

Cognition and Behavior

Differential Involvement of Three Brain Regions during Mouse Skill Learning

Aldis P. Weible,¹ Michael I. Posner,² and Cristopher M. Niell³<https://doi.org/10.1523/ENEURO.0143-19.2019>¹Institute of Neuroscience, University of Oregon Eugene, Oregon, 97403, ²Department of Psychology, University of Oregon Eugene, Oregon, 97403, and ³Department of Biology, University of Oregon, Eugene, Oregon, 97403

Abstract

Human skill learning is marked by a gradual decrease in reaction time (RT) and errors as the skill is acquired. To better understand the influence of brain areas thought to be involved in skill learning, we trained mice to associate visual-spatial cues with specific motor behaviors for a water reward. Task acquisition occurred over weeks and performance approximated a power function as often found with human skill learning. Using optogenetics we suppressed the primary visual cortex (V1), anterior cingulate cortex (ACC), or dorsal hippocampus (dHC) on 20% of trials at different stages of learning. Intermittent suppression of the V1 greatly reduced task performance on suppressed trials across multiple stages but did not change the overall rate of learning. In accord with some recent models of skill learning, ACC suppression produced higher error rates on suppressed trials throughout learning the skill, with effects intensifying in the later stages. This would suggest that cognitive influences mediated by the anterior cingulate continue throughout learning. Suppression of the hippocampus only modestly affected performance, with largely similar effects seen across stages. These results indicate different degrees of V1, ACC, and dHC involvement in acquisition and performance of this visual-spatial task and that the structures operate in parallel, and not in series, across learning stages.

Key words: anterior cingulate cortex; hippocampus; skill learning; visual cortex

Significance Statement

Mice resemble humans with improvements in accuracy and speed during skill learning. Through optogenetics, we can suppress different regions of the mouse brain at different stages of training to better understand when each region contributes to learning. Here we found that primary visual cortex (V1) suppression reduced accuracy across all training stages. Suppressing anterior cingulate cortex (ACC), a region thought to be important for attention early in training, also reduced accuracy throughout learning. Suppressing the hippocampus, a structure critically involved in associative learning, affected performance more modestly. These findings reveal parallel, rather than serial, involvement of these three structures in a mouse model of skill learning.

Introduction

The study of human skill learning examines both the processes involved in individual trials as well as the

stages through which the person passes from initial acquisition to final performance. For example, [Sternberg](#)

Received April 15, 2019; accepted July 25, 2019; First published August 1, 2019.

The authors declare no competing financial interests.

Author contributions: A.P.W. performed research; A.P.W. analyzed data; A.P.W. and M.I.P. wrote the paper; M.I.P. and C.M.N. designed research.

This work was supported by the Office of Naval Research Grant N00014-15-1-2148.

Correspondence address should be addressed to Aldis P. Weible at aweible@uoregon.edu.

<https://doi.org/10.1523/ENEURO.0143-19.2019>

Copyright © 2019 Weible et al.

This is an open-access article distributed under the terms of the [Creative Commons Attribution 4.0 International license](#), which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.

(1969) proposed that an individual trial involved the processes of encoding input, searching memory, selecting the response and executing the response, a sequence that is supported by recent imaging data (Tenison et al., 2016). On the broader scale, Fitts and Posner (1967) proposed that acquiring expertise in a skill involved a progression through cognitive, associative, and autonomous stages. This influential three-stage model has received significant experimental support, and was most recently validated using the ACT-R model developed by John Anderson and associates (Tenison and Anderson, 2016).

According to these models, stages of skill learning progress in serial order. But how do different brain regions engage during learning to facilitate the transition from the novice to the expert? One possibility is that involvement of different brain areas might, like the three-stage model proposed by Fitts and Posner, occur in a strictly serial order. Alternatively, different structures might operate fully in parallel. For example, multiple nodes of the neural circuit might be engaged at the start of training, with connections strengthening over time with trial repetition (Frey and Morris, 1997, 1998; Redondo and Morris, 2011). A third possibility is that, while different regions might be maximally active at different times, their patterns of activation overlap, thus representing a fusion of the serial and parallel models (Formisano and Goebel, 2003; Tenison et al., 2016). The goal of the present study was to compare serial versus parallel involvement of three brain regions, the primary visual cortex (V1), anterior cingulate cortex (ACC), and dorsal hippocampus (dHC) in a mouse model of skill learning.

To accomplish this goal, we trained mice running on a spherical treadmill to respond to the location of a visual stimulus (above or below center on a computer monitor) with a motor action (running to the left or right) for a water reward. This skill requires several weeks for mice to attain mastery. We examined the improvement of the mice across sessions in accuracy, bias, and reaction time (RT) to derive an overall view of learning. We used optogenetic control of parvalbumin-expressing interneurons (PV-INs) to suppress excitatory activity at different stages of training in V1, ACC, and dHC. Because PV-INs provide broad and potent inhibitory control over pyramidal cells, their activation effectively inhibits global levels of excitatory output (Lien and Scanziani, 2011; Atallah et al., 2012). We used a PV-Cre driver line together with a Cre-dependent channelrhodopsin-2 (ChR2) line (Boyden et al., 2005) to provide cell-type specific depolarization with a high degree of temporal specificity. In the PV-ChR2 cross, light delivered through optic fibers increases PV-IN firing, thus suppressing regional excitatory activity.

We first considered how average improvement in performance compared with human skill learning. Next, we sought to determine how each region of interest was involved at different stages of training. We suppressed activity in each of the three brain areas on 20% of the trials during early, middle and late stages of learning based on specific performance criteria, as well as a final over-trained stage. By suppressing performance on only a

small number of trials we sought to compare performance between suppressed and unsuppressed trials within stages while not impairing task acquisition.

V1 was selected for suppression because of the use of visual cues in the task. An effect of V1 suppression would confirm that the task was forebrain dependent, a result which would in turn be crucial to interpreting the presence or absence of suppression effects in either of the other two regions. The ACC was selected because it plays a central role in high level attention and top-down cognitive control (Botvinick et al., 2001; Petersen and Posner, 2012), and has been implicated previously in different forms of skill learning and performance (Raichle, 1998; Procyk et al., 2000; Poldrack and Gabrieli, 2001; Karim et al., 2017). The dHC, which also exhibits changes in activation with skill learning (Poldrack and Gabrieli, 2001; Karim et al., 2017), was selected because it is crucial for many forms of associative learning (Brasted et al., 2003; Suzuki, 2007; Squire and Wixted, 2011). Our approach should shed light on whether, when, and how each region is involved in a visual-spatial discrimination mouse model of skill learning.

Materials and Methods

Mice

All procedures were conducted in accordance with the ethical guidelines of the National Institutes of Health and were approved by the Institutional Animal Care and Use Committee at the University of Oregon. Animals were maintained on a reverse 12/12 h light/dark cycle. Training and experiments were performed during the dark phase of the cycle. All mice were male, 8–12 weeks of age at the time of surgery, and were bred to the C57Bl6/J background strain. Optogenetic suppression of excitatory activity was performed in offspring (total $n = 41$; V1 $n = 8$, ACC $n = 9$, dHC $n = 10$, control $n = 14$) from a cross between homozygotic Pvalb-IRES-Cre (“PV”, 008069; The Jackson Laboratory) and Rosa26-CAG-LSL-ChR2H134R-eYFP (“ChR2”, Ai32, 012569; The Jackson Laboratory) lines. In these mice (PV-ChR2), ChR2 was expressed in PV-INs.

Surgery

We administered atropine (0.03 mg/kg) preoperatively to reduce inflammation and respiratory irregularities. Surgical anesthesia was induced and maintained with isoflurane (1.25–2.0%). Optogenetic suppression of neural activity was achieved by passing light through 200 μm in diameter optic fibers implanted bilaterally over the caudal ACC (AP: 0.6 mm, ML: 0.4 mm, DV: 0.4 mm, 15° angle toward the midline), V1 (AP: –3.5 mm, ML: 2.2 mm), or dHC (AP: –2.0 mm, ML: 2.0 mm, DV: 0.8 mm). A titanium cross-bar was also fixed to the skull to enable restraint of the mouse during behavior training. Fibers and cross-bar were cemented in place with Grip Cement. Carprofen (10 mg/kg) was administered postoperatively to minimize discomfort. Mice were housed individually after the surgery and allowed 7 d of postoperative recovery.

