

Learning alters neural activity to simultaneously support memory and action

Darby M. Losey^{1,2,3}, Jay A. Hennig^{1,2,3,+}, Emily R. Oby^{2,7,+}, Matthew D. Golub^{2,4,5}, Patrick T. Sadtler^{2,7}, Kristin M. Quick^{2,7}, Stephen I. Ryu^{5,8}, Elizabeth C. Tyler-Kabara^{2,9,10,11}, Aaron P. Batista^{2,7,*}, Byron M. Yu^{1,2,4,6,*}, Steven M. Chase^{1,2,6,*}

1 Neuroscience Institute, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA

2 Center for the Neural Basis of Cognition, Pittsburgh, Pennsylvania 15213, USA

3 Machine Learning Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA

4 Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA

5 Department of Electrical Engineering, Stanford University, Stanford, California 94305, USA

6 Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA

7 Department of Bioengineering, University of Pittsburgh, Pittsburgh, Pennsylvania 15213, USA

8 Department of Neurosurgery, Palo Alto Medical Foundation, Palo Alto, California 94301, USA

9 Department of Physical Medicine and Rehabilitation, University of Pittsburgh, Pittsburgh, Pennsylvania 15213, USA

10 Department of Neurological Surgery, University of Pittsburgh, Pittsburgh, Pennsylvania 15213, USA

11 Department of Neurosurgery, Dell Medical School, University of Texas at Austin, Austin, Texas 78712, USA

+ and * denote equal contributions, see author contributions table

Abstract

1 How are we able to learn new behaviors without disrupting previously learned ones?
2 To understand how the brain achieves this, we used a brain-computer interface (BCI)
3 learning paradigm, which enables us to detect the presence of a memory of one behav-
4 ior while performing another. We found that learning to use a new BCI map altered
5 the neural activity that monkeys produced when they returned to using a familiar
6 BCI map, in a way that was specific to the learning experience. That is, learning left
7 a “memory trace.” This memory trace co-existed with proficient performance under
8 the familiar map, primarily by altering dimensions of neural activity that did not
9 impact behavior. Such a memory trace could provide the neural underpinning for
10 the joint learning of multiple motor behaviors without interference.

11 Introduction

12 Suppose an experienced skier learns to snowboard. Skiing and snowboarding require
13 different sets of muscle activations, driven by different neural population activity
14 patterns, to achieve the same goal of getting down the mountain without falling. How
15 does the brain incorporate the knowledge about how to snowboard along with the
16 knowledge about how to ski? The first possibility is that after learning to snowboard
17 the neural activity used for skiing remains unchanged. In this scenario, the new
18 neural activity for snowboarding would only be recalled when snowboarding again.
19 Such context-dependent recall has been observed in certain learning settings, such
20 as the remapping of hippocampal place-fields between environments (Alme et al.,
21 2014), and has been proposed as a potential mechanism for motor memory storage
22 (Herzfeld et al., 2014).

23 The second possibility is that the neural activity used for skiing is altered by
24 the recently acquired ability to snowboard. Several studies have suggested that this
25 might be the case, as neural tuning has been observed to change as a result of motor
26 adaptation (Li et al., 2001, Arce et al., 2010, Cherian et al., 2013, Perich and Miller,
27 2017, Sun et al., 2022). There are two possible reasons for these neural activity
28 changes. First, these changes might constitute a memory of the learning experience.
29 That is, learning could lead to a “memory trace”, which we define as an alteration
30 of the population activity patterns used to perform familiar tasks in a manner that
31 renders them also appropriate for a newly learned task. Second, these changes could
32 be attributed to the many task-agnostic factors, such as changes in arousal (Cowley
33 et al., 2020, Hennig et al., 2021a), motivation (Roesch and Olson, 2004), posture
34 (Graziano, 2006), or altered arm dynamics (Cherian et al., 2013, Perich and Miller,
35 2017). Without a known causal link between neural activity and behavior, it is
36 difficult to determine if and how changes in neural activity after learning might
37 constitute a memory trace.

