JNeurosci
THE JOURNAL OF NEUROSCIENCE

Research Articles | Behavioral/Cognitive

Habit learning shapes activity dynamics in the central nucleus of the amygdala

https://doi.org/10.1523/JNEUROSCI.0751-26.2026

Received: 5 May 2026
Revised: 14 August 2026
Accepted: 20 August 2026

Title: Habit learning shapes activity dynamics in the central nucleus of the amygdala

Authors: Kenneth A. Amaya[1, 2], J. Eric Carmichael[1, 3], Jeffrey J. Stott[1], Erica S. Townsend[1], Elizabeth A. Bien[1], Jensen A. Palmer[1], and Kyle S. Smith[1]

Affiliations:

1. Department of Psychological and Brain Sciences, Dartmouth College, Hanover, NH 03755, USA
2. Department of Neuroscience, Tufts University School of Medicine, Boston, MA 02111, USA
3. Department of Psychiatry, Douglas Mental Health Research Institute, McGill University, Montreal, QC, Canada

Corresponding authors: Kenneth A. Amaya, Kyle S. Smith

Email: Kenneth.Amaya@tufts.edu; Kyle.S.Smith@dartmouth.edu

Keywords: instrumental, learning, electrophysiology, cFos, reinforcement

Abstract:

A habit develops as a motivated behavior that is strengthened by extended experience and feedback. While the performance component of a habit is relatively well studied, being linked to action-related neural dynamics in the basal ganglia and beyond, neural mechanisms for outcome feedback during habit formation, the reinforcement component, remain unknown. One candidate for this feedback mechanism is in the central nucleus of the amygdala (CeA). Here, we identify CeA neural firing dynamics in male and female rats that serve to promote habit formation. We used a novel maze task with rewards of differing identity and value to show that habits arise with task overtraining. After showing that overtraining engages CeA as shown through elevated cFos expression, we recorded in-vivo CeA activity across learning, and after outcome devaluation when habitual behavior is most identifiable. Neuronal activity changes tracked with habit formation. During learning, a group of recorded cells were significantly responsive at the choice point of maze trials while others encoded outcome consumption. By late training, neural activity exhibited a rapid depression at the choice point of the maze and excitation during reward receipt. In both cases, outcome magnitude, but not identity, was a significant modulator of neural activity. Additionally, a population of neurons tracked instantaneous changes in animal run speed.

These speed cells dramatically decreased in number as habits formed. Together, these findings support a role for the CeA in providing reinforcement for habitual behavior, offering a signal that marks successful performance, while identifying speed representations in the CeA.

Significance Statement

Habits enable efficient behavior but can also become inflexibly maladaptive. Although the neural circuits underlying habitual action execution have been studied extensively, the mechanisms by which outcome feedback reinforces habit formation remain poorly understood. Here, we show that neural activity in the central nucleus of the amygdala evolves with habit development, shifting from representations of choice and reward consumption to reinforcement-related signals during overtraining. Activity reflected outcome magnitude rather than identity, consistent with a role for the central amygdala in reinforcement. We also identify a previously unrecognized population of CeA neurons that track animal running speed. Together, these findings implicate the CeA as a source of reinforcement signals that promote habit formation.

Introduction

When repeatedly performing instrumental actions, animals can transition from reliance on goal-directed behaviors that rely upon associations between actions and outcomes, to habits, governed by associations between responses and antecedent stimuli (Balleine and Dickinson, 1998). Cortical, striatal, and dopaminergic contributions have been identified that underlie these two learning processes (Gasbarri et al., 2014; Amaya and Smith, 2018; Malvaez and Wassum, 2018; Balleine, 2019; Malvaez, 2020; Vandaele and Ahmed, 2021). In addition, there is evidence that the amygdala plays a key role in dictating the balance of goal-directed versus habitual strategies of behavior. Lesions to the central nucleus of the amygdala (CeA) abolish habit formation, as does functional disconnection of the CeA from the dorsolateral striatum (DLS) (Lingawi and Balleine, 2012). The CeA is connected within a wider habit circuit in the brain, including the infralimbic cortex and substantia nigra pars compacta (Steinberg et al., 2020). Moreover, CeA innervation of the dorsomedial striatum has been found to interfere with goal-directed processes to promote habits (Giovanniello et al., 2025).

Classical models of habits require the brain to link stimuli and responses, which are reinforced by positive outcomes and which, with time, can be executed automatically (Thorndike, 1927; Dickinson and Weiskrantz, 1985; Balleine, 2019). Therefore, brain dynamics for encoding of habits should reflect the development and execution of behavior as a habit. In accord with this idea, one common activity motif among habit-promoting brain areas is action “chunking”, such that activity evolves to represent action boundaries as action sequences become well-learned (Graybiel, 1998; Jog et al., 1999; Barnes et al., 2005; Jin and Costa, 2010; Smith and Graybiel, 2013). In several brain areas, action chunking aligns with the formation of habits defined by vigorous, consistent, and outcome-insensitive performance (Smith and Graybiel, 2016). There is evidence that habits arise not only from activity in habit brain regions, but also how they are balanced with activity in systems promoting goal-directed actions For example, habit learning has been linked to changes in the DMS activity and its innervation from cortical and subcortical areas (Thorn et al., 2010; Gremel and Costa, 2013; Seiler et al., 2022; van Elzelingen et al., 2022; Giovanniello et al., 2025).

To date, the majority of neural dynamics that represent habits are found to do so through engagement during the behavioral performance period. Yet, there is also evidence that neural activity during the post-performance window can be relevant, with, for example, neurons in the DLS representing rewards and errors

comprise a cell population that is wholly distinct from the population that represents performance itself (Smith and Graybiel, 2016; Cunningham et al., 2021). In this population, activity at the time of reward receipt or reward loss (from performance error) changes in alignment with the emergence of habits, specifically in the loss of error-related signals to, in principle, encourage continued habitual behavior despite negative feedback (Smith and Graybiel, 2016). There is also evidence that the CeA plays a role in mediating the reinforcing nature of outcomes for control over habits. Stimulation of the CeA at the time of reward occurrence can increase reward-seeking behaviors to the point of compulsion (Robinson et al., 2014; Tom et al., 2019; Warlow et al., 2020). Additionally, in population-level activity, the CeA can be engaged during addiction-like behaviors like ethanol-seeking, and CeA-to-dorsomedial striatum projection neurons are active at the end of reward seeking behaviors including habits following stress (Fraser et al., 2024; Giovanniello et al., 2025).

Given this background information, the CeA clearly contributes to habit development and performance, yet there are no studies that have examined how single-unit CeA neuronal activity relates to habits. To address this gap, we first detail a self-paced maze task to demonstrate that extended training produces a habit response. Using that task, we measure CeA recruitment as a product of extended training before recording electrophysiological activity in the CeA to relate neuronal activity there to aspects of habit formation.

Methods

Subjects. Individually housed wild-type Long-Evans rats (N = 17, 8 male and 9 female, behavioral; N = 9, 5 male and 4 female, electrophysiology), transgenic cFos-GFP Long-Evans rats (N = 12, 6 male and 6 female, cFos experiment) were maintained on a reverse light-dark cycle and within 85% of their *ad libitum* body weights. Wild-type Long-Evans rats were acquired from Charles River. The cFos-GFP transgenic rats were provided by the laboratory of Dr. Bruce Hope (NIDA). Animals were run in experiments during their dark (active) cycle. All procedures were in accordance with the National Institute of Health's Guide for the Care and Use of Laboratory Animals. Protocols were approved by the Dartmouth College Institutional Animal Care and Use Committee.

Plus maze habit task. Training began approximately 1 week following the commencement of food deprivation, once animals were at the target weight threshold of 85% of their *ad libitum* weights.

Rats first underwent an initial 60-minute acclimation session on a custom plus-shaped maze in a dimly lit room with ambient white noise playing (~20 dB) while wearing custom-made LED backpacks for position detection via an overhead camera. During this session, both grape-flavored and orange-flavored 20% sucrose solution were freely available at food sites located at the ends of each maze arm. The maze itself was constructed of aluminum sheets (12.1 cm in width x 61.3 cm in length) painted with black matte paint. Each aluminum sheet was connected to a square, black plastic piece (12.1 cm x 12.1 cm) that sat atop PVC pipe to elevate the maze 24.1 cm above the floor. Reward receptacles were affixed to the ends of each arm and were either connected to a pellet hopper (modified from MedAssociates GS-ENV-203M-20; cFos experiment) or liquid tubing connected to a 60 mL syringe operated by a variable speed pump (MedAssociates PHM-210; behavioral, electrophysiology experiments).