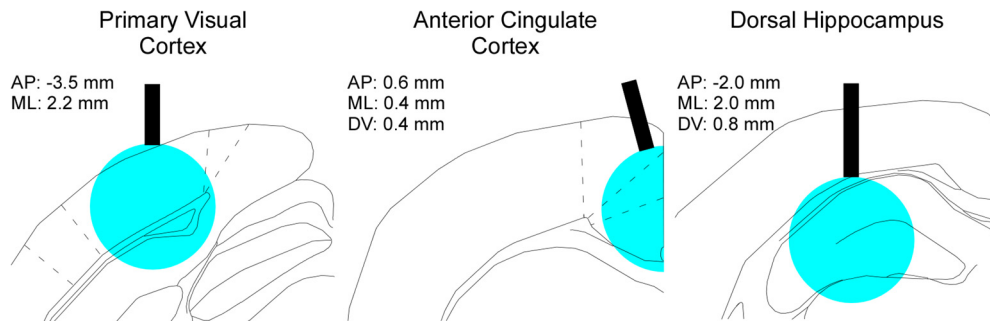


Figure 1. Chronically-implanted optic fibers enable light delivery to V1, ACC, and dHC. Optic fibers (200 μm in diameter; illustrated by the black rectangle) were implanted for delivery of 450-nm wavelength “blue” light to stimulate PV-INs, resulting in the suppression of excitatory output. The expected volume of suppression is a region of tissue 1.5 mm in diameter (illustrated by the blue circle; Weible et al., 2014). AP: anterior (+) or posterior (–) to bregma; DV: ventral to the cortical surface; ML: lateral to the midline.

Optogenetic suppression

We suppressed excitatory activity in PV-ChR2 mice using a 450-nm wavelength diode laser set to an output power of 6.3 mW. The resulting intensity of 200 mW/mm^2 , as measured at the tip of the 200 μm in diameter fiber, results in suppression through a volume ~ 1.5 mm across (Weible et al., 2014). The expected spread of light through each of the three target regions is illustrated in Figure 1. Rise/fall times of laser pulses were 5 μs . We measured laser power and rise/fall times with a Thorlabs PM100D power meter. Suppression was induced during 20% of trials in a pseudorandomly determined manner. Suppression was initiated at the moment the mouse successfully requested a trial, and continued until a movement was executed (or 10 s had elapsed).

Behavioral apparatus and training

Mice were trained to perform a visual-spatial discrimination task for a water reward. Training took place on a spherical treadmill based on designs by Dombeck et al. (2007) and Niell and Stryker (2010). Briefly, a hollow polystyrene ball (250 mm in diameter) was set in a 3D printed hemi-spherical base. Air delivered through the bottom of the base was adjusted to a rate of flow enabling free rotation of the ball. A two-tined fork, extending from the back and over the top of the ball, provided the anchor points to fix the mouse in position on the ball. Visual stimuli were presented on a ViewSonic VA2342 liquid crystal display monitor [28 \times 50 cm; linearized by eye to correct for gamma (mean luminance 35 cd/m^2)], set 25 cm in front of the mouse. Water rewards were delivered through a spout positioned immediately beneath the mouth of the mouse. An optical mouse set $\sim 15^\circ$ below the horizontal in front of the ball was used to detect movement.

Before training, mice were handled for 3 d, and weights for each mouse were collected to establish a baseline for the water restriction. Weights were subsequently collected daily throughout training, and mice falling below 75% of baseline weight were given supplemental water until weights rose back above the minimum 75% of baseline threshold. Having established baseline weights, mice were then placed on the ball for 2 d of familiarization to the apparatus (day 1: ~ 5 min; day 2: ~ 30 min).

Task performance required first “requesting” a trial by remaining immobile for 1 s, then responding based on the stimulus by moving the ball to the left or right to receive the water reward. Training was performed in two phases. During the first phase, mice learned to request trials (i.e., remaining immobile for 1 s) by receiving “stop” water rewards. This phase lasted 3 d, with stop rewards decreasing in volume across days. During the second phase, water rewards were given only for correct responses, and were reduced in volume as performance improved to maximize trial count per session. Visual cues determined the correct movement selection. Cues in the present task were 15.5 cm in diameter circles containing vertical or horizontal bars positioned at the top or bottom of the screen. Only the spatial location of the circle was predictive of water reward (bar orientation was irrelevant, and alternated pseudorandomly between horizontal and vertical). Rotation of the ball to the left was rewarded when stimuli were presented at the top of the screen (and to the right for stimuli at the bottom). Errors were signaled by a full-field bright flash (two-fold increase in luminance over gray background) that persisted for 1 s.

Stages of training

Training progressed through four stages. Stage 1 began with the cessation of stop rewards, when mice were rewarded only for correct responses. Stages 2 and 3 began after mice achieved specific criterion levels of performance. Stage 2 began the day after mice first performed $>65\%$ correct responses during a block of 50 consecutive trials. Stage 3 began the day after mice performed $>85\%$ correct responses in a 50-trial block, marking the point in training at which mice were considered to have “learned” the task. Following stage 3, the percentage of correction trials was gradually reduced to 0, with no laser suppression. When correction trials reached 0%, suppression trial data were collected for stage 4.

While these stages are somewhat arbitrary, they fit rather closely to mean RT and error data as shown in Figure 2. The mice differed greatly in how many days were required to reach criterion. To plot the average performance we used Vincent curves (Hilgard and Campbell, 1937) which plot the RT and accuracy in each tenth of the

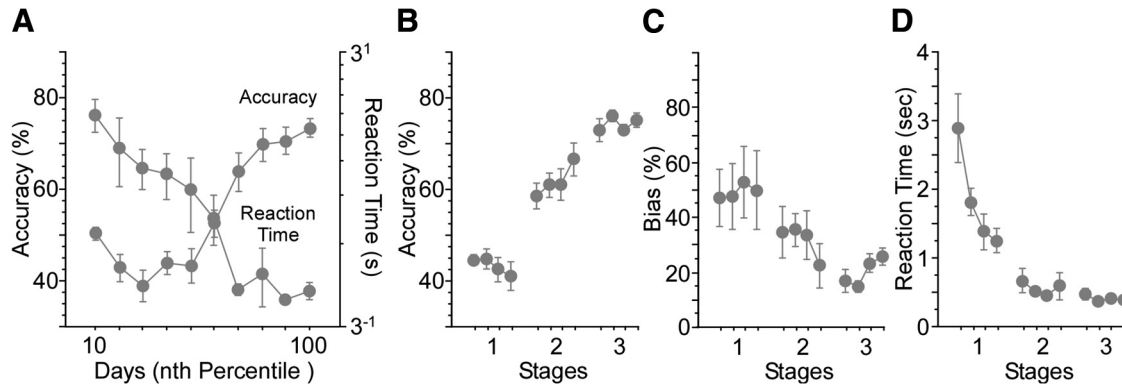


Figure 2. Accuracy and RT performance measures recapitulate changes observed with human skill learning. Control mice acquired the visual spatial discrimination task without interleaved suppression trials. **A**, The decrease in RT seen as accuracy increases approximates the power function found with human skill learning. **B–D**, Accuracy, bias, and RT measures are plotted by the 4 d of each stage. Response bias trended lower over the three stages, and RT decreased significantly (bias: $df = 11, \chi^2 = 18.5, p = 0.07$; RT: $df = 11, \chi^2 = 61.3, p < 0.0001$). Analyses for each stage included 4 d of training. Stage 1 included the first 4 d of training. Stage 2 included the first 4 d following performance of 65% correct responses in a 50-trial block. Stage 3 included the 4 d following performance of 85% correct responses in a 50-trial block.

overall trials to learning criterion. The first few values in the Vincent curve show marked improvement in RT but accuracy remains at or below chance. This is closely related to what is defined above as the first stage of training. In this stage the association is still unlearned but the general task shows an improvement in RT as different strategies appropriate to the task are selected (e.g., adopting a strong bias, increasing trial count, correctly associating stimulus position with rewarded movement). Later, both RT and accuracy improve. This includes the 65% accuracy defining stage 2 where the mice appear to be learning the correct association while still performing slower than the final RT. Stage 3 is defined by 85% accuracy and corresponds closely to points which show no further improvement in RT or in accuracy, suggesting that performance has reach a kind of asymptote. As defined, stage 4 occurs after mice have reach criterion and correction trials are suspended.