38 Here we overcome this difficulty by leveraging a brain-computer interface (BCI)
39 paradigm (Jarosiewicz et al., 2008, Ganguly and Carmena, 2009, Koralek et al., 2012,
40 Hwang et al., 2013, Sadtler et al., 2014, Gulati et al., 2017, Jeon et al., 2022, Oby
41 et al., 2019). A key advantage of a BCI for the study of motor memory is that the
42 relationship between neural activity and behavior (termed the BCI map) is specified
43 by the experimenter (Golub et al., 2016, Orsborn and Pesaran, 2017). This feature
44 of a BCI is crucial in enabling us to look for a memory trace because it allows us to
45 evaluate how changes in neural activity relate to a task that is not being performed.

46 We trained three monkeys to perform a BCI task. We used two different BCI maps
47 in each experimental session. Much like the example of an experienced skier learning
48 to snowboard, a monkey first controlled a computer cursor using a familiar *Map A*,
49 and then learned how to use a new *Map B*. Following learning, we reinstated *Map*
50 *A*. This allowed us to evaluate whether monkeys used different population activity
51 patterns to control *Map A* before and after learning *Map B*. Furthermore, to see if
52 neural activity showed a memory trace of having learned *Map B*, we evaluated how
53 well the neural activity produced by the monkey during the use of *Map A* would have

54 controlled the cursor through the offline Map B, comparing pre- versus post-learning.
55 We observed that, after learning Map B, the monkeys were subsequently able to
56 control the cursor using Map A, and yet the neural activity remained consistent
57 with improved performance using Map B. That is, learning left a memory trace by
58 altering the neural activity used to perform the familiar task, such that the neural
59 activity became more appropriate for the learned task. The memory trace coexisted
60 alongside proficient Map A performance by altering neural activity primarily along
61 dimensions that did not affect cursor movements under Map A. Overall, our results
62 reveal that learning can leave a memory trace in neural population activity that
63 need not interfere with the subsequent behavior. The formation of a memory trace
64 may thus provide a mechanism to facilitate the learning of multiple motor skills
65 without interference (Krakauer et al., 2005), instantaneous switching between tasks,
66 and rapid relearning of motor behaviors (“savings”).

67 Results

68 Here we study how learning to perform a new task affects the neural activity used
69 while performing a familiar task (Fig. 1a). We trained three monkeys to perform
70 an eight-target center-out task using a brain-computer interface (BCI). The mon-
71 key’s goal on each trial was to guide a computer cursor to an instructed target by
72 modulating his neural activity (Fig. 1b; see Methods). At each moment in time,
73 a BCI map determined the relationship between the neural activity, recorded from
74 ~90 neural units in primary motor cortex (M1), and the cursor’s 2D velocity. Each
75 experiment utilized two different BCI maps, *Map A* and *Map B*, presented across
76 three blocks of trials (Fig. 1c). During the first block (“Task A1”), we provided the
77 monkeys with Map A, which was an “intuitive” map calibrated that day to allow
78 for proficient cursor control without any learning. For the second block (“Task B”),
79 we changed the BCI map to Map B, which the monkey had never used before (see
80 Methods). This resulted in an initial decrement in the monkey’s performance, which
81 improved over the course of several hundred trials as he learned to control the cursor.
82 In the third block (“Task A2”), we reinstated Map A. This typically resulted in the
83 well-known aftereffect that typically follows a bout of motor learning after which
84 performance returns to level comparable to that of Task A1 (Shadmehr and Mussa-
85 Ivaldi, 1994). Data from the Task A1 and Task B periods have been examined in
86 prior work (Sadtlter et al., 2014, Golub et al., 2018, Hennig et al., 2018, 2021a). In
87 this study, we now focus on the neural activity recorded during Task A2, which is the
88 appropriate epoch to address our central question and which we have not reported
89 on previously.