Rats commenced training where a specific response (turn right or turn left) was required to earn outcomes at each arm-end. Task rules were counterbalanced across rats and never changed. The four arm ends delivered either large magnitude (0.3 mL) grape solution, large magnitude orange solution, small magnitude (0.1 mL) grape solution, or small orange solution for correct responses. Orange and grape flavoring (No-Sugar Kool-Aid, Kraft Heinz, Chicago, IL) was added to 20% sucrose at 2 g/L. For rats in the cFos experiment, chocolate or banana flavored sucrose pellet (Bioserv #F0299 and #F0059, Flemington, NJ) outcomes were used instead of sucrose solutions, as these animals were run in an old version of the maze task that was outfitted with pellet hoppers. The location of each magnitude-flavor pairing remained constant throughout the duration of each experiment. Sessions ended after either 60 minutes or 100 trials elapsed, whichever occurred first. Successful completion of a training session was defined as a response accuracy of at least 60% correct turns during the session.

The duration of training varied by a rat's group assignment. Key timepoints were defined based on prior literature (Smith and Graybiel, 2013), where Criteria training involved successful completion of 3 sessions and Overtraining involved successful completion of 10 sessions. In the Behavioral Experiment (Figure 1), rats were split into two groups (Criteria-trained, CT, N = 7; Overtrained, OT, N = 10). In the central amygdala cFos experiment (Figure 2), rats were split into three groups (Criteria-trained, CT, N = 4; Overtrained, OT, N = 4; and No-Run Control, N = 4). The No-Run group experienced the same handling techniques but no behavioral procedures to control for the frequent handling and backpack-wearing of the experimental groups.

To minimize the amount of time rats (N = 9) were implanted, maze training began before surgical implantation. Pre-surgery training was completed once animals finished 100 trials (including correct and incorrect response types) within the allotted 60-minute session time limit. Surgical implantation of silicon probes followed pre-training. Training resumed approximately 1.5 weeks following surgery, once animals hit their threshold of 85% of their *ad libitum* weights, where each rat went through Overtraining.

Following training, rats in Experiment 1 and rats implanted with recording electrodes underwent outcome devaluation by conditioned taste aversion (methods below). After devaluation, rats were tested on a brief 5-minute extinction probe, similar to habit tests from the literature, where no outcomes were presented for trial completion. The following 3 sessions were post-devaluation reacquisition tests, where correct responses were rewarded (as in the training phase). Post-devaluation reward consumption was recorded during reacquisition sessions.

Outcome devaluation procedures. Outcome devaluation involved free access to one reward flavor was given on two of the maze arms (North/South or East/West) and paired with a 0.3M injection of lithium chloride (10 mg/kg) to devalue one of the two reward identities (grape or orange). Devaluation proceeded over multiple sessions until animals rejected the delivered outcomes, as previously described (Amaya et al., 2020).

Behavioral data analysis. Animal position was recorded for the duration of each session during the experiment. Trial accuracy (percent correct turns) was used calculated for each training session. During and after outcome devaluation, outcome consumption was measured. Using custom MATLAB scripts, animal speed was calculated for each trial. Additionally, animal tracking allowed for quantification of the curvature of a rat's trajectory during each trial, and using the IdPhi metric, trials could be classified as either having or not having a deliberative vicarious trial error (VTE where IdPhi > 180) (Stott and Redish, 2014).

Animal learning was analyzed using generalized linear mixed models to model Correct Completion Percentage as predicted by Session, or Group, Session, and their interaction, where appropriate. Odds ratios, 95% confidence intervals, and p-values are reported. For measures of completion rate and animal speed, linear mixed models were used to model those variables as predicted by Session. There, effect estimates, 95% confidence intervals, and p-values are reported. Analyses and figure generation were conducted in R (R Core Team 2016, "stats", "lme4", "lmerTest", "car", "ggplot2"). Figures were stylized in Adobe Illustrator.

cFos experiment histological procedures and analysis. Animals were sacrificed 90 minutes following the completion of the final training session for each group to allow for expression of the immediate early gene cFos. To accomplish this, animals were deeply anesthetized via isoflurane then transcardially perfused with 0.9% saline followed by 10% formalin. Brains were extracted and stored in 20% sucrose for 2 days before being sectioned at a thickness of 60 microns, targeting the central nucleus of the amygdala at the coronal plane 1.8mm behind bregma, informed by prior literature (Lingawi and Balleine, 2012). Sections were mounted onto microscope slides and cover-slipped with DAPI-containing hard set mounting medium (Vectashield; Vector Laboratories, Burlingame, CA, USA). Bilateral GFP expression and DAPI labeling in the CeA was visualized via fluorescence microscopy and images were saved for cell counting.

cFos experiment data analysis. Animals contributed two sections each before cFos quantification. Counts of cells expressing cFos were collected in ImageJ by an experimenter blinded to the group identities of the animals. DAPI-labeled cells were also counted by the same experimenter and overlap between the two counts were used to verify that putative cFos-labelling were indeed cells. Counts from two sections were averaged to obtain a single cFos count measure for each animal. Data analysis proceeded by comparing animal-level expression of CeA cFos+ densities, separated into medial and lateral CeA subdivisions, between groups using a linear mixed model. Analyses and figure generation were conducted in R (R Core Team 2016, "stats", "lme4", "lmerTest", "car", "ggplot2"). Figures were stylized in Adobe Illustrator.

Electrophysiological recording experiment surgery. Rats used in the electrophysiological recording experiment were implanted with 16-channel, 50-μm-thick silicon probes (A2x2-tet-10mm-150-150-121; NeuroNexus, Ann Arbor, MI). Probes were ordered attached to a chronic microdrive (dDrive XL, NeuroNexus) that stood 15.5 mm in height and featured a maximum driving distance of 7.0 mm. During surgery, rats were anesthetized via isoflurane (5% induction, 1 –3% maintenance) and a non-steroidal anti-inflammatory agent was administered pre- and post-operatively to reduce pain and discomfort. Silicon probes were implanted unilaterally over the central nucleus of the amygdala (relative to bregma AP: -1.8 mm; ML: +4.0 mm, per Lingawi and Balleine 2012) and lowered via the stereotaxic arm until the base of the drive mechanism made contact with the surface of the skull (typical probe depth about 2.4 mm into the brain). Once at this resting depth, the drive mechanism was affixed to the surface of the skull via Metabond dental cement (Parkell, Edgewood, NY) and implanted bone screws that held down copper mesh, to be used later in the surgery to create an on-head faraday cage to reduce noise and protect the implant. A ground screw (over cerebellum, contralateral to implant site) and a reference screw (over cerebellum, ipsilateral to implant) were then connected to the ground and reference wires coming from the probe and cemented into place. Probes were lowered approximately 3.0 mm further during the surgery (final surgical probe depth: 5.2 – 5.6 mm). The on-head faraday cage was constructed and reinforced with dental cement and mated with a custom-made cap to allow for future access to the probe. Surgical procedures were adopted from past literature (Vandecasteele et al., 2012). Following surgery, animals were given post-operative intraperitoneal injections of 0.9% saline and the antibiotic Baytril for three days. Probes were lowered to the target depth (7.8 mm deep) over the course of 10 post-surgical days and were subsequently moved in < 0.02 mm steps.

Neuronal spiking data acquisition and preprocessing. Electrical signals from silicon probes were recorded via a head mounted HS-18 preamplifier and synchronized with behavioral tracking and task events using a Digital Lynx SX acquisition system (Neuralynx). The data was sampled at 32kHz, filtered between 600-6000Hz, and spike detection thresholds were set on a tetrode-by-tetrode basis before the animal began the task. When a spike was detected, 32 samples of the continuous signal representing the spike waveform were saved for offline sorting in MClust 4.3 (A. D. Redish et al.). Spike isolation was initially automated using KlustaKwik (Kadri et al. 2015) before manual refinement and curation. Only neurons with sufficient firing over the entire recording session (>0.1Hz) were included in the analyses. We report a total of 182 neurons recorded over all training and testing, with 153/182 of those neurons recorded before outcome-devaluation procedures took place. Neurons were classified into putative fast-spiking and regular-firing types using k-means clustering based on spiking statistics and waveform features (Viskontas et al., 2007).

