩E) selectively reduced response prediction accuracy to chance levels, suggesting these represent a *Training ensemble*, while omission of lever-period neurons *active* only during Extinction (E-T) disrupted no-response prediction, suggesting these represent an *Extinction ensemble*.

Finally, these opposite decoder prediction impairments after *Training* or *Extinction ensemble* omission persisted during Reinstatement, indicating that both ensembles are reengaged and maintain their respective roles. Overall, the results of our analyses indicate that operant training and extinction learning recruit distinct PL ensembles to support behavioral flexibility in response to changing reward contingencies.

Operant Training and Extinction recruit distinct PL ensembles. Previous in vivo recording studies have demonstrated diverse PL activity patterns related to cues, response execution/inhibition, and outcome/reward (5, 21, 24, 25, 35), and identified stable PL neuron populations (ensembles) during both operant and Pavlovian reward learning (14, 26, 29, 36). Consistent with these findings, we observed diverse PL neuron activity patterns during Training sessions, with similar numbers and high overlap of *active* neurons across sessions, regardless of trial period. However, contrary to the PL-go/IL-stop hypothesis' prediction of reduced PL engagement during Extinction, the number of *active* neurons increased in Extinction sessions. Notably, this increase was specific to the lever availability period and there was higher overlap in lever period *active* neurons within each session type than between Training and Extinction sessions, indicating that different ensembles encode Training vs. Extinction in the same PL subregion.

Furthermore, decoders trained on PL lever period activity patterns consistently predicted individual rats' response execution and inhibition across Training and Extinction sessions, suggesting that the PL does not simply support responding during reward availability and "go offline" during Extinction. These results align with previous in vivo electrophysiological recording studies using a discriminative stimulus (DS) procedure for sucrose (21, 22) and cocaine reward, which found that neuronal activity in PL and IL is not specifically locked to the DS or the action (i.e., going or stopping), but rather reflects the appropriate behavioral action based on context (e.g., responding during DS+ vs. inhibiting responses following extinction training).

Surprisingly, lever period *active* neurons from the Training session remained active during the Extinction session. Further, shared and overlapping neuron analysis between Training and Extinction sessions revealed that a majority of lever period active neurons from the Training session were *also active* during Extinction (T∩E), along with a second population *exclusively active* during Extinction (E-T). One possibility is that both populations are simultaneously active and that the extinction only (E-T) population suppresses response execution mediated by the reactivated Training population (T∩E). However, the response-specific bias in decoder prediction accuracy following targeted omission of these two subpopulations suggests another explanation.

Omission of the T∩E subpopulation alone selectively reduced decoder *response* prediction accuracy to chance levels, while omission of the E-T subpopulation specifically disrupted *no-response* prediction accuracy. This suggests that these two populations are likely not active at the same time or in opposition to each other but instead come online at different times to support separate outcomes. In this perspective, the T∩E subpopulation would represent a *Training* ensemble that is formed during Training and supports response execution in pursuit of the reward. This ensemble would then be reactivated during *response trials* in the Extinction sessions when rats continue to make responses (to earn food reward based on the previous contingency) as they learn the new contingency (i.e., extinction learning). The second E-T subpopulation would then represent the new *Extinction* ensemble that is recruited as rats learn to suppress behavior on *no-response trials* during the Extinction session.

It is possible that extended extinction training might lead to a reduction in the size of the T∩E population as rats learn to make fewer responses in pursuit of reward. We used a within day extinction procedure and chose four extinction sessions (based on preliminary behavioral experiments) to minimize the number of recording days and maintain stable across-session recordings. While additional experiments with longer extinction periods are needed, in support of this idea, we previously found that the number of *active* IL extinction session neurons, measured using the activity marker Fos, decreases between 2 and 7 d of extinction training (12). Furthermore, targeted ablation (37) of *active* training session neurons reduced food lever responding, while ablation of the *active* extinction session neurons increased lever responding (12). Overall, our data align with recent studies showing the development of stable PL ensembles (26, 29) during reward learning, and suggest that operant training and extinction learning recruit distinct PL ensembles whose *activity patterns* support behavioral flexibility in response to changing reward contingencies. We hypothesize that these distinct *Training* and *Extinction* ensembles are engaged by different patterns of excitatory inputs, responsive either to the expectation of food reward during training or to reward omission during extinction.