90 Our central question is: how does learning Map B affect the neural activity pro-
91 duced while using Map A? To illustrate the possibilities, we depict two dimensions
92 of neural population activity controlling 1D cursor movements (Fig. 1d). During
93 Task A1, the monkey produces neural activity appropriate for Map A, in that the
94 projection of neural activity onto Map A results in high cursor velocities toward the

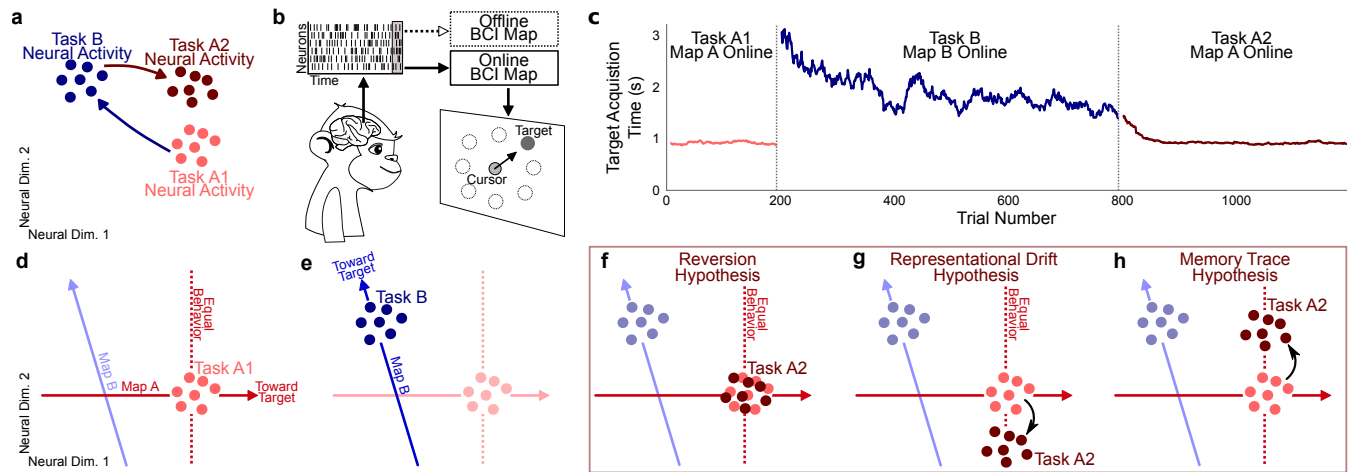


Figure 1. How learning a new task might change the neural activity used for a familiar task

(a) Schematic of how neural activity (colored dots) may change when performing different tasks. Performing Task A for the first time (light red; Task A1), then Task B (blue, Task B), then Task A again (dark red; Task A2) may all yield distinct neural activity patterns. (b) The activity of ~90 neural units, recorded using a Blackrock array implanted in the primary motor cortex (M1), were translated into the movement of the cursor through a brain-computer interface (BCI). A BCI directly relates neural activity to behavior (the horizontal and vertical velocities of a cursor on a computer screen) using a map specified by the experimenters. The online BCI map is the BCI map that dictates cursor movements. The same neural activity can also be interpreted with respect to an offline BCI map that did not determine cursor control movement. (c) Target acquisition times for an example session (N20160714). The initial period where the monkey used Map A to control the cursor is defined to be Task A1 (acquisition times shown in light red). For Task B, the map was switched to Map B. The monkey had to learn how to control the cursor with this new map through trial and error (dark blue). Acquisition times improved, showing that learning occurred. For Task A2, Map A was reinstated (dark red). For visualization, acquisition times were smoothed with a causal 25-trial moving window and are not shown for the first 8 trials of each task. Success rates were near 100% for all three tasks. (d-e) Schematics of how neural activity might look during the three tasks. For illustrative purposes, we show a 2D neural space, which was mapped to a 1D cursor velocity. In the actual experiments, the neural space was ~90D (one dimension per recorded unit), which was mapped to a 2D cursor velocity. (d) During Task A1, neural activity is appropriate for Map A. (e) During Task B, neural activity becomes appropriate for Map B. (f-h) We explore three possibilities for what neural activity might look like during Task A2. (f) Reversion hypothesis: Task A2 neural activity is similar to that used during Task A1. (g) Representational Drift Hypothesis: Task A2 neural activity is different from that used Task A1, but not in a manner that consistently retains high performance through Map B. (h) Memory Trace Hypothesis: Task A2 neural activity contains a memory trace, whereby neural activity is appropriate for both Map A and Map B.