Optogenetic suppression of principal neuron activity was conducted only for the first 4 d of each stage, as this manipulation might impact behavior not only on a trial-by-trial basis, but the overall rate of learning as well. Therefore, each “stage” includes data from 4 d of training. However, achieving each criterion level of performance typically took >4 d of training, and the total days trained varied between mice. Infrequently, mice achieved the next performance criterion in fewer than 4 d. In these cases, mice were considered to have transitioned to the next stage only after the four consecutive days of the previous stage had been completed. In a few instances, mice also failed to learn the task. Mice failing to transition to stage 3 within 50 d were excluded from analysis.

Correction trials, wherein the stimulus on the previous trial was repeated, were used to counter the development of response bias during training. At the start of training, 50% of incorrect responses were followed by a correction trial. This level was maintained or adjusted higher over the course of training, based on the strength of response bias observed, but was never adjusted during the four consec-

utive days of suppression, as suppression itself could directly bias response selection. To assess the effects of suppression in the absence of correction trials, following the four consecutive days of stage 3, correction trials were gradually reduced to 0%, and mice transitioned to stage 4: four consecutive days of training with interleaved suppression of principal neuron activity in the absence of correction trials. Mice in stage 4 are referred to as over-trained.

To test for effects of suppression on overall rates of learning, data were also collected from control mice that were implanted with fibers but did not undergo suppression during stages 1–3. In this way, we could compare trials-to-learning with and without interleaved suppression. Following the final day of stage 3, however, control mice transitioned to stage 4 with suppression trials. This enabled a comparison to determine whether suppression at different intervals during training impacted effects seen following the removal of correction trials.

Data analysis

Accuracy data were based on percentage correct responses for suppressed and non-suppressed trials for each session. Analysis of RT (defined as time between visual cue display and mouse movement) data were performed using median RT for each session. To determine bias, we calculated the percentage correct to left (L) and right (R), subtracted each from the percentage correct for the entire session (Session%), and added the absolute value of both.

$$\text{Bias}_{L,R} = \text{Session\%} - (\#\text{Correct}_{L,R} / \text{Total}_{L,R})$$

$$\text{Total Bias} = |\text{Bias}_L| + |\text{Bias}_R|.$$

The result yields a value where 0% represents no bias and 100% a complete bias to one side.

We chose to use non-parametric analyses in the present study because many of the comparisons involved data that were not normally distributed (verified with the

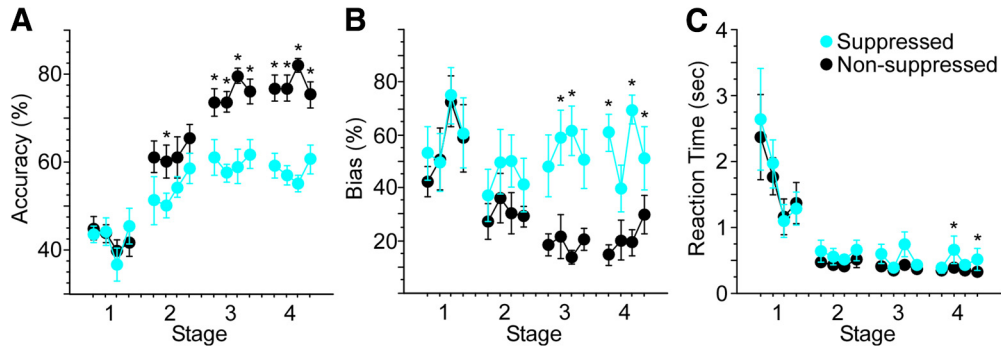


Figure 3. V1 suppression significantly reduces accuracy and increases response bias and RT. Light was delivered to V1 to suppress excitatory activity on 20% of trials through stage 4. **A**, Suppression reduced accuracy across stages compared with non-suppressed trials (Friedman's, $df = 1$, $\chi^2 = 30.3$, $p < 0.0001$). **B**, Response bias increased on suppressed versus non-suppressed trials (Friedman's, $df = 1$, $\chi^2 = 36.1$, $p < 0.0001$). **C**, Suppression produced a small but significant increase in RT across stages (Friedman's, $df = 1$, $\chi^2 = 6.6$, $p = 0.01$). Analyses for each stage included 4 d of training. Stage 1 included the first 4 d of training. Stage 2 included the first 4 d following performance of 65% correct responses in a 50-trial block. Stage 3 included the first 4 d following performance of 85% correct responses in a 50-trial block. Following stage 3, correction trials were reduced gradually to 0%, at which point mice received an additional 4 d of overtraining (stage 4); * $p < 0.05$, ** $p < 0.01$.

Lilliefors test), and because statistical power is comparable even when the underlying assumptions for the corresponding parametric analysis were met (Kitchen, 2009). Friedman's test was used to identify within-group effects of suppression across days, as well as changes across days separately for suppressed and non-suppressed trials. *Post hoc* analysis of differences on individual days was performed using a Wilcoxon paired sign test. Between-group comparisons were performed using the Kruskal–Wallis test. To determine whether differences were present between any pair of groups within a given stage (as a *post hoc* analysis of Kruskal–Wallis test results), days within the stage were compared using the Mann–Whitney *U* test.

Histology

Following stage 4, mice with fibers targeting the ACC or dHC were administered a lethal dose of sodium pentobarbital (100 mg/kg Euthasol) and transcardially perfused with phosphate buffered saline (0.9%) and the 4% paraformaldehyde (PFA). Brains were removed and post-fixed for a minimum of two additional days in 4% PFA, and then sectioned at 100 μm on a vibratome. Fiber placement was then verified with fluorescent microscopy. Two mice were excluded because fibers targeting the hippocampus were found to have terminated either within or beneath the level of the dentate gyrus. All other implanted mice had fiber tracks histologically verified as accurately targeting the ACC or dHC. No attempt was made to histologically verify the placement of fibers targeting V1 as these fibers were placed on the surface on the dura.

Results

In the present study, we used an animal model of skill learning, a top-bottom visual-spatial discrimination task performed by mice for a water reward, to examine whether and when V1, the ACC, and dHC were involved in task acquisition and performance. Specifically, we examined how optogenetic suppression of these structures influenced movement selection, response bias, and RT.

Relationship to human skill learning

Before considering the results of our suppression experiments, it is worth addressing whether the mouse data resembles that from human studies. In human skill learning, changes in performance and RT averaged over trials and subjects are often described as fitting a power function (Fitts and Posner, 1967; Newell and Rosenbloom, 1981; Anderson et al., 1999). For RT at least, this feature is recapitulated in mice performing the top-bottom discrimination task. We trained a group of fiber-implanted mice up through stage 3, without suppressing activity in any of the target regions. Figure 2A illustrates performance and RT data. Performance data on this scale exhibits a complex pattern, skewed early in training by transient response biases (as illustrated also in Fig. 2B,C). However, the changes in RT are quite similar to those reported in the human skill learning literature, suggesting at least a superficial similarity over the course of learning.

For the purposes of the present study, we organized the control data into the performance-based training stages used in the suppression experiments below. All of the control mice trained achieved the >85% criterion marking stage 3 (Fig. 2B). Response bias trended lower over the three stages, and RT decreased significantly (Friedman's, bias: $df = 11$, $\chi^2 = 18.5$, $p = 0.07$; RT: $df = 11$, $\chi^2 = 61.3$, $p < 0.0001$; Fig. 2C,D).