PL Training and Extinction Ensembles Are Reengaged During Reinstatement. *Training* and *Extinction* ensembles were reactivated during pellet-primed reinstatement and maintained their decoder selectivity for response and no-response prediction. At the population level, we observed that lever period *active* neurons from both Training and Extinction session were reactivated during the reinstatement

session (although to different degrees), and very few neurons were uniquely active during the same period in reinstatement. Notably, although the rat-specific decoders were never trained on reinstatement activity patterns, lever period (but not pre-lever period) decoder performance was significantly above chance levels when tested on reinstatement trials, suggesting that PL activity likely supports whether the rat makes or withholds the operant response (the food-seeking response) during reinstatement. Further, omission of active neuron subpopulations increased decoder error, and omission of the *Training* (T∩E) or *Extinction* (E-T) ensembles oppositely biased decoder prediction errors toward response vs. no-response, respectively. These findings suggest that behavior (i.e., whether the rat makes an operant response) during reinstatement is not driven by a distinct third ensemble but involves preferential recruitment of one of two response-outcome-specific ensembles acquired during operant training and extinction. The *Training* ensemble is engaged when rats reinstate lever pressing in response to the acute noncontingent reward provided by the priming pellet, while the *Extinction* ensemble supports response inhibition as they suppress responding when lever presses fail to result in food pellet delivery.

Methodological Limitations and Future Directions. In our study, we used only one behavioral measure — responses on the active lever — which limited our ability to identify ensembles encoding other behavioral features in our experiment. For example, it is possible that Training and Extinction lever-period ensembles also encode food receptacle approach, consummatory responses, or lever approach during no-response trials. The high number of active neurons during the post-lever period suggests additional PL ensembles may be engaged during this period to encode consummatory behaviors. While the current procedure revealed recruitment of a separate Extinction ensemble, future studies should incorporate trial-based operant discrimination procedures (21, 38–41) to test selective recruitment of ensembles during reinforced vs. nonreinforced trials *within the same session*. Integrating these complex tasks with our ensemble-level analyses and markerless pose-estimation and behavioral tracking approaches (29, 42, 43) will be critical for fully capturing the full contribution of these—and potentially additional—ensembles to appetitive and consummatory learned behaviors.

A key question raised by our findings is whether the same or distinct cellular and circuit mechanisms underlie Training vs. Extinction ensemble recruitment. Previous studies have identified glutamatergic synaptic plasticity (15, 44, 45), changes in intrinsic excitability (46, 47), and ensemble-specific cellular and synaptic alterations (48, 49) that likely support recruitment and reengagement of ensembles during learning (45).

Since Extinction ensembles reflect a second, although opposing, learned association, similar mechanisms likely govern their recruitment—potentially via a distinct pattern of excitatory inputs engaged by reward omission—to suppress the previously learned behavior. To test this, future studies should combine activity-dependent labeling approaches (50–53) with targeted single cell/nucleus transcriptomics (54–56), electrophysiological characterization (45–47), and synaptic profiling (49) to define the molecular identity and functional properties of these two ensembles. In parallel, brain-wide ensemble mapping (57) will be critical to determine whether these opposing ensembles are unique to PL, or part of a broader pattern of selective ensemble engagement across other regions (e.g., IL, ACC, BLA, NAc) involved in operant training, extinction, and reinstatement.

Conclusions. In summary, we show that two distinct PL ensembles are recruited over the course of operant reward training and extinction to support opposite behavioral responses, and these same ensembles are reengaged when rats reinstate operant behavior in response to a priming pellet. Our findings provide neuronal ensemble-based evidence supporting the idea that operant extinction reflects new learning, rather than a weakening of the response-outcome association (58, 59). Consistent with this view, we show that the reinstatement of operant behavior after extinction relies on the reengagement of preestablished ensembles rather than the recruitment of a distinct third ensemble. Overall, our findings challenge the *one primary role per region* premise of the PL-go/IL-stop hypothesis and demonstrate how the interplay of preestablished and newly recruited ensembles within reward-related brain regions enables behavioral flexibility in response to changing reward contingencies.