Spiking data analyses and statistics. All behavioral and spiking data were processed in MATLAB and exported to R for statistical analyses. To determine whether CeA neurons were responsive to task events, a peri-event firing rate average was taken and z-scored using 500 random samples of the same duration from across the recording session. The event-related responses of each neuron were calculated in 2-second windows around either the choice point (approach: -1.0 – 1.0 seconds relative to the rat exiting the center of the maze) or feeder arrival (outcome: 0 – 2 seconds following feeder arrival). Neurons were classified as event-responsive if the mean activity during the event window exceeded ±1.96 SD of the pseudo-baseline (van der Meer and Redish, 2009). For visualization, peri-event time average plots were constructed using 50ms bins and smoothed with a Gaussian (100ms). All figures were generated in MATLAB or R and exported to Adobe Illustrator for curation.

Linear decoding of animal run speed. To quantify the overall correlation between the running speed of the animal and CeA neuron activity, the firing rate was computed for each neuron using 30 ms bins and convolved with a gaussian window (250 ms, 25 ms sd). Both the firing rate and the animal's speed were interpolated to have corresponding vector lengths before computing the Pearson correlation on periods of immobility (≥ 2 cm/s)

(Kropff et al., 2015). Shuffle surrogates were computed for each neuron by applying a random circular shift to the speed vector relative to the firing rate and computing the correlation value. This was repeated 1000 times to generate a shuffle distribution. If the speed-firing rate correlation with neurons exceeded the 99th percentile of the shuffle distribution that neuron was considered to be 'speed modulated'. Linear decoding was performed in MATLAB using *polyfit* and *polyval*. Firing rate and locomotion (as speed) data were binned at 1-second intervals and used as inputs. The time series were smoothed as before and a 3rd degree polynomial was used to fit. A ten-fold cross validation was applied to the testing and training data and the resulting R^2 set was averaged to give a mean decoding error for each cell. Decoding error is expressed as the squared difference between the left out and predicted speed averaged across the ten folds. To establish chance level distribution of decoding performance, the firing rate data for each cell was randomly circularly shifted between 20%-50% of the vector length relative to the speed vector (Lopes et al., 2025). The model was then trained on this relationship and used to predict the speed given the shifted firing rate. This was repeated 500 times per cell, and the average of these shuffled surrogate scores was considered the chance-level. A single mean shuffled R^2 was thus generated per cell to create a surrogate pool.

Results

Extended training on a self-paced maze task facilitates development of habit phenotypes

We first developed a self-paced plus-maze task where groups of rats (total N = 17, 8 male and 9 female) were required to turn a certain direction (right or left) as they exited one of the maze arms to earn food outcomes of varying flavor and magnitude that were available at each arm end (Fig 1A). In each daily session, rats were allowed to complete up to 100 trials within 60 minutes, where a trial was defined as the exiting of one maze arm and entrance into another maze arm. Only 'correct' turn actions, as defined by the experimenter and counterbalanced across animals, were rewarded, though 'incorrect' actions counted towards the 100 trial limit in each session. The turn rule was never signaled to the animal and the rule never changed over the course of training and testing. There were no experimenter-set intertrial intervals and animals were allowed to begin the next trial at their own discretion. Experimental groups consisted of animals that were trained to criterion (CT; 60% correct responding for 3 days) and animals that were overtrained (OT; 3 criteria days plus an additional 7 days). Rats in both groups readily and similarly learned this task, as shown by a significant effect of Session (odds ratio (OR): 1.58; 95% confidence interval (CI): 1.45 – 1.72; $p < 0.001$), but no effects of Group (OR: 1.03; CI: 0.76 – 1.39; $p = 0.85$) nor an interaction (OR: 0.92; CI: 0.83 – 1.03; $p = 0.16$) (Fig 1B).

Habits are often referred to as automatic and inflexible behaviors or responses. Much of the literature has focused on the inflexible aspect of habit by employing outcome devaluation and post-devaluation behavioral testing, where an animal relying on habit responding to complete a task would maintain performance despite no longer valuing the outcome. Here, we conduct these tests, but we also present data to analyze the 'automaticity' of habit behaviors that can be gleaned from maze running, like the run speed and the decisiveness with which rats complete trials. These metrics have an added benefit of assessing habit-like behavior without disrupting task contingencies or outcome values. We compared mean trial run speeds at two timepoints (final training session, C3 or O7, vs. first completed training session, C1) for the two groups, and found that the speed difference in the OT rats (mean: 6.02 cm/s faster on O7 compared to C1) was significantly greater than the speed difference in Criteria-trained rats (mean: 1.57 cm/s slower on C3 compared to C1) ($t(15) = -3.459$, $p = 0.004$) (Fig 1C). Additionally, we measured the decisiveness with which animals executed trials to quantify the number of trials animals performed that included a vicarious trial error (VTE; threshold = 180 IdPhi). We compared VTE

incidences at the same two timepoints for the two groups (CT and OT) and found that OT rat performance was significantly less deliberative than CT rats ($t(15) = 2.397$, $p = 0.0301$) (Fig 1D). These data suggest that rats with extended training ran faster and more directly, supporting the idea that those rats were employing a habit-like strategy to solve the maze task by Session O7.

Finally, after animals completed either criteria training (CT) or overtraining (OT), one of the two outcome flavors was devalued by conditioned taste aversion. Rats decreased consumption of the devalued outcome over devaluation sessions (Fig 1E). Following outcome devaluation, rats underwent post-devaluation probe tests in extinction (no outcomes) and reacquisition (outcomes delivered for correct turns) to determine whether extended training was sufficient to render rats outcome devaluation insensitive. In the brief 5-minute extinction test, we did not observe a difference in trial accuracy between the two groups ($W = 35$, $p = 0.44$) (Fig 1F), which we believe is due to the limited test time. However, behavior during the subsequent reacquisition phase provided a clear distinction between CT and OT groups, reflective of habit development in the OT animals. As noted, animals were required to approach each goal (including the newly devalued goals) to initiate the next trial. Given that reacquisition session parameters were identical to training sessions (60 minute and 100 trial limits), we assessed the final training session for respective groups (C3 or O7) along with the first session of reacquisition. Here, we observed robust differences in response accuracy, with a significant main effect of Group (OR: 0.49; CI: 0.25 – 0.99; $p = 0.048$), Session (OR: 0.30; CI: 0.23 – 0.39; $p < 0.001$), and Group by Session interaction (OR: 1.81; CI: 1.28 – 2.55; $p = 0.001$) (Fig 1G). Additionally, we measured outcome consumption by devaluation condition during the first reacquisition test, which did not differ between the two groups; a two-way ANOVA revealed a significant main effect of Devaluation Condition ($F(1, 30) = 16.90$, $p < 0.001$), but no main effect of Group ($F(1, 30) = 0.005$, $p = 0.94$) nor was there an interaction between Group and Devaluation Condition ($F(1, 30) = 0.137$, $p = 0.71$). In other words, both groups similarly rejected the devalued outcome while continuing to consume the non-devalued outcome (Fig 1H), supporting the idea that outcomes were sufficiently devalued and that information was recalled by rats. Finally, we analyzed correct responding over all three reacquisition sessions where a generalized linear mixed model revealed main effects of Group (OR: 1.99; CI: 1.06 – 3.76; $p = 0.034$) and Session (OR: 1.26; CI: 1.12 – 1.42; $p < 0.001$), but no interaction between these variables (OR: 0.88; CI: 0.75 – 1.03; $p = 0.10$) (Fig 1I). Therefore, Overtrained animals responded correctly more than Criteria trained

animals overall. Collectively, our findings from this experiment suggest that the different training protocols produced differing levels of devaluation sensitivity, with Overtraining biasing animals towards habit responding based on (1) run speed increase, (2) reduction of VTE events, and (3) performance continuation during post-devaluation reacquisition sessions.