Suppression effects by group

All eight mice undergoing V1 suppression during training successfully acquired the task. Performance on suppressed trials was significantly lower than non-suppressed trials (Friedman's, $df = 1$, $\chi^2 = 30.3$, $p < 0.0001$; Fig. 3A). The impact of suppression was initially evident during stage 2, and grew more robust during stages 3 and 4. Suppression also influenced response bias. As illustrated in Figure 3B, strong response biases often developed early in training. These were addressed following stage 1 by increasing the percentage of correction trials. As mice learned the relationship between cue position and rewarded direction of movement, biases

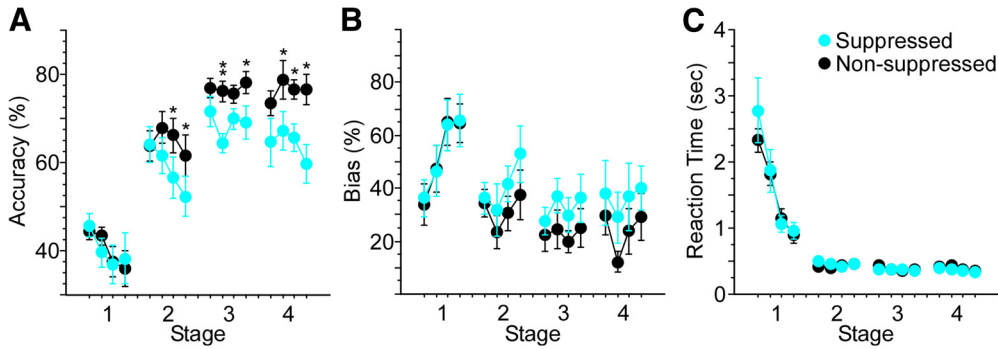


Figure 4. ACC suppression significantly reduces accuracy and increases response bias. Light was delivered to ACC to suppress excitatory activity on 20% of trials through stage 4. **A**, Suppression reduced accuracy across stages compared with non-suppressed trials (Friedman’s, $df = 1, \chi^2 = 9.6, p = 0.002$). **B**, Response bias increased on suppressed versus non-suppressed trials (Friedman’s, $df = 1, \chi^2 = 4.6, p = 0.033$). **C**, Suppression did not significantly affect RT across days of training. Analyses for each stage included 4 d of training. Stage 1 included the first 4 d of training. Stage 2 included the first 4 d following performance of 65% correct responses in a 50-trial block. Stage 3 included the first 4 d following performance of 85% correct responses in a 50-trial block. Following stage 3, correction trials were reduced gradually to 0%, at which point mice received an additional 4 d of overtraining (stage 4); * $p < 0.05$, ** $p < 0.01$.

decreased on non-suppressed trials, exhibiting a strong trend toward significance across all four stages of training (Friedman’s, $df = 15, \chi^2 = 25.0, p = 0.05$). As biases decreased, the percentage of correction trials was reduced. With suppression, however, biases remained elevated (Friedman’s, $df = 1, \chi^2 = 36.1, p < 0.0001$), and were on par with those seen during stage 1. Finally, suppression increased RT. Median RTs decreased significantly across stages for both non-suppressed (Friedman’s, $df = 15, \chi^2 = 40.6, p = 0.0001$) and suppressed (Friedman’s, $df = 15, \chi^2 = 29.5, p = 0.01$) trials. This decrease in RT was modestly, but significantly, less with suppression (Friedman’s, $df = 1, \chi^2 = 6.6, p = 0.01$).

All nine mice undergoing ACC suppression during training successfully acquired the task. Performance on suppressed trials was significantly lower than non-suppressed trials (Friedman’s, $df = 1, \chi^2 = 9.6, p = 0.002$; Fig. 4A). As with suppression in V1, the impact of ACC suppression was evident first during stage 2, and grew more robust during stages 3 and 4. ACC suppression also influenced response bias (Fig. 4B). While biases de-

creased across days for both non-suppressed and suppressed trials (Friedman’s, $df = 15, \chi^2 = 52.3, p < 0.0001$ and $\chi^2 = 32.2, p = 0.006$, respectively), response bias remained greater with suppression (Friedman’s, $df = 1, \chi^2 = 4.6, p = 0.033$). RT decreased significantly across stages during both non-suppressed and suppressed trials (Friedman’s, $df = 15, \chi^2 = 97.2, p < 0.0001$ and $\chi^2 = 96.8, p < 0.0001$, respectively; Fig. 4C), and was not affected by suppression ($p = 0.54$).

Eight of 10 mice undergoing suppression of dHC during training successfully acquired the task (the 2 mice failing to reach stage 3 criterion performance within 50 d of training were excluded from analysis). Performance on suppressed trials was lower than non-suppressed trials across stages (Friedman’s, $df = 1, \chi^2 = 4.8, p = 0.03$; Fig. 5A). No differences were observed for individual days (Wilcoxon signed rank). Biases for both non-suppressed and suppressed trials decreased significantly across days (Friedman’s, $df = 15, \chi^2 = 41.8, p = 0.0002$ and $\chi^2 = 25.5$ and $p = 0.04$, respectively; Fig. 5B), with a weak trend toward a difference between conditions ($p = 0.15$). RT for

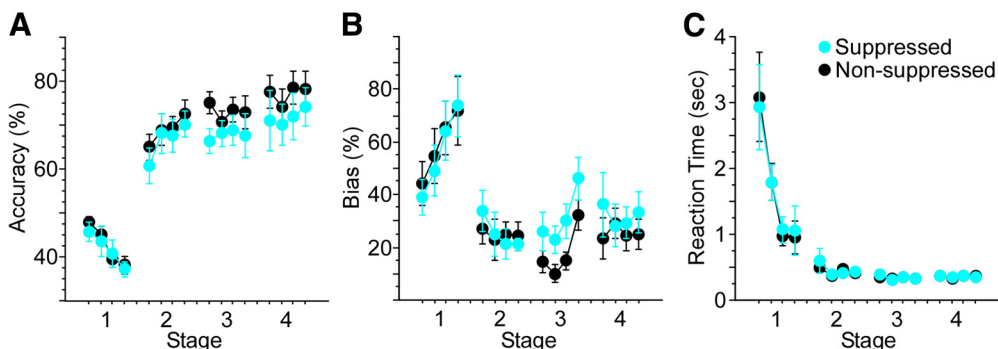


Figure 5. Suppression of dHC modestly reduces response accuracy. Light was delivered to dHC to suppress excitatory activity on 20% of trials through stage 4. **A**, Suppression reduced accuracy across stages compared with non-suppressed trials (Friedman’s, $df = 1, \chi^2 = 4.8, p = 0.03$). Neither response bias (**B**) nor RT (**C**) changed significantly with suppression. Analyses for each stage included 4 d of training. Stage 1 included the first 4 d of training. Stage 2 included the first 4 d following performance of 65% correct responses in a 50-trial block. Stage 3 included the first 4 d following performance of 85% correct responses in a 50-trial block. Following stage 3, correction trials were reduced gradually to 0%, at which point mice received an additional 4 d of overtraining (stage 4).

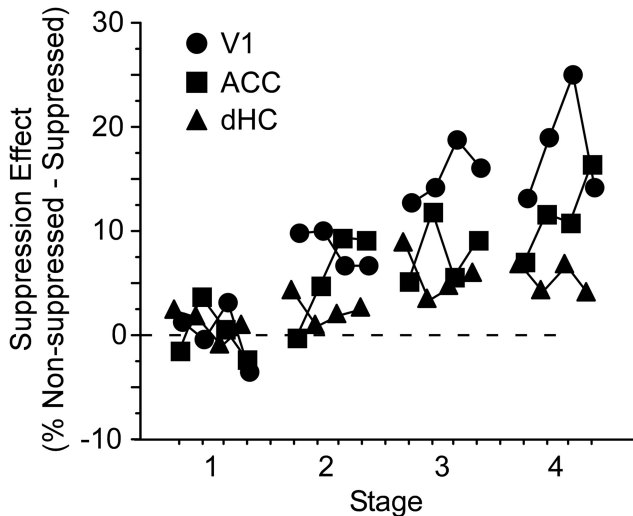


Figure 6. Decreases in accuracy grow over time with suppression of ACC and V1, but not dHC. Suppression effect size was calculated by subtracting accuracy of suppressed trials from non-suppressed trials. The size of the suppression effect increased significantly across days for V1 and ACC mice (Friedman's, $df = 15$, $\chi^2 = 62.4$, $p < 0.0001$ and $df = 15$, $\chi^2 = 25.7$, $p = 0.04$, respectively), but not dHC mice. Differences in effect size were significant between groups during stages 3 and 4 (Kruskal–Wallis, $df = 2$, $p = 0.002$ and $df = 2$, $p < 0.0001$, respectively), with effect size in V1 exceeding that of ACC during stages 3 and 4 (Mann–Whitney, $z = -2.7$, $p = 0.006$ and $z = -2.4$, $p = 0.02$, respectively), and ACC exceeding that of dHC during stage 4 (Mann–Whitney, $z = -2.4$, $p = 0.02$).

both non-suppressed and suppressed trials decreased significantly across days (Friedman's, $df = 15$, $\chi^2 = 73.4$, $p < 0.0001$ and $\chi^2 = 76.9$, $p < 0.0001$, respectively; Fig. 5C) with no significant difference between conditions ($p = 0.98$).