Materials and Methods

We combined a trial-based operant behavioral procedure with in vivo calcium imaging to investigate PL neuron activity dynamics during different stages of palatable food taking and seeking behavior in rats. A schematic overview of the study design is shown in Fig. 1*A*. Timelines for surgical procedures, behavioral training, and imaging sessions are provided in *SI Appendix*, Fig. S1 *A* and *B*. We first trained rats to lever press for 45 mg palatable food pellets (32) using a trial-based procedure. We then used miniature epifluorescent microscopes to track PL activity patterns during operant food self-administration (Training), as they learned to suppress responding when the reward was omitted (Extinction), and when they reinstated responding after receiving a noncontingent priming pellet (Reinstatement). We kept the microscopes attached for the entire duration of the imaging study (seven sessions over 4 d) to maintain the same imaging field of view and allow for longitudinal tracking of PL neuron activity patterns. A brief description of experimental subjects, surgery, behavioral and imaging procedures and apparatus, data analysis pipelines, and statistical analyses is provided below, and additional details pertaining to experimental subjects, surgical procedures, behavioral apparatus, and image processing pipelines are described in *SI Appendix*.

Subjects. We used male Long Evans rats (Charles River, Raleigh, n = 24), weighing 300 to 350 g at the start of the experiment. All procedures (see *SI Appendix* for details) were approved by the NIDA IRP Animal Care and Use Committee and followed the guidelines outlined in the Guide for the Care and Use of Laboratory Animals (60). We excluded 4 rats based on initial training performance, 10 rats based on poor imaging field of view, and 1 rat due to hardware failure and excessive motion artifacts. We report the number of rats included in each analysis in the corresponding figure legend.

Surgery.

Intracranial virus surgery. We performed intracranial surgery (see *SI Appendix* for details) to deliver adeno-associated viruses (AAVs) to express the genetically encoded calcium indicator GCaMP6 (AAV2/1-hSyn1-GCaMP6s, 1.9E13 GC/mL) in the prelimbic cortex (unilateral, right hemisphere) and allowed rats to recover for at least 7 d before initial behavioral training.

Intracranial GRIN lens implantation surgery. Following initial behavioral training, we performed intracranial surgery (see *SI Appendix* for details) to implant a 1 mm diameter, 0.95 pitch graded refractive index (GRIN) lens (ILW-100-P, Go!Foton) above the prelimbic cortex (unilateral, right hemisphere). We placed a custom protective cap above the GRIN lens and allowed rats to recover for up to 3 wk before continuing behavioral training. We removed the custom protective cap once every two weeks and imaged under isoflurane anesthesia using a custom epifluorescent microscope to monitor scar tissue regression and determine the appropriate time for miniature microscope optical alignment and baseplate installation.

Miniscope alignment and baseplate installation. Following scar tissue regression, we performed a final noninvasive procedure (see *SI Appendix* for details) to align the Miniscope to the GRIN lens and installed a custom threaded baseplate for stable optical coupling during behavior.

Habituation to head-mounted Miniscopes. Following recovery, we used dummy Miniscopes that had the same weight and physical tethering properties of the recording Miniscope to habituate rats to imaging equipment during behavior. We continued tethered training with dummy scopes till rats showed stable food self-administration before starting the imaging sessions.

Miniscope attachment for imaging. One day before the first imaging session, we lightly anesthetized rats using isoflurane gas to align the Miniscope sSee *SI Appendix* for details). We allowed rats to recover from anesthesia, then tested the Miniscope + acquisition equipment for optical alignment and stability as animals freely explored the behavior chambers. After confirming alignment, we disconnected the data acquisition cable and housed the rats in the behavior chambers for the rest of the experiment. For each rat, we kept the Miniscope position and alignment fixed for all seven imaging sessions and attached/disconnected the data acquisition cable before/after the session to minimize cable damage.

Behavioral Apparatus. We used standard Med Associates self-administration chambers for palatable food (TestDiet, USA; Catalog # 1811155, 12.7% fat, 66.7% carbohydrate, and 20.6% protein) self-administration training and for in vivo calcium imaging. Each chamber was enclosed in a ventilated, sound-attenuating cabinet with blacked out windows and equipped with a stainless-steel grid floor, a red houselight, a retractable active lever, a nonretractable inactive lever, a pellet dispenser, and a food receptacle (see *SI Appendix* for details).