Overtraining recruits CeA neurons

We next used Fos-GFP+ Long Evans rats (total N = 12, 6 male and 6 female) to observe CeA engagement as a product of limited and extended training on the maze task (Fig 2A). Here, experimental groups included animals that were trained to criterion (CT), were overtrained (OT), or a No-Run control group that was handled in the same manner as OT rats. Again, animals in both groups similarly learned the task, as shown by a significant effect of Session (OR: 1.37; CI: 1.21 – 1.54; $p < 0.001$), but no effects of Group (OR: 1.13; CI: 0.59 – 2.17; $p = 0.71$) or an interaction between the two variables (OR: 0.90; CI: 0.79 – 1.01; $p = 0.07$) (Fig 2B). Shortly after completion of their final training session (C3 for Criteria rats; O7 for Overtrained rats), animals were sacrificed, and tissue was processed to visualize cFos expression in the CeA (Fig 2C). As shown by a linear mixed model, no difference in cFos density was observed between the No-Run control and Criteria groups (est: -6.60; CI: -46.88 – 33.49; $p = 0.73$) while Overtrained group showed significantly elevated cFos expression compared to controls (est: 110.42; CI: 64.02 – 156.82; $p < 0.001$). Interestingly, there was a significant interaction between subnucleus and group (est: 102.58: 36.96 – 168.20; $p = 0.004$), though there was no main effect of subnucleus on cFos density (est: 40.28; CI: -6.12 – 86.68; $p = 0.084$) (Fig 2D). Together these data show that extended training engaged the lateral subdivision of the CeA.

Action performance and outcome receipt differentially modulate CeA neuron activity over habit learning

The above experiments established that overtraining on our maze task creates habitual behaviors, as defined by performance automation measures and devaluation-insensitivity (in the reacquisition phase), and that this overtraining period is accompanied by a sharp increase in CeA engagement. To better understand what aspects of habitual behavior CeA neurons encode, we next implanted silicon recording probes targeting the CeA

(N = 5 male, 4 female) and recorded 153 units. These units had an average firing rate of 7.088 Hz throughout training and overtraining on the maze task (Fig 3A, 3B, Methods). In the task, outcomes again varied in identity and size across the maze's four end-arms and included large magnitude (0.3 mL) grape solution, large magnitude orange solution, small magnitude (0.1 mL) grape solution, or small orange solution. Animals readily learned the task, reflected by their performance above the experimenter-set 60% correct criteria for the duration of the experiment with a significant effect of session on the proportion of correct/reinforced trials (OR: 1.12; CI: 1.10 – 1.14; $p < 0.001$) (Fig 3C). Further, we observed a significant effect of session on correct trial rate (est: 0.38; CI: 0.31 – 0.46; $p < 0.001$) (Fig 3D). Thus, animals increased the rate at which they made correct responses across sessions. As we did in the behavioral experiment above, we used tracking data to approximate the decisiveness with which animals completed each trial. Trials were classified as either containing or not containing a VTE (Fig 3E). Overall, the decisiveness of trial runs increased over training, as shown by a significant decrease in mean IdPhi over session (est: -6.07; CI: -10.95 – (-1.19); $p = 0.015$) (Fig 3F). Similarly, the proportion of laps that contained a VTE decreased over training, as shown by a significant effect of session on the proportion of VTE laps observed (OR: 0.92; CI: 0.90 – 0.94; $p < 0.001$) (Fig 3G). Rats were also more engaged in task performance as they spent less time around the feeder before starting the next trial, as shown by a decrease in feeder dwell time over sessions (est: -0.22; CI: -0.33 – (-0.11); $p < 0.001$) (Fig 3H). Taken together, the behavioral measures from this experiment demonstrate that rats increased the rate at which they completed the task, were more decisive, and became increasingly engaged as training progressed. These data were all in alignment with the findings from the experiment above, suggesting that animals were responding habitually by the end of overtraining. Behavior after outcome devaluation, reported below, further confirmed this conclusion.

Recorded neural activity in the CeA was dynamic during training, with populations of neurons showing modulation by trial performance and outcome consumption. Two example units show increased activity during distinct periods of trial completion on the maze task, such as during action performance (Fig 4A) or during outcome receipt/consumption (Fig 4B). To visualize unit responses over time, we created peri-event time averages of z-scored firing rates for each of the recorded units, where time = 0 is either when the approached the choice point (Fig 4C) or arrived at the reward feeder (Fig 4E). Cells were classified as approach-modulated or outcome-modulated if their z-scored activity in their respective 2-second windows (approach: peri-choice

point; outcome: peri-feeder arrival) exceeded 1.96x their baseline, which was calculated by averaging 500 randomly-sampled 2-second windows from that cell's recording. For approach-modulated (Fig 4D) and outcome-modulated cells (Fig 4F), we reconstructed the estimated medial-lateral and dorsal-ventral locations of the recording probes to show how these cell populations were distributed.

To understand whether modulated populations of cells changed over learning, we quantified proportions of unit responsiveness early (C2 – O2) and late (O4 – O7) in training. Early in training, the proportion of approach-responsive cells totaled 28.75% of recorded units, split as 20% activated and 8.75% inhibited around the choice point. Late in training, there were 16.28% of recorded units modulated at the choice point, all of which displayed inhibition (Fig 4G). It is possible that features of the coming outcome influenced neural activity at the choice point, leading us to analyze z-scored neural activity among all neurons during early and late learning (Fig 4J). There, we found that z-scored activity was not significantly predicted by outcome flavor (est: -0.19; CI: -0.72 – 0.35; $p = 0.51$), but was significantly predicted by outcome magnitude (est: -0.72; CI: -1.25 – (-0.19; $p = 0.008$) and learning phase (est: -1.34; CI: -2.22 – (-0.45); $p = 0.003$), and there were no significant interactions between any of these variables. This data demonstrates that choice point related neural activity decreases over learning, with activity while approaching to the large outcome magnitude being lesser than activity while approaching the small magnitude outcome.

We examined neural activity during the outcome consumption window of each trial over the course of habit learning. There was limited overlap between approach- and outcome-modulated units, as only 12 recorded units of the 153 had significantly greater activity during both windows compared to baseline (Fig 4H). Using the same learning phases, we observed that the proportion of recorded cells classified as outcome-responsive increased from early learning (20.25%; 18.75% excited) to late learning (32.56%; 30.23% excited) (Fig 4I), aligning with the increase in cFos expression reported earlier. We used the same analysis as above to determine if outcome features may predict z-scored outcome-related neural activity among recorded neurons. Here, we did not observe an effect of flavor (est: 0.44; CI: -0.15 – 1.03; $p = 0.145$) nor learning phase (est: 0.60; CI: -0.29 – 1.49; $p = 0.184$) on neural activity. However, we did observe a significant effect of outcome magnitude (est: 1.18; CI: 0.59 – 1.76; $p < 0.001$) and there were no significant interactions between any of these variables (Fig 4K).

These findings were strikingly similar to those from the approach z-scored neural activity, suggesting that outcome magnitude is a strong modulator of CeA single-unit responding across habit formation.

A subpopulation of CeA neurons dynamically tracks speed over habit learning

Another prominent behavioral feature that appeared to modulate the activity of a subset of CeA neurons was run speed. This was a surprise as, to our knowledge, speed cells have not yet been reported in the amygdala. Behaviorally, rats increased their run speed over each training session (est: 0.99; CI: 0.71 – 1.26; $p < 0.001$) (Fig 5A). However, speed was not modulated by outcome identity (est: -0.04; CI: -0.93 – 0.86; $p = 0.938$), nor magnitude (est: 0.57; CI: -0.32 – 1.46; $p = 0.208$), with a possible ceiling effect (Fig 5B, 5C). Still, we recorded a subpopulation of units that tracked instantaneous changes in animal speed. We extended time windows to plot speed (gray trace) and neuronal firing rate (trace in color) of three example units: one that is speed correlated (Fig 5D top, $r = 0.54$), one that is speed anti-correlated (Fig 5D middle, $r = -0.32$), and one that is not correlated with speed (Fig 5D bottom, $r = 0.00$). Speed correlations were calculated for each recorded unit, revealing that 46 out of 182 (25.3%) units were significantly correlated with animal speed (Fig 5E) with firing rates that were widely distributed (Fig 5F). Of our identified speed units (total n = 46), most were modulated by approach (n = 24), though some also displayed elevated activity during the outcome analysis window (n = 14) (Fig 5G top). Given that CeA units, generally, were sensitive to outcome magnitude (Fig 4I), we were interested in whether this value sensitivity extended to speed modulated units. We found that 35 of 46 (76.1%) speed modulated cells were responsive during runs to the large magnitude outcome (Fig 5G, bottom), indicating a sensitivity of speed cells generally to the value of the outcome goal. We next used spike timing (burst index and firing rate) and waveform properties (spike width) to classify all recorded units into two putative types: regular-spiking (n = 175) and fast-spiking (n = 7) (Fig 5H top). Of the 175 regular spiking units, 22.8% were speed modulated (Fig 5H bottom left, green) while 85.7% of the fast-spiking units were speed modulated (Fig 5H bottom right, pink).