Comparisons between groups

As illustrated in Figures 3–5, the size of the suppression effect on accuracy (accuracy on non-suppressed minus suppressed trials) varied both across days and between groups. Suppression of V1 and ACC had an increasingly large effect on accuracy over time (Friedman's, $df = 15$, $\chi^2 = 62.4$, $p < 0.0001$ and $df = 15$, $\chi^2 = 25.7$, $p = 0.04$, respectively; Fig. 6). The effect size with dHC suppression did not change significantly across days ($p = 0.90$).

The difference in effect size between groups that was first apparent during stage 2 grew to significance during stages 3 and 4 (Kruskal–Wallis, $df = 2$, $p = 0.002$ and $df = 2$, $p < 0.0001$, respectively). Effects of V1 suppression on accuracy were greater during stages 3 and 4 than those seen with ACC suppression (Mann–Whitney, $z = -2.7$, $p = 0.006$ and $z = -2.4$, $p = 0.02$, respectively), and the size of effect seen with ACC suppression during stage 4 exceeded that seen in dHC (Mann–Whitney, $z = -2.4$, $p = 0.02$). A difference in percentage correction was observed between suppressed groups during stage 3 (Kruskal–Wallis, $df = 2$, $p = 0.0002$; Table 1). However, this is unlikely to explain the differences in effect size as no correction was employed during stage 4.

Table 1. % Correction trials

	V1	ACC	dHC	Control
Stage 1	50	50	50	50
Stage 2	69.5 ± 0.09	65.4 ± 0.02	65.5 ± 0.01	78.1 ± 0.02
Stage 3	47.5 ± 0.01	26.7 ± 0.04	35.8 ± 0.03	30.0 ± 0.03
Stage 4	0	0	0	0

All values % ± SE.

While suppression of spiking activity imposed in each of these regions was intended to be trial specific, it is nonetheless possible that this manipulation could have had broader consequences on learning and performance. To test whether suppression impacted rates of acquisition, we compared trials-to-criterion for control mice and V1, ACC, and dHC suppressed mice. No significant effect of group was observed (Kruskal–Wallis, $p = 0.29$; Fig. 7A). Accuracy on non-suppressed trials across the first three stages also did not differ between the four groups (Kruskal–Wallis, $p = 0.83$).

These data indicate that trial-specific suppression did not influence rates of acquisition more broadly. It is also possible that the efficacy of suppression might change over time. Because control mice were implanted with fibers, we were able to suppress activity in these mice for the first time at stage 4 and compare the effects with those seen at stage 4 in V1, ACC, and dHC mice (i.e., mice that had undergone suppression during the previous three stages). Suppression of V1 or ACC in control mice significantly reduced accuracy at stage 4 (Friedman's, $df = 1$, $\chi^2 = 9.1$, $p = 0.004$ and $df = 1$, $\chi^2 = 16.9$, $p < 0.0001$, respectively). The size of these effects was comparable to those seen with V1 and ACC mice that had undergone suppression through stage 3 (Fig. 7B,C). Finally, accuracy of control mice was unaffected by suppression of dHC during stage 4 (Fig. 7D).

RT and suppression

In contrast to accuracy and response bias, RT across days only changed with suppression of V1 (Fig. 3). However, RT on any given trial is often influenced by the outcome of the preceding trial (Laming, 1979; Rabbitt, 1966). We therefore examined first the relationship between error trials and RT in the present data, and then tested for any impact of suppression in V1, ACC, and/or dHC following an incorrect choice.

We compared correct trial RTs on trial $N + 1$ following correct and incorrect (error) responses on trial N . RTs were consistently greater on trial $N + 1$ following an error, for both non-suppressed and suppressed trials (Wilcoxon, $z = -16.7$, $p < 0.0001$ and $z = -12$, $p < 0.0001$, respectively). Suppression lengthened trial $N + 1$ RTs following errors (Wilcoxon, $z = -4.1$, $p < 0.0001$) but had no effect following correct responses ($z = -0.7$, $p = 0.51$). These patterns could be broken down further by region of suppression and stage of training. In V1, suppression lengthened trial $N + 1$ RTs following errors during stage 2 and stage 4 (Wilcoxon, $z = -2.7$, $p = 0.006$ and $z = -2.4$, $p = 0.01$, respectively; Fig. 8A, dashed black lines). In ACC, post-error suppression produced longer RTs during stage 2 (Wilcoxon, $z = -2.3$, $p = 0.02$; Fig. 8B, dashed

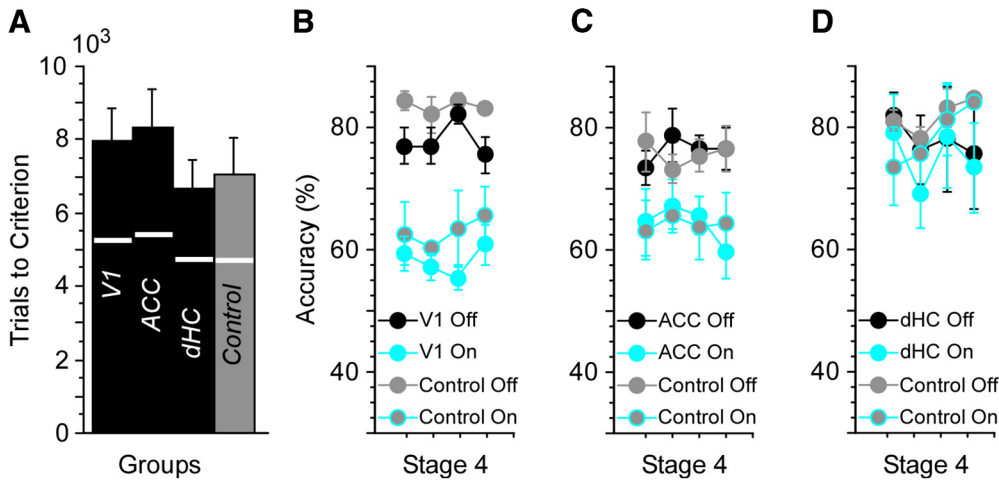


Figure 7. Suppression effects on accuracy are trial specific. **A**, Control mice and mice undergoing V1, ACC, and dHC suppression learned to stages 2 and 3 criteria in a comparable number of trials (accuracy criteria: >65% and >85% correct responses in a 50-trial block, respectively). Trials to stage 2 indicated by the white line bisecting each bar. Control mice did not undergo suppression over the course of training, but were implanted with fibers targeting V1, ACC, or dHC. During stage 4, suppression was performed in V1, ACC, or dHC mice and compared with effects seen in controls with matching fiber placements undergoing suppression for the first time. The results demonstrate that suppression over the course of training in V1 (**B**), ACC (**C**), or dHC (**D**) did not have a subsequent effect on suppression performed after mice reached the >85% criterion.

black line). In dHC, suppression following an error produced longer RTs during stage 3 (Wilcoxon, $z = -2.8, p = 0.004$; Fig. 8C, dashed black line). In no region did suppression increase RTs following correct choices. We also examined the effect of suppression during trial N to determine whether mice had a smaller RT during trial $N + 1$ following an error when the error might have been due to suppression. This could happen particularly with suppression of V1 since mice would be less influenced by visual

input. However, in no condition did we find that suppression of trial N influenced RT on trial $N + 1$. This may be because the cues introduced after the subject response gave the mice sufficient information about whether they were correct even with suppression.