Behavioral Procedures. We trained rats on a trial-based palatable food self-administration procedure and used in vivo calcium imaging to monitor activity dynamics of PL neurons during the last two sessions of food self-administration (Training sessions; T1, T2), four consecutive sessions of extinction (Extinction sessions; E1–E4), and one session of palatable food primed reinstatement (Reinstatement session; R) behavior. Details of the individual stages of training and behavioral procedures during imaging are described in *SI Appendix*.

Analysis of Behavioral Data. We conducted statistical analysis on two behavioral measures–1) the total number of trials with at least one food-paired lever press (denoted as *Trials*) and 2) the total number of active and inactive lever responses made during the session (denoted as *Lever presses*). We used repeated measures ANOVAs with Greenhouse–Geisser correction as appropriate and followed up on statistically significant main effects or interactions with post hoc tests as described below. Alpha (significance) level was set at 0.05, two-tailed. Because some of our models yielded multiple main effects and interactions, we report only those that are critical for data interpretation.

For trial-based palatable food self-administration training sessions prior to imaging, we performed statistical analysis on each phase (Initial training, and postsurgery training phases 1, and 2) separately. We analyzed *Trials* using one-way ANOVA with within-subject factor of training session and *Lever presses* using two-way ANOVA with within-subject factors of training session and lever type (active, inactive). We used Bonferroni's multiple comparison correction for pairwise comparisons between active and inactive lever presses within a training session.

For imaged behavioral sessions (two training sessions, four extinction sessions, one reinstatement session), we analyzed *Trials* using one-way ANOVA with within-subject factor of session (seven sessions) and *Lever presses* using two-way ANOVA with within-subject factors of session (seven sessions) and lever type (active, inactive). For both measures we used Bonferroni's multiple comparison correction for pairwise comparisons between imaging sessions.

Calcium Imaging Data Processing. We concatenated the imaging data across all imaged sessions (7 sessions × 100 trials/session × 799 images/trial for each rat), applied a sequence of image processing routines to extract fluorescence time courses and spatial footprints of single neurons and deconvolved each neuron's temporal component to estimate spiking (see *SI Appendix* for details).

Calcium Imaging Data Analyses.

Identification of "active" neurons by session. For each rat, we used spike data to calculate average firing rates of all neurons detected across all seven sessions. We computed separate average firing rates within the following periods of interest–1) trial period: start to end of behavioral trial, 2) pre-lever period: trial start to lever presentation, 3) lever period: lever presentation to retraction, and 4) post-lever period: lever retraction to trial end. For each session, *a* neuron was deemed *active* in a period if its firing rate within that period was significantly different from that of the baseline period (5″ window prior to start of the trial) using the two-sided paired *t* test. We corrected for multiple comparisons using Benjamini and Hochberg's false discovery rate (FDR) method and set FDR at 0.05.

Quantification of shared vs. unique active neuron populations. For each rat, we identified *active* neurons in each period for all seven imaging sessions and used custom scripts to identify shared and unique *active* neuron populations between pairs of sessions. We used two measures to quantify the proportion of shared/unique *active* neuron populations between pairs of sessions (session A, session B):

(1) %Overlap = number of neurons that were active in *both* Sessions A and B/number of neurons that were active in Sessions A or B = $|A \cap B|/|A \cup B|$
(2) $\%Unique_A$ = number of neurons that were active *only in* Session A *but not* Session B/number of neurons that were active in Session A = $|A - B|/|A|$

For each rat, we also combined identified *active* neuron populations from the three session types (training, extinction, or reinstatement) and performed the same analysis, to identify the proportion of shared and unique *active* neuron populations between pairs of session-types.

Decoding trial-wise behavioral response from neuronal activity. For each rat, we trained linear decoders predicting the trial-wise behavioral outcome (response/no-response) from the vector of average firing rates of a set of neurons during a period of interest. We trained separate decoders for two distinct periods of interest: 1) lever period and 2) pre-lever period. For each rat, we trained each decoder on a randomly selected subset of the data containing 75% of training and extinction trials. We reported the metrics resulting from applying the decoder to two separate test sets: 1) the remaining held-out 25% subset of training and extinction trials not used in training and 2) all reinstatement trials unless otherwise stated. Note that no reinstatement trials were used in training the decoder. We resampled balanced samples (equal number of response and no-response trials) for the test set and reinstatement session to obtain a 50% chance level.