Next, we quantified the proportion of speed cells over the course of habit formation, divided into the defined learning phases: early learning and, late learning, and post-devaluation. The proportion of observed

speed units decreased overall as sessions progressed (OR: 0.20; CI: 0.08 – 0.47; $p < 0.001$). When considering whether the units were positively or negatively correlated with animal run speed, the decrease in observed speed cells was not distributed equally across the two types of correlation. Specifically, early in learning, there were many more positively correlated speed units than there were negatively correlated units. The number of negatively correlated units steadily decreased over learning. However, by late learning, there was a precipitous drop in the number of positively correlated speed cells (Fig 5I). After outcome devaluation, there were almost no speed correlated units recorded, where only one negatively correlated unit was identified.

Finally, we further validated speed coding by CeA neurons by using a linear decoder. For an example cell, we show the rat's actual speed (Fig 5J, black line), the rat's decoded speed (red line), and the rat's decoded speed across shuffled surrogates (gray line). Speed cells (mean R^2= 0.046 +/- 0.010) outperformed non-speed cells (mean R^2 = 0.014 +/- 0.002) and shuffled surrogates (mean R^2 = 0.004 +/- 0.00015) in decoding accuracy based on a one-way analysis of variance (ANOVA) revealing a significant effect of condition on decoding accuracy ($F(2,361) = 40.41$; $p < 0.001$) (Fig 5K). Tukey-adjusted post hoc comparisons showed that decoding accuracy was significantly higher among Speed cells than Non-speed cells (mean difference = 0.033, $t(361) = 6.624$, $p < 0.001$) and shuffles (mean difference = 0.043, $t(361) = 8.98$, $p < 0.001$). Also, decoding accuracy was higher in Non-speed cells than shuffles (mean difference = 0.0102, $t(361) = 3.102$, $p = 0.0059$). This is likely due to some non-speed cells still encoding task events like outcome receipt, or non-speed related approach encoding. Regular spiking neurons (mean R^2 = 0.044 +/- 0.012) and fast spiking neurons (mean R^2 = 0.057 +/- 0.022) that were identified as speed cells displayed similar decoding accuracy ($t(44) = -0.492$, $p = 0.625$; Fig 5L) as did speed-correlated (mean R^2 0.046 +/- 0.012) and speed anti-correlated cells (mean R^2 = 0.026 +/- 0.007; $t(57) = 1.124$, $p = 0.266$; Fig 5M). Collectively, these findings indicate that a sizable population of speed-correlated cells exists within the CeA, and their activity is dynamic and inversely associated with habit learning.

Task performance persists following outcome devaluation despite altered outcome consumption and CeA activity.

After completion of the final overtraining session (O7), animals underwent outcome devaluation, where one reward flavor (e.g., grape) was paired with lithium chloride-induced nausea until that outcome was rejected by the animals (Fig 6A). Following outcome devaluation, animals underwent post-devaluation probe tests in extinction (no outcomes) and reacquisition (outcomes delivered). When outcomes were delivered during reacquisition sessions, rats rejected the devalued outcome as compared to the non-devalued outcome, confirming the efficacy of the outcome devaluation procedure (Fig 6B). In post-devaluation sessions, however, rats continued to perform the task at accuracy rates like those observed before devaluation, as there was no significant effect of session on correct responding (OR: 1.04; CI: 0.92 – 1.17; p = 0.536) (Fig 6C). We observed that animal running speeds were sensitive to outcome devaluation in an outcome-magnitude-specific manner, such that animal running speeds were faster when running towards the large, non-devalued magnitude outcome (est: 2.90; CI: 0.51 – 5.29; p = 0.018), but no main effect of outcome devaluation was observed on run speeds (est: 0.42; CI: -1.97 – 2.80; p = 0.725), nor was there a significant interaction between outcome devaluation and magnitude (est: -2.35; CI: -5.72 – 1.03; p = 0.168) (Fig 6D). In all, rats continued to perform the task but modulated their run speeds to the large magnitude outcomes in a devaluation-specific manner. While the design of the task involved approach to now-devalued arms of the maze, a rat operating to maximize the number of non-devalued outcomes could theoretically make intentional errors to subvert the rules and increase the output of the large, non-devalued outcome; we did not observe this among our rats. Instead, these animals maintained task performance despite earning and then mostly rejecting outcomes of reduced value on those trials.

For neural analyses in this section, we report activity from the three post-devaluation reacquisition sessions. We were unable to isolate single-unit activity during the brief 5-minute extinction probe. Generally, CeA unit responsiveness to task events deteriorated following devaluation. We created peri-event time averages of z-scored firing rates for each of the recorded units of this phase to visualize their responding relative to time points of interest like the choice point of a trial (Fig 6E top) or feeder arrival (Fig 6E bottom). We again classified cells as approach- or outcome-responsive if their activity exceeded ±1.96 SD above their z-scored baseline. In this post-devaluation task period, only 3.45% of neurons were approach-responsive while only 20.69% (13.79% excited) of neurons were modulated by outcome (Fig 6F). With data from all neurons, we created a linear mixed model to predict z-scored neural activity. First, when assessing choice point related activity (Fig 6G), we did not

observe any significant effects for outcome magnitude (est: 0.18; CI: -2.04 – 2.51; p = 0.87), outcome devaluation status (whether the animal approached the devalued arm; est: -0.56; CI: -2.78 – 1.66; p = 0.62), reacquisition session (est: -0.41; CI: -1.14 – 0.33; p = 0.27), nor any significant interactions between these variables. Similarly, when assessing outcome-related activity, we did not observe significant effects for outcome magnitude (est: 0.45; CI: -2.54 – 3.45, p = 0.77), outcome devaluation status (est: -0.43; CI: -3.42 – 2.56; p = 0.78), reacquisition session (est: -0.58; CI: -1.69 – 0.52; p = 0.30), nor any significant interactions thereof (Fig 6H). Together these data indicate that the robust approach- and outcome/magnitude-related encoding observed of CeA neurons during habit formation was effectively lost following outcome devaluation and subsequent task performance.

Discussion

Across extended training on a self-paced maze task, rats transitioned from goal-directed to habitual responding, characterized by increased run speed, reduced deliberation, reduced dwell times at reward locations, and diminished sensitivity to outcome devaluation. These behavioral changes were accompanied by a marked increase CeA engagement and by systematic changes in single-unit activity. Habit-related changes emerged in CeA activity related to reward approach, reward consumption, and movement speed. Together, these findings suggest that CeA activity is dynamically shaped as habits come to dominate behavior.

The dorsolateral striatum (DLS) and infralimbic cortex encode well-learned actions that are executed habitually (Coutureau and Killcross, 2003; Smith and Graybiel, 2013, 2016; Balleine, 2019). Given that the CeA interacts with these areas to encourage habit formation (Lingawi and Balleine, 2012), and that the CeA sends projections to the substantia nigra (Gonzales and Chesselet, 1990; Fudge and Haber, 2000; Lingawi et al., 2016; Steinberg et al., 2020), we hypothesized that CeA neurons might exhibit task-related activity patterns like those reported in habit circuits. Consistent with this, a substantial proportion of recorded units were sensitive at the choice point or to outcome receipt, with stronger responses observed for large-magnitude outcomes. This pattern suggests that CeA subpopulations may provide signals that are distinctly involved in reinforcing the performance of behavior.

At the choice point, CeA activity decreased with extended training, both in the proportion of modulated units and in overall firing changes relative to baseline. In contrast to the reduction in approach-related activity, outcome-related responses became more represented with habit learning. By foundational and modern accounts, appetitive habits are acquired incrementally through positive feedback for a behavior. These are iterative processes, strengthening behaviors that lead to positive outcomes and weakening those that do not. Within these models, however, there is no clear expectation that outcomes ought to gain a surge of influence over behaviors as they transition from goal-directed to habitual. One way to reconcile this issue is to consider that even the most ingrained habits require feedback that they ‘worked’, and that there may be specialized brain processes to accomplish this information transfer. A proportional increase of CeA neurons responsive to outcome may serve this reinforcing function for habit development.