95 target. During Task B, the monkey learns to produce neural activity that is appropriate for Map B (Fig. 1e). Finally, Map A is reinstated during Task A2, and the monkey's neural activity needs to once again become appropriate for Map A.

98 We consider three possibilities for what neural activity might look like after behavior stabilizes during Task A2. One possibility is that, after learning, the population activity patterns produced during Task A2 are similar to those produced during Task A1. We call this the reversion hypothesis (Fig. 1f). Reversion has been observed in various different contexts, such as reaching tasks (Perich and Miller, 2017, Cherian et al., 2013), BCI tasks in visual cortex (Jeon et al., 2022), and in the remapping of hippocampal place fields (Alme et al., 2014). This would indicate that the neural activity we observed in M1 during performance of a task can be unaffected by an intervening learning experience.

107 A second possibility is that neural activity changes in a manner agnostic to the learning experience. We call this the representational drift hypothesis (Fig. 1g; see Druckmann and Chklovskii (2012), Rule et al. (2019), Mau et al. (2020), Deitch et al. (2021), Schoonover et al. (2021)). Representational drift could occur alongside proficient task performance due to many activity patterns corresponding to the same behavioral output (Kaufman et al., 2014, Hennig et al., 2018). This drift could be attributed to any number of uncontrolled factors, such as arousal (Cowley et al., 2020, Hennig et al., 2021a).

115 A third possibility is that changes in neural activity are directly related to the learned task. We consider the possibility that neural activity changes to maintain the memory of the learned task (Task B), while simultaneously supporting accurate cursor movement (i.e., action) during Task A2. We call this the memory trace hypothesis (Fig. 1h). Neural activity changing in this manner could help facilitate the formation of new memories without leading to interference with subsequent behavior. While prior work has observed changes in neural activity as a result of an intervening learning experience and speculated that these changes reflect a memory trace (Li et al., 2001, Arce et al., 2010), with a BCI we know the causal relationship between neural activity and behavior and thus are now able to disambiguate between the representational drift and memory trace hypotheses.

126 We commence our analyses by considering the reversion hypothesis. If the reversion hypothesis were true, we would expect the tuning of individual neural units to remain the same between Tasks A1 and A2. To test this, we fit cosine tuning curves in each of these task periods and measured the change in preferred direction between them. We found many neurons exhibited substantial tuning changes (Fig. 2a). Overall, these tuning curve changes confirm that neural activity produced during Task A2 is distinct from that of Task A1 at the single unit level (Fig. 2b). A lack of support of the reversion hypothesis is also evident when we consider the population of neurons together (Fig. 2c). We observed that, for many targets, neural activity during Tasks A1 and A2 occupied different regions within the neural population space (Fig. 2d), in contradiction to the schematic in Fig. 1f. Thus, our data are not consistent with the reversion hypothesis.

138 Although we can rule out the reversion hypothesis, our analyses to this point




Figure 2. Learning a new task changes the neural representation of a familiar task.

(a) Tuning curves relating cursor-to-target direction to the firing rate for an example neural unit. A cosine tuning curve was fit separately for each of the three task periods. This unit (unit 37 from session L20131205) changes its tuning (measured by a change in preferred direction, ΔPD) between Tasks A1 and Tasks A2. Shading indicates a 95% confidence interval. (b) Many units show a change in tuning between Tasks A1 and A2 ($P < 10^{-10}$, two-sided paired Wilcoxon signed-rank test, $n=3461$ neural units). Black shows the absolute change in preferred direction for units across all sessions. Grey indicates the prediction of the reversion hypothesis (that is, no change in PD other than that due to sampling error). This was estimated using a shuffle control in which labels for Task A1 and A2 were randomly permuted across trials (see Methods). (c) A view of the population neural activity for one example target (J20120601; target 270°) across all three task periods. We applied linear discriminant analysis (LDA) to find the plane which best separates the neural activity from the three task periods. Activity is projected onto that plane, with mean and covariances across timesteps shown. (d) Population activity is different between Task A1 and Task A2 ($P < 10^{-10}$, two-sided paired Wilcoxon signed-rank test, $n=172$ targets). Black shows the distance between the Task A1 and Task A2 means in the 10D population activity space. Grey indicates the prediction of the reversion hypothesis, obtained using a shuffle control (see Methods).