Contribution of Training and Extinction session active subpopulations to lever-period decoder accuracy. For each rat, we measured the lever period decoder accuracy in the absence of six specific *active* neuron subpopulations, to determine how much they contribute to decoder performance. The six subpopulations of *active* neurons considered were as follows:

1. Training (T)–*active* in the two training sessions
2. Extinction (E)–*active* in the four extinction sessions
3. Training exclusive (T-E)–*active* in training but not extinction sessions
4. Extinction exclusive (E-T)–*active* in extinction but not training sessions
5. T or E (T∪E)–*active* in training, or extinction sessions
6. T & E (T∩E)–*active* in training and extinction sessions

We made the subpopulation neurons absent in the decoder (already trained as described in the previous section) by setting their corresponding weights to 0. For each of the subpopulations, we also obtained a control accuracy by computing the average accuracy obtained leaving out randomly chosen sets of neurons of the same size. We used two-way repeated measures ANOVA with within-subjects factors of *Population* (1–6) and *Drop type* (specific, random) to assess the effect of *active* population exclusion on overall decoder prediction accuracy. We used two-way repeated measures ANOVA with within-subjects factors of *Drop type* (specific, random) and *Trial type* (response, no-response) to assess the effect of specific population exclusion on decoder prediction accuracy and followed up with Bonferroni post hoc tests where appropriate (alpha level was set at 0.05).

Data, Materials, and Software Availability. All data generated or analyzed during this study, and needed to evaluate the conclusions in the paper, are included in the manuscript (Figs. 1–3) and supplementary materials (*SI Appendix*). Raw imaging data files are not archived in a repository as it is too large and will be made available by the corresponding author (bhope@intra.nida.nih.gov or rajtarun.madangopal@nih.gov) upon reasonable request.

ACKNOWLEDGMENTS. We thank the Genetically-Encoded Neuronal Indicator and Effector Project and the Janelia Research Campus of the Howard Hughes Medical Institute for generously allowing the use of GCaMP6 in our research. We thank members of the Hope Lab, Shaham Lab, and the Machine Learning Core of the National Institute of Mental Health Intramural Research Program for their feedback on study design and analysis and for thoughtful comments during the writing of this manuscript. We thank Dr. David Moorman for his valuable insights into data analysis and interpretation. Research was supported by the Intramural Research Programs of the National Institute on Drug Abuse (Projects ZIA-DA000467, and ZIA-DA000434, and the National Institute of Mental Health (project ZIC-MH002968). The contributions of the NIH author(s) are considered Works of the United States Government. The findings and conclusions presented in this paper are those of the author(s) and do not necessarily reflect the views of the NIH or the U.S. Department of Health and Human Services.

Author affiliations: [a]Behavioral Neuroscience Research Branch, Intramural Research Program, National Institute on Drug Abuse, Baltimore, MD 21224; [b]Machine Learning Core, Intramural Research Program, National Institute of Mental Health, Bethesda, MD 20892; and [c]Cellular and Neurocomputational Systems Branch, Intramural Research Program, National Institute on Drug Abuse, Baltimore, MD 21224

Author contributions: R.M., D.C., Y.S., and B.T.H. designed research; R.M., C.H., B.L., G.B., L.E.K., S.J.W., D.J.T., Y.G., D.Q.P., K.E.S., B.L.W., D.C., M.V., J.M.B., L.A.R., and H.P.J. performed research; R.M., Y.Z., C.H., J.Z., B.L., G.B., K.C.L., K.E.S., J.M.B., H.P.J., G.S., D.-T.L., and F.P. contributed new reagents/analytic tools; R.M., Y.Z., C.H., J.Z., B.L., G.B., K.C.L., L.E.K., S.J.W., D.J.T., Y.G., D.Q.P., K.E.S., B.L.W., D.C., M.V., J.M.B., L.A.R., G.S., D.-T.L., Y.S., F.P., and B.T.H. analyzed data; and R.M., Y.Z., K.E.S., L.A.R., Y.S., F.P., and B.T.H. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[5]Present address: Institute Pasteur Italia—FondazioneCenci Bolognetti—Department of Physiology and Pharmacology, SapienzaUniversity of Rome, Rome 00161, Italy.

[6]Present address: Santa Lucia Foundation (La Fondazione Santa Lucia di Roma è un Istituto di Ricovero e Cura a Carattere Scientifico, Fondazione Santa Lucia), Rome 00179, Italy.