Summary

In the present study, we sought to examine serial versus parallel involvement of three brain regions of interest

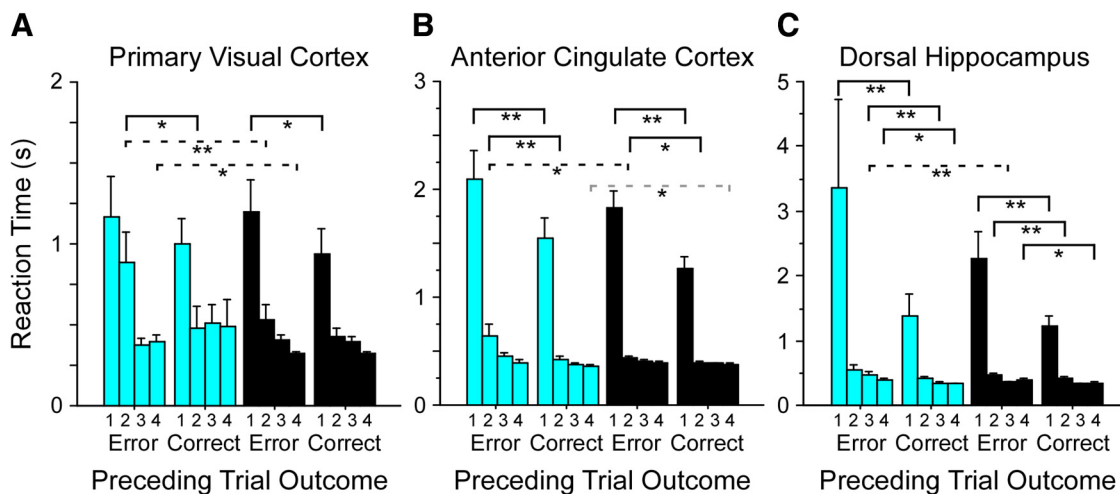


Figure 8. Suppression increases RT following errors. Errors during training were signaled by an increase in screen luminance that persisted for 1 s. RT on correct responses was longer following errors than following correct responses. Suppression further lengthened RTs following errors but had no effect on RT following correct responses. These patterns were observed for mice from each suppression group, with performance on suppressed trials indicated by blue bars, and non-suppressed trials by black bars. **A**, In V1, suppression lengthened RTs following errors during stages 2 and 4 (dashed black lines; Wilcoxon, $z = -2.7, p = 0.006$ and $z = -2.4, p = 0.01$, respectively). **B**, In ACC, suppression produced longer RTs during stage 2 (dashed black line; Wilcoxon, $z = -2.3, p = 0.02$). **C**, In dHC, suppression following an error produced longer RTs during stage 3 (dashed black line; Wilcoxon, $z = -2.8, p = 0.004$). Dashed gray line: increased RT on non-suppressed trials. Solid black line: increased RT following error versus correct trials, within either suppressed or non-suppressed conditions; * $p < 0.05$, ** $p < 0.01$.

in skill learning, using a visual-spatial discrimination task performed by mice. The impact of suppression of excitatory spiking activity in each region was analyzed during three training stages distinguished by varying levels of accuracy, and a fourth, over-trained stage added here to test the effects of suppression in the absence of correction trials. Mice showed improvements in speed and accuracy over days that were qualitatively similar to changes seen with human skill learning. With suppression, no changes were seen during stage 1. However, effects seen during stages 2 and 3 resembled a reversion to stage 1, suggesting that the absence of a stage 1 effect reflects a behavioral floor in performance. Suppression of V1 reduced accuracy and enhanced bias, effects that increased across stages. This increase is likely attributable to the overall improvement in performance, with the result that mice had further to drop back toward the behavioral floor. Suppression of ACC produced effects on accuracy and bias that were qualitatively similar to those seen in V1, but of lesser magnitude. Effects of suppression on accuracy in dHC mice were modest in comparison to those seen in V1 and ACC mice, and appeared to be largely restricted to stage 3. An overall effect on RT was only observed with suppression in V1. This effect was modest, and most evident in the late stage of training (Fig. 3). However, RTs were consistently longer following errors, and this effect was enhanced with suppression in all three regions during different stages of training.

Discussion

The goal of the present study was to compare serial versus parallel involvement of three brain regions, the V1, ACC, and dHC, in a mouse model of skill learning. Suppression of excitatory activity in V1 and ACC significantly decreased accuracy and increased response bias across stages of training. Similar, though less robust, effects were observed with dHC suppression. While overall effects on RT were only seen with V1 suppression, RTs were lengthened with suppression in all three regions following errors at different stages of training. Taken together, these data are generally consistent with a model of parallel involvement of these three regions in our visual-spatial discrimination mouse model of skill learning.

Cortical areas

We examined the effects of V1 suppression as a control to determine whether cortical processing was necessary for the visual-spatial discrimination. We assumed that suppression of the visual system would influence performance throughout the task. The ACC and dHC, both of which have been found to exhibit changes in activation with skill learning and performance (Raichle, 1998; Procyk et al., 2000; Poldrack and Gabrieli, 2001; Karim et al., 2017), were also hypothesized to play a role based on specific task features. Performance of the visual-spatial discrimination required mice to transform vertically-oriented (top-bottom) cues to horizontal (left-right) movements. An action based on this type of transformation is more cognitively demanding than one involving a cue directing movement in the same plane and direction (e.g., a leftward cue directing a leftward movement). In addition

to transforming cue orientation to plane of movement, mice also had to inhibit any pre-existing natural bias which could conflict with the appropriate movement. Because the ACC is involved in higher-level attention (Petersen and Posner, 2012) and conflict resolution (Botvinick et al., 2001) we hypothesized that it would be involved in the early stage of the task as mice sought the optimal strategy for reward. Our task also required forming numerous different associations. Because the dHC is critically involved in associative learning (Brasted et al., 2003; Suzuki, 2007; Squire and Wixted, 2011), we hypothesized that it would be involved during early and middle stages as the associations among stimulus location, movement, and reward delivery were developed and strengthened. While our task is conceptually simple, mice clearly found forming and acting on the associations specific to the optimal strategy challenging as evidenced by the thousands of trials required to reach criterion performance. This lends further credence to a proposed role for the dHC in our task as forming even the simplest associations may require the hippocampus if stimulus parameters are changed to increase task difficulty (Beylin et al., 2001).

Serial and parallel views

Both serial and parallel models of performance of individual trials during skill learning have been proposed. Sternberg (1969) provided a serial model and developed the additive factors method for determining which independent variables influenced any of the successive stages needed to carry out the task. Anderson and colleagues (Anderson, 2005; Tenison et al., 2016) have developed serial models using ACT-R to simulate complex problem solving tasks. Parallel models of individual trials often use a race horse analogy in which responses compete and the first to reach a threshold determines the response (Logan, 2002; Ratcliff et al., 2016). Because stored items accumulated over time, later trials were faster than earlier ones. Logan's model, however, argued for a continuous process, not for stages that overlapped in time.

Most cognitive models of the stages of learning have used a serial viewpoint to describe the stages involved. Fitts and Posner proposed that skill learning progressed through three stages (Fitts and Posner, 1967; Tenison et al., 2016). The first stage was cognitive and involved evaluation of conflicting strategies in accord with past experience. The second stage was associative and involved connecting the strategies into an overall solution. The third stage was autonomous and involved practice to achieve a highly efficient program to execute the skill. The likely reason that such a serial process has been widely used is that early and late parts of the learning curve clearly do not overlap.

Data from the present study are consistent with a parallel view of information processing across brain areas during skill learning. Nevertheless, behavioral changes in the present study were reminiscent of those one might expect based on the serial three-stage model proposed by Fitts and Posner. For example, different strategies are available to mice to obtain water rewards, and the effec-

tiveness of some of these strategies depends on external factors including reward size and the percentage of correction trials. Before stage 1, mice are rewarded simply for standing still on the ball. At the start of stage 1, this shaping step is removed, and mice must first request a trial (stand still) and then execute a movement. Only when the movement is in the correct direction will a reward be delivered. At this point, mice typically engage a bias strategy, as seen during stage 1 in [Figures 3–5](#). This phenomenon is both consistent and robust, with 76% of mice from the present study developing a >50% bias during stage 1, and 64% exhibiting biases between 75% and 100%. The experimenter gradually increases the % correction, such that the bias strategy become increasingly inefficient for obtaining water rewards. It is, however, still possible to obtain sufficient water simply by executing more trials. In extreme cases, highly biased mice may engage this second strategy and execute 1000+ trials in a single session. We combat this by decreasing the reward size. At some point, the strategy shifts to one of correctly executing movements based on cue position. Bias, trial count, and % correction trials all are decreased as accuracy increases past 65% in a 50-trial block. Therefore, stage 1 in the present study includes a period during which there are several conflicting strategies available for obtaining the water reward, and the selection of the optimal strategy typically only follows rejection of less efficient ones based on previous experience. Stage 2 includes the period when we first see the specific associations of the optimal strategy executed effectively to achieve the reward. Finally, stage 3 includes the highly efficient and reliable performance of the optimal strategy.

The cognitive stage 1 of Fitts and Posner is generally considered a period when task goals are recognized and that information is used to develop an appropriate sequence of actions for achieving those goals, a process seen as involving conscious evaluation of explicit knowledge. Whether a mouse can be said to engage in such a process cannot be directly ascertained from the current data. However, this question is secondary to our demonstration that suppression of V1, ACC, and dHC impact behavior throughout training in this mouse model of skill learning.