Figure 3. Learning leaves a memory trace. (a) During Task A1 and Task A2, neural activity drives the cursor through Map A (red trajectories, with dots denoting cursor positions at each timestep). The same neural activity can also be projected through Map B in an offline analysis (blue arrows). During Task A2, the Map B velocities more directed toward the target than during Task A1. Both trials come from the 225° target from session N20160329. For visualization purposes, the data are rotated and the velocities are scaled. (b) Task performance is similar between Task A1 and Task A2 (see [Extended Data Fig. 1](#)). For this target, there is no significant difference in progress (i.e., the component of velocity that points toward the target), through Map A ($P = 0.80$, two-sided unpaired Wilcoxon rank-sum test). Dots on the horizontal axis denote the average progress for the trials shown in (a). Triangles above the histograms denote the mean of each distribution. (c) Velocities through the offline Map B. The difference in average progress defines the memory trace for that target. For this target, there is significantly higher progress through Map B during Task A2, relative to Task A1 ($P = 0.0077$, two-sided unpaired Wilcoxon rank-sum test) yielding a memory trace of 14.49 mm/s. Same conventions as in (b). (d) All three monkeys showed a memory trace for well-learned targets (Monkey J, $P < 10 \times 10^{-10}$, two-sided paired Wilcoxon signed-rank test, $n=88$ targets; Monkey N, $P = 1.14 \times 10^{-7}$, $n=48$ targets; Monkey L $P = 0.0020$, $n=36$ targets). For a small fraction of targets, the measured memory trace is negative. This arises when progress through Map B is worse during Task A2 than Task A1. When also including unlearned targets, a memory trace is still evident for Monkey’s J and N, but not Monkey L (Monkey J, $P = 1.57 \times 10^{-4}$, two-sided paired Wilcoxon signed-rank test, $n=176$; Monkey N, $P = 1.99 \times 10^{-6}$, $n=96$; Monkey L $P = 0.61$, $n=72$; see [Extended Data Fig. 3](#)). Monkey L showed a smaller memory trace than Monkeys J and N, likely due to less learning occurring ([Extended Data Fig. 4](#)). Triangles denote the average memory trace for each monkey. The white tick mark on the horizontal axis of the middle histogram denotes the example target illustrated in (b) and (c).

139 does not distinguish the memory trace hypothesis from the representational drift
140 hypothesis. To do so, we must evaluate how the observed changes in neural activity
141 relate to the previously-learned behavior. Our BCI approach makes this possible
142 because we can quantify whether the neural activity is suitable for a BCI map that
143 is not currently being used by the monkey. To illustrate this process, we compare
144 neural activity from a single trial during each of Task A1 and Task A2 corresponding
145 to the same target (Fig. 3a top). For each population activity pattern, we can
146 evaluate its “progress” through Map A as the extent to which it moves the cursor
147 toward the target (see Methods). During both Tasks A1 and A2, Map A determines
148 cursor velocity, and the monkeys showed proficient control of the cursor during both
149 tasks (Fig. 3b; Extended Data Fig. 1).