[7]Present address: Department of Neurobiology, University of Maryland School of Medicine, Baltimore, MD 21201.

[8]Y.S., F.P., and B.T.H. contributed equally to this work.

1. A. Burgos-Robles, H. Bravo-Rivera, G. J. Quirk, Prelimbic and infralimbic neurons signal distinct aspects of appetitive instrumental behavior. *PLoS ONE* **8**, e57575 (2013).
2. J. W. Dalley, R. N. Cardinal, T. W. Robbins, Prefrontal executive and cognitive functions in rodents: Neural and neurochemical substrates. *Neurosci. Biobehav. Rev.* **28**, 771–784 (2004).
3. D. R. Euston, A. J. Gruber, B. L. McNaughton, The role of medial prefrontal cortex in memory and decision making. *Neuron* **76**, 1057–1070 (2012).
4. J. G. Howland, R. Ito, C. C. Lapish, F. R. Villaruel, The rodent medial prefrontal cortex and associated circuits in orchestrating adaptive behavior under variable demands. *Neurosci. Biobehav. Rev.* **135**, 104569 (2022).
5. B. Kaminska, J. P. Caballero, D. E. Moorman, "Chapter Three: Integration of value and action in medial prefrontal neural systems" in *International Review of Neurobiology*, A. T. Brockett, L. M. Amarante, M. Laubach, M. R. Roesch, Eds. (Academic Press, 2021), vol. 158, pp. 57–82.
6. D. E. Moorman, M. H. James, E. M. McGlinchey, G. Aston-Jones, Differential roles of medial prefrontal subregions in the regulation of drug seeking. *Brain Res.* **1628**, 130–146 (2015).
7. R. T. LaLumiere, K. E. Niehoff, P. W. Kalivas, The infralimbic cortex regulates the consolidation of extinction after cocaine self-administration. *Learn. Mem.* **17**, 168–175 (2010).
8. J. Peters, R. T. LaLumiere, P. W. Kalivas, Infralimbic prefrontal cortex is responsible for inhibiting cocaine seeking in extinguished rats. *J. Neurosci.* **28**, 6046–6053 (2008).
9. J. Peters, J. Vallone, K. Laurendi, P. W. Kalivas, Opposing roles for the ventral prefrontal cortex and the basolateral amygdala on the spontaneous recovery of cocaine-seeking in rats. *Psychopharmacology (Berl)* **197**, 319–326 (2008).
10. S. L. Gourley, J. R. Taylor, Going and stopping: Dichotomies in behavioral control by the prefrontal cortex. *Nat. Neurosci.* **19**, 656–664 (2016).
11. B. L. Warren *et al.*, Separate vmPFC ensembles control cocaine self-administration versus extinction in rats. *J. Neurosci.* **39**, 7394–7407 (2019).
12. B. L. Warren *et al.*, Distinct Fos-expressing neuronal ensembles in the ventromedial prefrontal cortex mediate food reward and extinction memories. *J. Neurosci.* **36**, 6691–6703 (2016).
13. J. M. Bossert *et al.*, Ventral medial prefrontal cortex neuronal ensembles mediate context-induced relapse to heroin. *Nat. Neurosci.* **14**, 420–422 (2011).
14. L. S. Brebner *et al.*, The emergence of a stable neuronal ensemble from a wider pool of activated neurons in the dorsal medial prefrontal cortex during appetitive learning in mice. *J. Neurosci.* **40**, 395–410 (2020).
15. L. R. Whitaker *et al.*, Associative learning drives the formation of silent synapses in neuronal ensembles of the nucleus accumbens. *Biol. Psychiatry* **80**, 246–256 (2016).
16. L. R. Whitaker *et al.*, Bidirectional modulation of intrinsic excitability in rat prelimbic cortex neuronal ensembles and non-ensembles after operant learning. *J. Neurosci.* **37**, 8845–8856 (2017).
17. S. Pfarr *et al.*, Losing control: Excessive alcohol seeking after selective inactivation of cue-responsive neurons in the infralimbic cortex. *J. Neurosci.* **35**, 10750–10761 (2015).
18. N. Suto *et al.*, Distinct memory engrams in the infralimbic cortex of rats control opposing environmental actions on a learned behavior. *eLife* **5**, e21920 (2016).
19. B. L. Warren, N. Suto, B. T. Hope, Mechanistic resolution required to mediate operant learned behaviors: Insights from neuronal ensemble-specific inactivation. *Front. Neural Circuits* **11**, 28 (2017).