Evidence in the present study of involvement of different regions across different stages of learning is not unexpected. Using multivoxel neuroimaging techniques [Tenison et al. \(2016\)](#) showed it was possible to examine the extent to which several brain areas were active during stages of a problem-solving task. While prefrontal areas were most active during the early stage, and motor areas in the late stage, there was overlap between the neural areas and the stages. This is consistent with other neuroimaging ([Formisano and Goebel, 2003](#)) and animal studies ([Redondo and Morris, 2011](#)) that have shown that brain areas overlap with respect to their involvement in the learning process, supporting a parallel view of information processing. But how might we interpret the effects observed in the present study, especially those associated with suppression of the ACC and dHC?

Effects of ACC suppression mirrored the time course of V1 suppression, decreasing accuracy throughout training. While our data are consistent with a parallel model, it is possible that the ACC performs different functions at different stages of training which may, or may not, directly relate to skill learning. The ACC is clearly involved in attention and top-down control, processes that would be critical early in training in the identification and adoption of an optimal strategy for task performance ([Botvinick et al., 2001](#); [Petersen and Posner, 2012](#)). At the other end of the spectrum, the ACC has been implicated in recall of remote memories ([Bontempi et al., 1999](#)), with the time course for learning in the present study easily accommodating that required for the structural changes seen in the ACC accompanying memory consolidation ([Vetere et al., 2011](#)). Potentially representing a bridge between these two roles, [Procyk et al. \(2000\)](#) described trial-to-trial transitions in the primate ACC between search-related activity involving increased attentional load and repetition-related activity based on memory of a motor program when attention to action was reduced. While their study does not speak to skill learning per se (all motor behaviors and task solutions had been learned before recording), it demonstrates responses of different neurons in the same region associated with either attention or recall. It is possible that the route and/or direction of information flow also change with stage of learning. The ACC connects with the hippocampus through both the thalamus ([Xu and Sudhof, 2013](#)) and entorhinal cortex ([Anderson et al., 2016](#)). Information also travels in both directions between the ACC and dHC. Determining whether and how these observations relate to ACC involvement throughout training, could be addressed by future studies.

While dHC suppression did modestly reduce accuracy, it did not affect the overall rate of acquisition. The relatively small effect of hippocampal suppression may support the idea that learning in this task involves the striatum rather than the hippocampus ([Makino et al., 2016](#)). Suppression may also have been more effective were it to have been performed on a higher percentage of trials. It is also possible that the root of the modest effects seen here are methodological in nature. Based on the literature, PV-IN stimulation should be an effective means to suppress principal neuron activity in the hippocampus. Hippocampal PV-INs have extensive axonal arbors yielding an enormous number of synaptic contacts ([Sik et al., 1995](#); [Norenberg et al., 2010](#)). Stimulation of PV-INs expressing ChR2 effectively suppresses neighboring CA1 pyramidal neuron activity ([Ognjanovski et al., 2017](#)). Thus, we are confident that this approach accomplished the desired physiologic effect in the hippocampus within the illuminated region. We cannot, however, fully discount the complex geometry of the dHC, and the possibility that targeting multiple sites along the septal-temporal axis may have yielded more robust effects. Regardless, our dHC data are in line with a parallel model where the computations in each brain area occurred throughout learning rather than being specific to any stage.

A major unanticipated result of this study was that, while suppression decreased accuracy, it had relatively

little overall effect on RT. This was a surprising result since increasing task difficulty and distracting attention have been shown to increase RT and error in both rodents and humans (Beane and Marrocco, 2004; for review, see Posner and Niell, 2019). Manipulating neural activity through optogenetics is, of course, not something done in humans, thus complicating the comparison. However, less invasive efforts to interfere with network performance by transcranial magnetic stimulation (Capotosto et al., 2009) or electrical stimulation (Zmigrod et al., 2016) show increases in RT during performance in humans.

It should not be concluded that RT is an insensitive measure in the mice. As shown in Figure 2, RT declines regularly with practice in mice much as in humans. Moreover, mice show longer RTs following error than following a correct response. This effect was even greater when the error occurred during a suppressed trial in a stage-dependent and region-dependent manner (Fig. 8B,D). In humans, slowing following an error even when there is no explicit error feedback is often taken as evidence that the person was aware of the error (Wessel, 2012). However, our experimental design differed from many human studies in a critical way that may explain the effects we observed: errors triggered a distinct visual cue (a full-field bright flash), which may have made the error salient to the mice. However, as we do not have physiologic data necessary to determine whether the error trials were accompanied by the error related negativity found in humans, directly testing this hypothesis is beyond the scope of the present study (Gehring et al., 1995, 2018).

In the present study, the goal was to “suppress” activity in different regions of the brain to examine how each was involved in skill learning and performance at different stages of proficiency. The direct effect of the laser was confined to the specific area (e.g., V1). However, an important caveat in the interpretation of optogenetic effects is that the impact on overall brain activity is not limited to the region of direct inactivation. Unlike other lesion methods, the optogenetic suppression is nearly immediate, giving no time for other areas of the brain to adjust to the loss of the suppressed areas. For example, the suppression of activity in V1 is propagated upstream and downstream to other brain areas, resulting in perturbation of these areas that may be greater than just the loss of visual input. Thus, the temporal specificity of action and maintenance of structural integrity seen as major advantages of optogenetics over conventional techniques may in fact expose adjacent elements of the circuit to unintended disruption (Hong et al., 2018). One potential alternative interpretation of the effect seen with ACC suppression on performance involves the direct, reciprocal projections linking the ACC and V1 (Zhang et al., 2014). Effects of ACC and V1 suppression in the present study were qualitatively similar. It is possible that ACC suppression impacted behavior indirectly through the V1 projection, resulting in the similar, albeit smaller, impairment in accuracy and increased bias. It is worth noting, however, that RT did not increase with suppression of the ACC as it did with suppression of V1, suggesting that the observed similarities may be correlative and not causative. Assum-

ing the effects on performance of ACC or V1 suppression were not due to their interconnections, we believe the most likely effect of suppression of V1 is to impair the visual coding of the stimulus. Mice have “requested” a trial, but do not see a cue, and thus resort to guessing, an interpretation that is consistent with the chance level of performance reported. In contrast, the smaller performance decrement seen with ACC suppression suggests that the cue is visible, but that suppression causes a lapse of attention (Petersen and Posner, 2012) or decrease in salience (Seeley et al., 2007; Menon, 2015; Chen et al., 2016; Peters et al., 2016) and thus a reversion to responding in the previously preferred (biased) direction.

Two of 10 mice with fibers targeting the dHC failed to learn within 50 d of training, whereas all mice from ACC, V1, and control groups successfully acquired the task. It seems unlikely that these two mice failed to acquire the task because of the optogenetic suppression, as the effects observed in the remaining eight were the smallest among the three groups undergoing suppression during training, and mean trials to criterion was the lowest for all four groups. It also is unlikely that the failure was attributable to the greater invasiveness of the surgical procedure (fibers targeting the dHC were implanted deeper than those targeting the ACC or V1). In addition to the eight mice that learned from this group, an additional five control mice with dHC fibers successfully acquired the task. We believe, therefore, that the most parsimonious explanation for the failure of these two mice to acquire the task was some reason other than the optogenetic manipulation or the surgery itself, and that it is simply chance that both mice were from one group.

Overall, our evidence supports that our mouse skill learning task is greatly influenced by V1 and ACC throughout learning. It shows a more modest influence from the suppression of the dHC. Notably, inactivating the ACC influences errors without a change in RT. Since this last effect is different from would be expected based on data from humans it requires further studies to explore its significance.

References

- Anderson JR (2005) Human symbol manipulation within an integrated cognitive architecture. *Cogn Sci* 29:313–341.
- Anderson JR, Fincham JM, Douglass S (1999) Practice and retention: a unifying analysis. *J Exp Psychol Learn Mem Cogn* 25:1120–1136.
- Anderson MC, Bunce JG, Barbas H (2016) Prefrontal-hippocampal pathways underlying inhibitory control over memory. *Neurobiol Learn Mem* 134:145–161.
- Atallah BV, Bruns W, Carandini M, Scanziani M (2012) Parvalbumin-expressing interneurons linearly transform cortical responses to visual stimuli. *Neuron* 73:159–170.
- Beane M, Marrocco R (2004) Cholinergic and noradrenergic inputs to the posterior parietal cortex modulate the components of exogenous attention. New York: Guilford Press.
- Beylin AV, Gandhi CC, Wood GE, Talk AC, Matzel LD, Shors TJ (2001) The role of the hippocampus in trace conditioning: temporal discontinuity or task difficulty? *Neurobiol Learn Mem* 76:447–461.
- Bontempi B, Laurent-Demir C, Destrade C, Jaffard R (1999) Time-dependent reorganization of brain circuitry underlying long-term memory storage. *Nature* 400:671–675.