150 Since we are using a BCI, progress can also be calculated for Map B, even when
151 the animal is using Map A to control the cursor. Progress under Map B measures the
152 extent to which a given neural activity pattern *would have moved* the cursor toward
153 the target, had Map B been instantiated. During Task A1, the monkeys exhibited low
154 progress through Map B, as the velocities through Map B are small and haphazardly
155 oriented relative to the target (Fig. 3a, bottom, Task A1). This is expected because
156 the monkey had not yet experienced Map B, and Map B was selected to be difficult to
157 control using Map A’s neural activity (Sadler et al., 2014). In contrast, during Task
158 A2 the velocities through Map B are larger and more directed toward the target
159 than they were during Task A1 (Fig. 3a, bottom, Task A2), that is, they show
160 higher progress (Fig. 3c). This occurs despite the fact that Map B has no influence
161 on behavior during Task A2 and thus the monkeys have no external incentive while
162 performing Task A2 to maintain high progress through Map B. We define the memory
163 trace as the average increase in the progress toward a given target when projecting
164 the neural activity patterns through Map B during Task A2, relative to Task A1.
165 Across all three monkeys, we found that average progress through Map B was larger
166 during Task A2 than Task A1 (Fig. 3d). This finding supports the memory trace
167 hypothesis (Fig. 1h), but not the representational drift hypothesis, which does not
168 predict this organization.

169 We next assessed the robustness of the memory trace with two tests. First, we
170 showed the memory trace was still present when using a different performance metric,
171 namely, angular error (Extended Data Fig. 2). Second, we quantified the consistency
172 of the effect by showing that the majority of targets from each session exhibited a
173 memory trace, and that the average of the memory traces per session is positive
174 (Extended Data Fig. 3).

175 We next considered whether the memory trace possesses three desirable properties
176 of useful memories. The first property is that a memory should *persist*, meaning that
177 it is present in neural activity without dissipating as time passes. To test this, we
178 examined the later trials of the sessions with the longest Task A2 blocks (Fig. 4a).
179 Specifically, we split sessions into two groups. The first group contained the sessions
180 with at least 300 Task A2 trials, while the second group contained sessions with
181 fewer than 300 Task A2 trials (see Methods). For the group with the longer Task
182 A2 period, we excluded the first 200 trials from analysis in order to quantify the

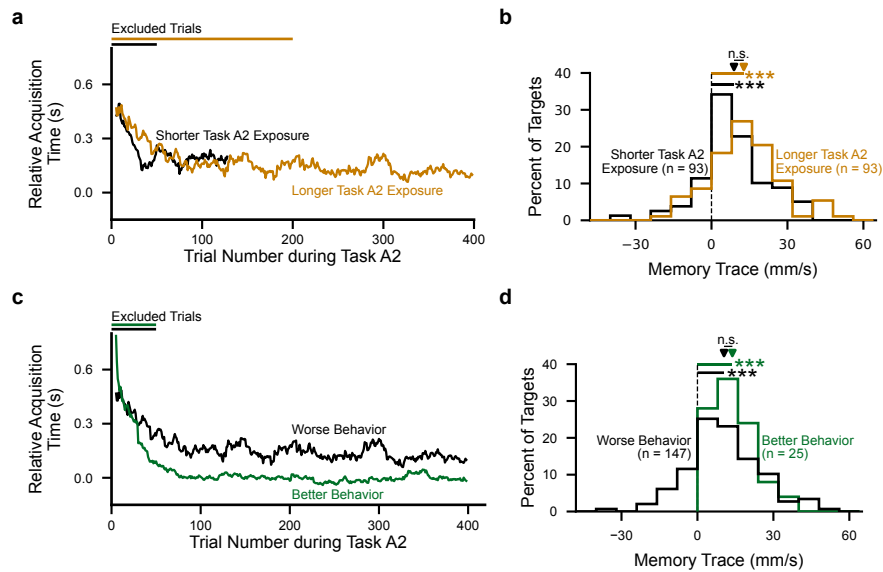


Figure 4. The memory trace persists over time and coexists alongside proficient task performance.