20. L. Kane *et al.*, Fos-expressing neuronal ensemble in rat ventromedial prefrontal cortex encodes cocaine seeking but not food seeking in rats. *Addict. Biol.* **26**, e12943 (2020).
21. D. E. Moorman, G. Aston-Jones, Prefrontal neurons encode context-based response execution and inhibition in reward seeking and extinction. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 9472–9477 (2015).
22. D. E. Moorman, G. Aston-Jones, Prelimbic and infralimbic medial prefrontal cortex neuron activity signals cocaine seeking variables across multiple timescales. *Psychopharmacology* **240**, 575–594 (2023).
23. S. Riaz *et al.*, Prelimbic and infralimbic cortical inactivations attenuate contextually driven discriminative responding for reward. *Sci. Rep.* **9**, 3982 (2019).
24. Y. M. Peters, P. O'Donnell, R. M. Carelli, Prefrontal cortical cell firing during maintenance, extinction, and reinstatement of goal-directed behavior for natural reward. *Synapse* **56**, 74–83 (2005).
25. E. A. West, M. P. Saddoris, E. C. Kerfoot, R. M. Carelli, Prelimbic and infralimbic cortical regions differentially encode cocaine-associated stimuli and cocaine-seeking before and following abstinence. *Eur. J. Neurosci.* **39**, 1891–1902 (2014).
26. R. I. Grant *et al.*, Specialized coding patterns among dorsomedial prefrontal neuronal ensembles predict conditioned reward seeking. *eLife* **10**, e65764 (2021).
27. J. M. Otis *et al.*, Prefrontal cortex output circuits guide reward seeking through divergent cue encoding. *Nature* **543**, 103–107 (2017).
28. B. Liang *et al.*, Distinct and dynamic on and off neural ensembles in the prefrontal cortex code social exploration. *Neuron* **100**, 700–714.e709 (2018).
29. Y. Zhang *et al.*, Detailed mapping of behavior reveals the formation of prelimbic neural ensembles across operant learning. *Neuron* **110**, 674–685.e676 (2022).
30. G. Barbera *et al.*, Spatially compact neural clusters in the dorsal striatum encode locomotion relevant information. *Neuron* **92**, 202–213 (2016).
31. L. Zhang *et al.*, Miniscope GRIN lens system for calcium imaging of neuronal activity from deep brain structures in behaving animals. *Curr. Protoc. Neurosci.* **86**, e56 (2019).
32. D. J. Calu, Y.-W. Chen, A. B. Kawa, S. G. Nair, Y. Shaham, The use of the reinstatement model to study relapse to palatable food seeking during dieting. *Neuropharmacology* **76**, 395–406 (2014).
33. M. Venniro, D. Caprioli, Y. Shaham, Animal models of drug relapse and craving: From drug priming-induced reinstatement to incubation of craving after voluntary abstinence. *Prog. Brain Res.* **224**, 25–52 (2016).
34. S. G. Nair *et al.*, Role of dorsal medial prefrontal cortex dopamine D1-family receptors in relapse to high-fat food seeking induced by the anxiogenic drug yohimbine. *Neuropsychopharmacology* **36**, 497–510 (2011).
35. D. R. Sparta *et al.*, Activation of prefrontal cortical parvalbumin interneurons facilitates extinction of reward-seeking behavior. *J. Neurosci.* **34**, 3699 (2014).
36. L. S. Brebner *et al.*, Extinction of cue-evoked food-seeking recruits a GABAergic interneuron ensemble in the dorsal medial prefrontal cortex of mice. *Eur. J. Neurosci.* **52**, 3723–3737 (2020).
37. E. Koya *et al.*, Targeted disruption of cocaine-activated nucleus accumbens neurons prevents context-specific sensitization. *Nat. Neurosci.* **12**, 1069–1073 (2009).
38. A. Ishikawa, F. Ambroggi, S. M. Nicola, H. L. Fields, Contributions of the amygdala and medial prefrontal cortex to incentive cue responding. *Neuroscience* **155**, 573–584 (2008).
39. R. Madangopal *et al.*, Discriminative stimuli are sufficient for incubation of cocaine craving. *eLife* **8**, e44427 (2019).