- Botvinick MM, Braver TS, Barch DM, Carter CS, Cohen JD (2001) Conflict monitoring and cognitive control. *Psychol Rev* 108:624–652.
- Boyden ES, Zhang F, Bamberg E, Nagel G, Deisseroth K (2005) Millisecond-timescale, genetically targeted optical control of neural activity. *Nat Neurosci* 8:1263–1268.
- Brasted PJ, Bussey TJ, Murray EA, Wise SP (2003) Role of the hippocampal system in associative learning beyond the spatial domain. *Brain* 126:1202–1223.
- Capotosto P, Babiloni C, Romani GL, Corbetta M (2009) Frontoparietal cortex controls spatial attention through modulation of anticipatory alpha rhythms. *J Neurosci* 29:5863–5872.
- Chen T, Cai W, Ryali S, Supekar K, Menon V (2016) Distinct global brain dynamics and spatiotemporal organization of the salience network. *PLoS Biol* 14:e1002469.
- Dombeck DA, Khabbazi AN, Collman F, Adelman TL, Tank DW (2007) Imaging large-scale neural activity with cellular resolution in awake, mobile mice. *Neuron* 56:43–57.
- Fitts PM, Posner MI (1967) Human performance. Belmont, CA: Brooks/Cole.
- Formisano E, Goebel R (2003) Tracking cognitive processes with functional MRI mental chronometry. *Curr Opin Neurobiol* 13:174–181.
- Frey U, Morris RG (1997) Synaptic tagging and long-term potentiation. *Nature* 385:533–536.
- Frey U, Morris RG (1998) Synaptic tagging: implications for late maintenance of hippocampal long-term potentiation. *Trends Neurosci* 21:181–188.
- Gehring WJ, Coles MG, Meyer DE, Donchin E (1995) A brain potential manifestation of error-related processing. *Electroencephalogr Clin Neurophysiol Suppl* 44:261–272.
- Gehring WJ, Goss B, Coles MGH, Meyer DE, Donchin E (2018) The error-related negativity. *Perspect Psychol Sci* 13:200–204.
- Hilgard ER, Campbell AA (1937) Vincent curves of conditioning. *J Exp Psychol* 21:310–319.
- Hong YK, Lacefield CO, Rodgers CC, Bruno RM (2018) Sensation, movement and learning in the absence of barrel cortex. *Nature* 561:542–546.
- Karim HT, Huppert TJ, Erickson KI, Wollam ME, Sparto PJ, Sejdić E, VanSwearingen JM (2017) Motor sequence learning-induced neural efficiency in functional brain connectivity. *Behav Brain Res* 319:87–95.
- Kitchen CM (2009) Nonparametric vs parametric tests of location in biomedical research. *Am J Ophthalmol* 147:571–572.
- Laming D (1979) Choice reaction performance following an error. *Acta Psychol (Amst)* 43:199–224.
- Lien AD, Scanziani M (2011) In vivo labeling of constellations of functionally identified neurons for targeted in vitro recordings. *Front Neural Circuits* 5:16.
- Logan GD (2002) An instance theory of attention and memory. *Psychol Rev* 109:376–400.
- Makino H, Hwang EJ, Hedrick NG, Komiyama T (2016) Circuit mechanisms of sensorimotor learning. *Neuron* 92:705–721.
- Menon V (2015) Salience network. In: *Brainmapping. An encyclopedic reference, Vol 2* (Toga AW, ed), pp 597–611. San Diego: Elsevier.
- Newell A, Rosenbloom PS (1981) Mechanisms of skill acquisition and the law of practice. Hillsdale, NJ: Erlbaum.
- Niell CM, Stryker MP (2010) Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* 65:472–479.
- Norenberg A, Hu H, Vida I, Bartos M, Jonas P (2010) Distinct nonuniform cable properties optimize rapid and efficient activation of fast-spiking GABAergic interneurons. *Proc Natl Acad Sci USA* 107:894–899.
- Ognjanovski N, Schaeffer S, Wu J, Mofakham S, Maruyama D, Zochowski M, Aton SJ (2017) Parvalbumin-expressing interneurons coordinate hippocampal network dynamics required for memory consolidation. *Nat Commun* 8:15039.
- Peters SK, Dunlop K, Downar J (2016) Cortico-striatal-thalamic loop circuits of the salience network: a central pathway in psychiatric disease and treatment. *Front Syst Neurosci* 10:104.
- Petersen SE, Posner MI (2012) The attention system of the human brain: 20 years after. *Annu Rev Neurosci* 35:73–89.
- Poldrack RA, Gabrieli JD (2001) Characterizing the neural mechanisms of skill learning and repetition priming: evidence from mirror reading. *Brain* 124:67–82.
- Posner MI, Niell CM (2019) Illuminating the neural circuits underlying orienting of attention. *Vision* 3.
- Procyk E, Tanaka YL, Joseph JP (2000) Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nat Neurosci* 3:502–508.
- Rabbitt PM (1966) Errors and error correction in choice-response tasks. *J Exp Psychol* 71:264–272.
- Raichle ME (1998) The neural correlates of consciousness: an analysis of cognitive skill learning. *Philos Trans R Soc Lond B Biol Sci* 353:1889–1901.
- Ratcliff R, Smith PL, Brown SD, McKoon G (2016) Diffusion decision model: current issues and history. *Trends Cogn Sci* 20:260–281.
- Redondo RL, Morris RG (2011) Making memories last: the synaptic tagging and capture hypothesis. *Nat Rev Neurosci* 12:17–30.
- Seeley WW, Menon V, Schatzberg AF, Keller J, Glover GH, Kenna H, Reiss AL, Greicius MD (2007) Dissociable intrinsic connectivity networks for salience processing and executive control. *J Neurosci* 27:2349–2356.
- Sik A, Penttonen M, Ylinen A, Buzsáki G (1995) Hippocampal CA1 interneurons: an in vivo intracellular labeling study. *J Neurosci* 15:6651–6665.
- Squire LR, Wixted JT (2011) The cognitive neuroscience of human memory since H.M. *Annu Rev Neurosci* 34:259–288.
- Sternberg S (1969) The discovery of processing stages. *Acta Psychol (Amst)* 30:276–315.
- Suzuki WA (2007) Making new memories: the role of the hippocampus in new associative learning. *Ann NY Acad Sci* 1097:1–11.
- Tenison C, Anderson JR (2016) Modeling the distinct phases of skill acquisition. *J Exp Psychol Learn Mem Cogn* 42:749–767.
- Tenison C, Fincham JM, Anderson JR (2016) Phases of learning: how skill acquisition impacts cognitive processing. *Cogn Psychol* 87:1–28.
- Vetere G, Restivo L, Cole CJ, Ross PJ, Ammassari-Teule M, Josselyn SA, Frankland PW (2011) Spine growth in the anterior cingulate cortex is necessary for the consolidation of contextual fear memory. *Proc Natl Acad Sci USA* 108:8456–8460.
- Weible AP, Moore AK, Liu C, DeBlander L, Wu H, Kentros C, Wehr M (2014) Perceptual gap detection is mediated by gap termination responses in auditory cortex. *Curr Biol* 24:1447–1455.
- Wessel JR (2012) Error awareness and the error-related negativity: evaluating the first decade of evidence. *Front Hum Neurosci* 6:88.
- Xu W, Sudhof TC (2013) A neural circuit for memory specificity and generalization. *Science* 339:1290–1295.
- Zhang S, Xu M, Kamigaki T, Hoang Do JP, Chang WC, Jenvay S, Miyamichi K, Luo L, Dan Y (2014) Selective attention. Long-range and local circuits for top-down modulation of visual cortex processing. *Science* 345:660–665.
- Zmigrod S, Zmigrod L, Hommel B (2016) Transcranial direct current stimulation (tDCS) over the right dorsolateral prefrontal cortex affects stimulus conflict but not response conflict. *Neuroscience* 322:320–325.