Although we did not observe a robust devaluation effect during the extinction probe in criteria-trained animals, reacquisition data indicate that goal-directed control was preserved. This discrepancy may reflect limitations in adapting extinction-based post-devaluation assays to self-paced maze tasks rather than a failure of devaluation or absence of habit reliance per se. These results point to a useful framework for identifying habits using post-devaluation reacquisition, as habits may not be slower to extinguish than goal-directed actions as one may intuit (Bouton, 2024). Past work has successfully distinguished trained from overtrained behavior using post-devaluation reacquisition tests in mazes (Smith and Graybiel, 2013), underscoring the value of behavioral markers of habit in contexts where outcomes remain available. It is curious that CeA outcome encoding essentially disintegrated during the post-devaluation reacquisition phase, when behavior was still seemingly habitual. One possibility is that, at that point in learning, the CeA had finished training output structures to maintain habitual responding, thereby becoming less necessary as an input for them to operate in a manner that promotes acquired habits.

Recipient areas of CeA projections include the dorsomedial striatum (DMS) and substantia nigra (SN) (Steinberg et al., 2020; Giovanniello et al., 2025). As the CeA is composed largely of GABAergic neurons (Ehrlich et al., 2009), their inhibitory projections could clamp down reward-related signaling in areas that might otherwise promote goal-directed performance, thereby tipping the balance to favor habitual responding (Giovanniello et al., 2025). Specifically, this inhibitory influence at the time of reward could prevent outcome expectations thought to be encoded in the DMS from being updated if those outcomes change. Additionally, the CeA interfaces with the SN, where dopamine neurons are essential for habitual behavior (Faure et al., 2005; Jin and Costa, 2010; Shindou et al., 2019). Our present demonstration that extended task training produced elevated cFos expression in the lateral subdivision of the CeA aligns well with histology that shows what may be the lateral CeA as a key input to the SN (Steinberg et al., 2020). While our recording electrodes targeted the central portion of the CeA, with a potential preferential sampling of lateral aspects of the CeA given the orientation of the implanted probe (Figure S1), we cannot be certain which CeA subdivision isolated units were from, which is a limitation of this study and an avenue for future work. Still, the timing of recorded CeA dynamics correspond well with the timing of lateral CeA cFos changes.

Our reported suppression of CeA activity as animals approach the choice point during habit performance could be contextualized as disinhibition of nigral cells which then encode the termination of the action sequence

(Jin and Costa, 2010) and promote activity in downstream dorsolateral striatum cells (Jog et al., 1999; Barnes et al., 2005; Smith and Graybiel, 2013). One caveat to note is that CeA neurons preferentially contact GABAergic neurons in the SN, allowing for CeA projections to activate dopaminergic neurons via disinhibition (Steinberg et al., 2020). In this case, perhaps the outcome-modulated neurons (which become proportionally larger as a population late in learning) synapse onto local GABAergic neurons that then release dopaminergic signaling. CeA influence over habit networks at the time of reward receipt also touches on the problem of credit assignment, otherwise framed as the question of how a behavior is credited as successful when its outcome occurs only after the behavior has been completed. The dynamics of CeA activity here could provide this function by, potentially, offering a stamp of success to the preceding habitual actions that have occurred in CeA-recipient brain areas. Mechanistically, recipient nigral dopamine cells are again of potential relevance. These have been suggested to offer a credit assignment signal to the dorsal striatum where, it is argued, habitual action plans are ultimately sculpted (Yagishita et al., 2014; Shindou et al., 2019; Tang et al., 2024; Magnard et al., 2025).

An alternative, but not mutually exclusive, interpretation of our data is that CeA activity reflects the evolving motivational significance of behavioral outcomes. Given that the CeA is involved in outcome-related learning and cue-potentiated feeding behavior (Lee et al., 2010; Holland and Hsu, 2014), it is plausible that activity in the region serves motivational processes in some capacity. Indeed, manipulating CeA activity can strongly amplify reward seeking in a manner that is resistant to punishment, fortifying the notion that the area plays a role in encoding outcome salience (Robinson et al., 2014; Esber et al., 2015; Warlow et al., 2017, 2020). Recordings in other brain regions like the DLS have delineated outcome representation changes over habit learning (Smith and Graybiel, 2013, 2016), and habits themselves are notably sensitive to changes in motivational need-states like hunger (Dickinson et al., 1995). The increased outcome-related activity observed here may therefore reflect a process by which rewards acquire enhanced motivational value with extended experience, promoting persistent and habitual responding for them.

We also identified a population of CeA neurons whose activity tracked instantaneous changes in movement speed during maze runs. While speed-related signals have been described in striatal and dopaminergic systems (Smith and Graybiel, 2013; Rueda-Orozco and Robbe, 2015; Dudman and Krakauer, 2016), they have not, to our knowledge, been reported yet in the amygdala. Interestingly, the prevalence of these speed-correlated neurons declined with extended training, despite increases in overall run speed. It is possible

that the late-training shift in reward approach activity described above is related to the diminished presence of speed cells, as many of the significantly-modulated approach cells were also speed cells (24/33 units). One possible explanation is that early in learning, CeA activity constrains performance by modulating movement-related signals, be it as animals approach the choice point or modulate their speed, whereas with extended training, this inhibitory influence diminishes, allowing other habit-related areas to dominate control over action execution.

Our results are consistent with a broader role for the CeA in shaping motivational aspects of behavior during habit formation, especially given other work implicating the brain region in governing general motivational learning like in the Pavlovian-to-Instrumental Transfer test (Corbit and Balleine, 2005). The presently reported changes in outcome-related and movement-related activity of CeA neurons may reflect complementary processes through which the CeA influences both the reinforcing properties of outcomes as well as the vigor of responding. In this sense, CeA activity may provide a bridge between associative learning mechanisms and the expression of motivated behavior resulting from those associations, which contributes to the transition from goal-directed behavior to habit reliance.

In summary, our findings highlight a specific decrease in CeA activity at outcome approach, an increase at outcome occurrence, and a decrease in speed-related neural populations as potential contributors to habitual behavior. Future studies using cell-type and projection-specific dissections of CeA subpopulations will be necessary to determine the precise circuit mechanisms underlying these effects and to clarify how CeA activity interacts with established midbrain-corticostriatal habit systems.

Acknowledgements

KAA and KSS designed the study, JJS designed and built the maze, KAA, EST, EAB, and JAP collected the data, JEC and JJS provided technical assistance, KAA, JEC, and JJS analyzed the data, and KAA and KSS wrote the paper.

Funding

This work was supported by NSF IOS1557987 (KSS) and NIH F99NS115270 (KAA).

Disclosures

The authors have no competing interests to declare.

References

Amaya KA, Smith KS (2018) Neurobiology of habit formation. Current Opinion in Behavioral Sciences 20:145–152 Available at: http://dx.doi.org/10.1016/j.cobeha.2018.01.003.

Amaya KA, Stott JJ, Smith KS (2020) Sign-tracking behavior is sensitive to outcome devaluation in a devaluation context-dependent manner: implications for analyzing habitual behavior. Learn Mem 27:136–149.

Balleine BW (2019) The Meaning of Behavior: Discriminating Reflex and Volition in the Brain. Neuron 104:47–62.

Balleine BW, Dickinson A (1998) Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. Neuropharmacology 37:407–419.

Barnes TD, Kubota Y, Hu D, Jin DZ, Graybiel AM (2005) Activity of striatal neurons reflects dynamic encoding and recoding of procedural memories. Nature 437:1158–1161.

Bouton ME (2024) Habit and persistence. J Exp Anal Behav 121:88–96.

Corbit LH, Balleine BW (2005) Double dissociation of basolateral and central amygdala lesions on the general and outcome-specific forms of pavlovian-instrumental transfer. J Neurosci 25:962–970.

Coutureau E, Killcross S (2003) Inactivation of the infralimbic prefrontal cortex reinstates goal-directed responding in overtrained rats. Behav Brain Res 146:167–174.

Cunningham PJ, Regier PS, Redish AD (2021) Dorsolateral striatal task-initiation bursts represent past experiences more than future action plans. J Neurosci 41:8051–8064.

Dickinson A, Balleine B, Watt A, Gonzalez F, Boakes RA (1995) Motivational control after extended instrumental training. Anim Learn Behav 23:197–206.

Dickinson A, Weiskrantz L (1985) Actions and habits: the development of behavioural autonomy. Philos Trans R Soc Lond B Biol Sci 308:67–78.

Dudman JT, Krakauer JW (2016) The basal ganglia: from motor commands to the control of vigor. Curr Opin Neurobiol 37:158–166.