40. S. M. Nicola, I. A. Yun, K. T. Wakabayashi, H. L. Fields, Firing of nucleus accumbens neurons during the consummatory phase of a discriminative stimulus task depends on previous reward predictive cues. *J. Neurophysiol.* **91**, 1866–1882 (2004).
41. R. Madangopal *et al.*, Inactivation of the infralimbic cortex decreases discriminative stimulus-controlled relapse to cocaine seeking in rats. *Neuropsychopharmacology* **46**, 1969–1980 (2021).
42. N. L. Goodwin *et al.*, Simple behavioral analysis (SimBA) as a platform for explainable machine learning in behavioral neuroscience. *Nat. Neurosci.* **27**, 1411–1424 (2024).
43. A. Mathis *et al.*, DeepLabCut: Markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* **21**, 1281–1289 (2018).
44. E. Koya *et al.*, Silent synapses in selectively activated nucleus accumbens neurons following cocaine sensitization. *Nat. Neurosci.* **15**, 1556–1562 (2012).
45. L. R. Whitaker, B. T. Hope, Chasing the addicted engram: Identifying functional alterations in Fos-expressing neuronal ensembles that mediate drug-related learned behavior. *Learn. Mem.* **25**, 455–460 (2018).
46. L. R. Whitaker *et al.*, Bidirectional modulation of intrinsic excitability in rat prelimbic cortex neuronal ensembles and non-ensembles after operant learning. *J. Neurosci.* **37**, 8845 (2017).
47. J. J. Ziminski *et al.*, Changes in appetitive associative strength modulates nucleus accumbens, but not orbitofrontal cortex neuronal ensemble excitability. *J. Neurosci.* **37**, 3160 (2017).
48. F. J. Rubio *et al.*, Context-induced reinstatement of methamphetamine seeking is associated with unique molecular alterations in Fos-expressing dorsolateral striatum neurons. *J. Neurosci.* **35**, 5625 (2015).
49. F. J. Rubio *et al.*, Flow cytometry of synaptoneurosomes (FCS) reveals increased ribosomal S6 and calcineurin proteins in activated medial prefrontal cortex to nucleus accumbens synapses. *J. Neurosci.* **43**, 4217 (2023).
50. W. E. Allen *et al.*, Thirst-associated preoptic neurons encode an aversive motivational drive. *Science* **357**, 1149–1155 (2017).
51. L. A. DeNardo *et al.*, Temporal evolution of cortical ensembles promoting remote memory retrieval. *Nat. Neurosci.* **22**, 460–469 (2019).
52. J. H. Hyun *et al.*, Tagging active neurons by soma-targeted Cal-Light. *Nat. Commun.* **13**, 7692 (2022).
53. B. Moeyaert *et al.*, Improved methods for marking active neuron populations. *Nat. Commun.* **9**, 4440 (2018).
54. M. Hagemann-Jensen, C. Ziegenhain, R. Sandberg, Scalable single-cell RNA sequencing from full transcripts with Smart-seq3xpress. *Nat. Biotechnol.* **40**, 1452–1457 (2022).
55. S. Picelli *et al.*, Full-length RNA-seq from single cells using Smart-seq2. *Nat. Protoc.* **9**, 171–181 (2014).
56. K. E. Savell *et al.*, MultipleXed population selection and enrichment single nucleus RNA sequencing (XPoSE-seq) enables sample identity retention during transcriptional profiling of rare populations. bioRxiv [Preprint] (2023). https://doi.org/10.1101/2023.09.27.559834 (Accessed 6 March 2025).
57. R. Madangopal *et al.*, Incubation of palatable food craving is associated with brain-wide neuronal activation in mice. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2209382119 (2022).
58. M. E. Bouton, *Context, Time, and Memory Retrieval in the Interference Paradigms of Pavlovian Learning* (American Psychological Association, 1993), pp. 80–99.
59. M. E. Bouton, S. Maren, G. P. McNally, Behavioral and neurobiological mechanisms of Pavlovian and instrumental extinction learning. *Physiol. Rev.* **101**, 611–681 (2021).
60. National Research Council, *Guide for the Care and Use of Laboratory Animals* (National Academies Press, Washington, DC, 2011), 10.17226/12910.