Ehrlich I, Humeau Y, Grenier F, Ciocchi S, Herry C, Lüthi A (2009) Amygdala inhibitory circuits and the control of fear memory. Neuron 62:757–771.

Esber GR, Torres-Tristani K, Holland PC (2015) Amygdalo-striatal interaction in the enhancement of stimulus salience in associative learning. Behav Neurosci 129:87–95.

Faure A, Haberland U, Condé F, El Massioui N (2005) Lesion to the nigrostriatal dopamine system disrupts stimulus-response habit formation. J Neurosci 25:2771–2780.

Fraser KM, Kim TH, Castro M, Drieu C, Padovan-Hernandez Y, Chen B, Pat F, Ottenheimer DJ, Janak PH (2024) Encoding and context-dependent control of reward consumption within the central nucleus of the amygdala. iScience 27:109652.

Fudge JL, Haber SN (2000) The central nucleus of the amygdala projection to dopamine subpopulations in primates. Neuroscience 97:479–494.

Gasbarri A, Pompili A, Packard MG, Tomaz C (2014) Habit learning and memory in mammals: behavioral and neural characteristics. Neurobiol Learn Mem 114:198–208.

Giovanniello JR, Paredes N, Wiener A, Ramírez-Armenta K, Oragwam C, Uwadia HO, Yu AL, Lim K, Pimenta JS, Vilchez GE, Nnamdi G, Wang A, Sehgal M, Reis FM, Sias AC, Silva AJ, Adhikari A, Malvaez M, Wassum KM (2025) A dual-pathway architecture for stress to disrupt agency and promote habit. Nature Available at: http://dx.doi.org/10.1038/s41586-024-08580-w.

Gonzales C, Chesselet MF (1990) Amygdalonigral pathway: an anterograde study in the rat with Phaseolus vulgaris leucoagglutinin (PHA-L). J Comp Neurol 297:182–200.

Graybiel AM (1998) The basal ganglia and chunking of action repertoires. Neurobiol Learn Mem 70:119–136.

Gremel CM, Costa RM (2013) Orbitofrontal and striatal circuits dynamically encode the shift between goal-directed and habitual actions. Nat Commun 4:2264.

Holland PC, Hsu M (2014) Role of amygdala central nucleus in the potentiation of consuming and instrumental lever-pressing for sucrose by cues for the presentation or interruption of sucrose delivery in rats. Behav Neurosci 128:71–82.

Jin X, Costa RM (2010) Start/stop signals emerge in nigrostriatal circuits during sequence learning. Nature 466:457–462.

Jog MS, Kubota Y, Connolly CI, Hillegaart V, Graybiel AM (1999) Building neural representations of habits. Science 286:1745–1749.

Kropff E, Carmichael JE, Moser M-B, Moser EI (2015) Speed cells in the medial entorhinal cortex. Nature 523:419–424.

Lee HJ, Gallagher M, Holland PC (2010) The central amygdala projection to the substantia nigra reflects prediction error information in appetitive conditioning. Learn Mem 17:531–538.

Lingawi NW, Balleine BW (2012) Amygdala central nucleus interacts with dorsolateral striatum to regulate the acquisition of habits. J Neurosci 32:1073–1081.

Lingawi NW, Dezfouli A, Balleine BW (2016) The psychological and physiological mechanisms of habit formation. In: The Wiley Handbook on the Cognitive Neuroscience of Learning, pp 409–441. Chichester, UK: John Wiley & Sons, Ltd.

Lopes PH, Tavares LCS, Tort ABL, Blanco W (2025) Speed encoding in the rat striatum. PLoS One 20:e0334601.

Magnard R, Cheng Y, Zhou J, Province H, Thiriet N, Janak PH, Vandaele Y (2025) Sequence termination cues drive habits via dopamine-mediated credit assignment. bioRxivorg Available at: http://dx.doi.org/10.1101/2024.10.16.618735.

Malvaez M (2020) Neural substrates of habit. J Neurosci Res 98:986–997.

Malvaez M, Wassum KM (2018) Regulation of habit formation in the dorsal striatum. Curr Opin Behav Sci 20:67–74.

Robinson MJF, Warlow SM, Berridge KC (2014) Optogenetic excitation of central amygdala amplifies and narrows incentive motivation to pursue one reward above another. J Neurosci 34:16567–16580.

Rueda-Orozco PE, Robbe D (2015) The striatum multiplexes contextual and kinematic information to constrain motor habits execution. Nat Neurosci 18:453–460.

Seiler JL, Cosme CV, Sherathiya VN, Schaid MD, Bianco JM, Bridgemohan AS, Lerner TN (2022) Dopamine signaling in the dorsomedial striatum promotes compulsive behavior. Curr Biol 32:1175-1188.e5.

Shindou T, Shindou M, Watanabe S, Wickens J (2019) A silent eligibility trace enables dopamine-dependent synaptic plasticity for reinforcement learning in the mouse striatum. Eur J Neurosci 49:726–736.

Smith KS, Graybiel AM (2013) A dual operator view of habitual behavior reflecting cortical and striatal dynamics. Neuron 79:608.

Smith KS, Graybiel AM (2016) Habit formation coincides with shifts in reinforcement representations in the sensorimotor striatum. J Neurophysiol 115:1487–1498.

Steinberg EE, Gore F, Heifets BD, Taylor MD, Norville ZC, Beier KT, Földy C, Lerner TN, Luo L, Deisseroth K, Malenka RC (2020) Amygdala-Midbrain Connections Modulate Appetitive and Aversive Learning. Neuron Available at: http://dx.doi.org/10.1016/j.neuron.2020.03.016.

Stott JJ, Redish AD (2014) A functional difference in information processing between orbitofrontal cortex and ventral striatum during decision-making behaviour. Philos Trans R Soc Lond B Biol Sci 369:20130472.

Tang JCY, Paixao V, Carvalho F, Silva A, Klaus A, da Silva JA, Costa RM (2024) Dynamic behaviour restructuring mediates dopamine-dependent credit assignment. Nature 626:583–592.

Thorn CA, Atallah H, Howe M, Graybiel AM (2010) Differential dynamics of activity changes in dorsolateral and dorsomedial striatal loops during learning. Neuron 66:781–795.

Thorndike EL (1927) The law of effect. Am J Psychol 39:212.

Tom RL, Ahuja A, Maniates H, Freeland CM, Robinson MJF (2019) Optogenetic activation of the central amygdala generates addiction-like preference for reward. Eur J Neurosci 50:2086–2100.

van der Meer MAA, Redish AD (2009) Covert expectation-of-reward in rat ventral striatum at decision points. Front Integr Neurosci 3:1.

van Elzelingen W, Warnaar P, Matos J, Bastet W, Jonkman R, Smulders D, Goedhoop J, Denys D, Arbab T, Willuhn I (2022) Striatal dopamine signals are region specific and temporally stable across action-sequence habit formation. Curr Biol 32:1163-1174.e6.

Vandaele Y, Ahmed SH (2021) Habit, choice, and addiction. Neuropsychopharmacology 46:689–698.

Vandecasteele M, Royer S, Belluscio M, Berényi A, Diba K, Fujisawa S, Grosmark A, Mao D, Mizuseki K, Patel J, Stark E, Sullivan D, Watson B, Buzsáki G (2012) Large-scale recording of neurons by movable silicon probes in behaving rodents. J Vis Exp Available at: http://dx.doi.org/10.3791/3568-v.

Warlow SM, Naffziger EE, Berridge KC (2020) The central amygdala recruits mesocorticolimbic circuitry for pursuit of reward or pain. Nat Commun 11:2716.

Warlow SM, Robinson MJF, Berridge KC (2017) Optogenetic central amygdala stimulation intensifies and narrows motivation for cocaine. J Neurosci 37:8330–8348.

Yagishita S, Hayashi-Takagi A, Ellis-Davies GCR, Urakubo H, Ishii S, Kasai H (2014) A critical time window for dopamine actions on the structural plasticity of dendritic spines. Science 345:1616–1620.

Figure Captions

Figure 1. Extended training on a self-paced maze task promotes habitual responding. A) Animals were trained on a plus maze task where performance of a turn rule (e.g. turn right) resulted in delivery of a 20% sucrose reward (two flavors: orange or grape, two magnitudes: 0.1mL or 0.3mL). Rats (N = 17) were divided into two groups, a Criteria-trained (CT) group and an Overtrained (OT) group that received 7 more training sessions than their CT counterparts. Training consisted of initial acquisition (variable in number of sessions, until rats hit a 60% correct responding criteria), Criteria training (3 sessions), and Over-training (7 sessions for OT rats only). All animals underwent outcome devaluation, and post-devaluation probe tests in extinction (no outcomes delivered) and reacquisition (outcomes delivered). B) Correct responding over training sessions by group. C) Change in run speed by group, comparing run speed on the final day of training (CT animals: Session C3; OT animals: Session O7) to the first completed Criteria session (Session C1). D) Change in proportion of laps in each session that contained a VTE, comparing VTE incidence on the final day of training (C3 or O7) to VTE incidence on the first completed Criteria session (C1). E) Outcome consumption over devaluation sessions by group. F) Correct responding in extinction by group. G) Correct responding by session and group, showing final training session (C3 or O7) compared to the first reacquisition session. H) Outcome consumption by devaluation status (non-devalued or devalued) by group during the first reacquisition session. I) Correct responding over reacquisition session by group. Rat sex is indicated by shape (circle: female; triangle: male).

Figure 2. Overtraining increases cFos expression in the central nucleus of the amygdala. A) Rats (N = 8) were trained on a plus maze task where performance of a turn rule (e.g. turn right) resulted in delivery of a food reward (two flavors: banana or chocolate flavored sucrose pellet, two magnitudes: 1 or 3 pellets). B) Rats were divided into two groups, a Criteria-trained group (CT; N = 4) and an Overtrained group (OT; N = 4) that received 7 training sessions more than their Criteria counterparts. C) Animals were sacrificed 90 minutes after the conclusion of their final training session (C3 for Criteria or O7 for Overtrained rats) to allow for the expression and subsequent quantification of cFos. Left: no-run control representative histology; Center: Criteria-trained representative histology; Right: Overtrained representative histology. Top: grayscale DAPI images; Center: grayscale cFos images; Bottom: overlay images. Scale bar: 50 μm. D) Mean cFos density by training group and CeA subdivision. Error bars represent ± SEM. Rat sex is indicated by shape (circle: female; triangle: male).

Figure 3. Recording rats similarly formed fast and automatic responses on the maze task. A) Rats (N = 9) were trained on a plus maze task where they were required to turn a certain direction to earn outcomes, with unique outcome magnitudes (0.3 or 0.1 mL) and identities (orange or grape flavored 20% sucrose) delivered at each arm end. Training consisted of brief acquisition training, pre-implant, followed by Criteria training and Over-training before animals advanced through outcome devaluation and post-devaluation behavior probes. B) Silicon probes were implanted targeting the central nucleus of the amygdala. Colored bars represent the recording tracks from rats, colored by rat ID, with horizontal crossbars representing the approximate recording depth over sessions. The number of units isolated from each of these rats is displayed at the top of the panel to show distribution of cells across rats. C) Percentage of correct trials performed by session, plotted as the mean with shading representing the standard error of the mean (SEM). The dashed line reflects a 60%-correct response learning threshold. D) Correct trial rate by training session, plotted as the mean +/- SEM. E) Example movement paths on two trial types, one with and one without a vicarious trial error. Animal run speed is color coded. F) Mean IdPhi plotted over sessions with shading representing SEM. G) Proportion of laps (trials) that had a vicarious trial error, plotted as the mean +/- SEM over sessions. H) Feeder dwell time, in seconds, per trial shown as a mean +/- SEM over training sessions.

Figure 4. Central amygdala unit activity is dynamic over habit learning. A, B) Example peri-event time histograms from units at different stages of training. Colors represent differing outcome identities (purple: grape; orange: orange) and dot size represents outcome magnitude, overall responding shown in black. Left: O1 (approach) Right: C2 (outcome). C) Peri-event time average of normalized CeA unit activity, constructed from abutted +/- 200 ms peri-event periods across all

learning (pre-devaluation). Timepoint 0 represents the point at which the rat made a turn decision (peri-approach, left). Colored x's represent peak deviation from pseudo-baseline, with color indicating activation (blue) or inhibition (red). Early learning (red vertical line) and late learning (blue vertical line) indicate units that contributed to subsequent early and late learning analyses. D) Schematic depicting estimated recording locations for approach-modulated neurons. Vertical lines represent the probe track and horizontal lines indicate the estimated recording depths. Numbers above vertical probe lines indicate the total number of responsive units isolated. Colors represent individual rats. E) Peri- event time average of normalized CeA unit activity, constructed from abutted +/- 200 ms peri-event periods across all learning (pre-devaluation). Timepoint 0 represents the point at which the rat arrived at the feeder area (peri-outcome, right). F) Schematic depicting estimated recording locations for outcome-modulated neurons. G) During the approach time window, proportion of units that were non-modulated (gray), whose activity increased (green), or whose activity decreased (pink); data split into early learning (left) and late learning (right) bins. H) Units that were approach-modulated (21/153), outcome-modulated (36/153), modulated by both (12/153), or modulated by neither (84/153). I) During the outcome time window, proportion of units that were non-modulated (gray), whose activity increased (green), or whose activity decreased (pink); data split into early learning (left) and late learning (right) bins. J) Mean (± SEM) z-scored neural activity during the approach time window over learning phase shown by outcome flavor (left) and outcome magnitude (right). K) Mean (± SEM) z-scored neural activity during the outcome time window over learning phase shown by outcome flavor (left) and outcome magnitude (right). (** $p < 0.01$; *** $p < 0.001$)

Figure 5. A subset of central amygdala units are modulated by animal run speed. A) Mean animal speed over training sessions. B) Animal mean speeds plotted by outcome magnitude and outcome flavor. No significant effects of outcome flavor or magnitude on animal run speed. C) Density plot of animal speeds by outcome magnitude. D) Example plots showing animal speed (gray) and unit firing rate (color). Top: speed correlated unit; Middle: speed anti-correlated unit; Bottom: unit not modulated by animal speed. E) Histogram of units with speed-modulated units in red. F) Scatterplot showing that speed correlated units (n = 46) have a wide range of firing rates. G) Identified speed cells that were also approach-modulated (17/46), outcome-modulated (7/46), modulated by both (7/46), or neither (15/46) (top). Proportion of speed cells that were significantly modulated by outcome magnitude (76.1%) (bottom). H) Cell classification of all recorded units (n = 182) into two major cell types: regular-spiking (n = 175, green) and fast-spiking (n = 7, pink). Bottom left: proportion of regular-spiking units that were speed modulated (22.80%). Bottom right: proportion of fast-spiking units that were speed modulated (85.7%). I) Bar plot showing proportion of units recorded during each training phase that were observed to be positively (blue) or negatively (red) speed correlated. J) Speed decoding. Actual speed (black), decoded speed (orange), and shuffled + decoded speed (gray) traces over a 100-second test window for a representative CeA speed neuron. K) Decoding accuracy (log R^2) for speed cells, non-speed cells, and shuffled surrogates. L) Decoding accuracy (log R^2) for fast-spiking and regular-spiking speed cells. M) Decoding accuracy (log R^2) for speed-correlated and speed-anticorrelated cells. (** $p < 0.01$; *** $p < 0.001$)

Figure 6. Post devaluation behavioral measures and neural correlates. A) Behavioral schematic depicting outcome devaluation procedure, followed by post-devaluation testing sessions. B) Mean percentage of outcome consumed by post- devaluation session (no rewards available in Extinction) and outcome type (non-devalued vs devalued). C) Post-devaluation task performance plotted as mean percentage of correct trials by session. Sessions to the right of the dashed line are post- devaluation extinction (E1) and reacquisition (R1 – R3, rewarded). D) Animal run speed plotted as density plots by outcome value (gray: non-devalued; purple: devalued) and outcome magnitude (top: small magnitude; bottom: large magnitude). E) Peri-event time averages of recorded neural activity (n = 29), presented as z-scored means. Top: time(0) = turn choice point. Bottom: time(0) = feeder arrival. F) Proportion of record cells that were approach-modulated (top) or outcome-modulated (bottom). Non-modulated neurons in gray, increase in neural activity in green, and decrease in activity in pink. G) Peri-decision point z-scored neural activity over reacquisition sessions, divided by approach to the non-devalued (left) or devalued (right) outcome arms and outcome magnitude (darker gray = larger outcome). H) Peri-outcome z-scored neural activity over reacquisition sessions, divided by non-devalued (left) and devalued (right) outcome identities, with the larger magnitude outcome in darker gray.