(a) To examine the influence of longer Task A2 exposure on the memory trace, we split sessions into two groups. The first group contains sessions with fewer Task A2 trials (behavioral performance for an example session shown in black), while the second group contains sessions with more Task A2 trials (example session in gold). For the group of shorter sessions, we excluded the first 50 Task A2 trials (black bar). For the longer group, we excluded the first 200 trials (gold bar). Acquisition times are plotted relative to the Task A1 period, where zero represents the an acquisition time equal to the average acquisition time for that target during Task A1. (b) The size of the memory trace was not different for shorter versus longer Task A2 exposure ($P = 0.11$, two-sided unpaired Wilcoxon rank-sum test). The memory trace was still evident for both the longer sessions (gold; $P < 10^{-10}$, two-sided paired Wilcoxon sign-rank test), and the shorter sessions (black; $P = 2.02 \times 10^{-8}$). (c) Behavioral performance during Task A2 from two example sessions, one session with better behavior (faster acquisition time; green) and the other with worse behavior (slower acquisition times; black). Note that the session with worse behavior is the same session as that shown in (a) for the longer Task A2 exposure. (d) To evaluate the influence that behavioral performance during Task A2 has on the size of the memory trace, we split targets into two groups. The first group contained targets where the mean target acquisition time during Task A2 was less than the mean target acquisition time during Task A1 (“better behavior”; green; see [Extended Data Fig. 1](#)), The second group contained targets where the mean target acquisition time during Task A2 was greater than during Task A1 (“worse behavior”; black). The size of the memory trace was not different between the worse behavior and better behavior groups ($P = 0.12$, two-sided unpaired Wilcoxon sign-rank test). Moreover, the memory trace was still evident in the group of targets with better behavior ($P = 5.96 \times 10^{-8}$, two-sided paired Wilcoxon rank-sum test) and worse behavior ($P < 10^{-10}$).

183 memory trace after extended usage of Map A. We found that the memory trace at
184 the end of these longer sessions was not detectably different from the memory trace
185 of the sessions with fewer Task A2 trials (Fig. 4b).

186 The second desirable property of a memory is that it should *coexist* alongside
187 proficient performance of other tasks. To address this, we examined whether the
188 size of the memory trace was contingent on how proficient the behavior was during
189 Task A2 (Fig. 4c). If the instances with worse behavioral performance during Task
190 A2 had the largest memory trace, it could suggest that the memory trace arises due
191 to a trade-off between performance through the two BCI maps. Alternatively, if
192 the memory trace were present even when behavioral performance during Task A2
193 returned to the levels seen during Task A1, it would suggest that the memory trace
194 can coexist without hindering the monkey’s ability to perform the familiar task. We
195 found that the memory trace for the targets with the best behavioral performance
196 during Task A2 showed an average memory trace whose strength was not significantly
197 different from the average memory trace of targets with worse behavioral performance
198 during Task A2 (Fig. 4d). These results indicate the memory trace coexists alongside
199 proficient behavioral performance of the familiar task, and does not represent a
200 compromise between the two learned behaviors.

201 The final property is that more learning should lead to more memory. We found
202 that to be the case, as the size of the memory trace was positively correlated with
203 the amount of learning during Task B (Extended Data Fig. 4). As Monkey L showed
204 less learning than the other two monkeys, this could explain why its memory trace
205 tended to be smaller (Fig. 3d, Monkey L).

206 How can a memory trace coexist without degrading behavioral performance dur-
207 ing Task A2? To understand this, we considered how the changes in neural activity
208 induced by learning Map B relate to Map A. Because we map the activity of ~90 neu-
209 ral units to two-dimensional BCI cursor movements (see Methods), not all changes
210 in neural activity affect cursor movement. We refer to changes in neural activity that
211 affect cursor movement as “output-potent” with respect to that map, and changes
212 that do not as “output-null” (Kaufman et al., 2014). Because Map A and Map B
213 do not share the same output-potent space, it is possible to have neural changes
214 that affect cursor movement through one map without impacting cursor movements
215 through the other.

216 We examined whether the memory trace of Map B (Fig. 5a) resides in the output-
217 potent or output-null space of Map A (Fig. 5b), by decomposing it into its output-
218 potent and output-null components (Fig. 5c). We found that the memory trace
219 resides predominantly in the output-null space of Map A (Fig. 5d), rather than
220 in the output-potent space of Map A (Fig. 5e). This means the memory trace is
221 primarily “stored” in dimensions that do not influence task performance (Extended
222 Data Fig. 5). Since neural activity in dimensions output-null to Map A do not
223 influence cursor velocities during Task A2, this explains how the memory trace can
224 co-exist with proficient behavioral performance.

225 Last, we considered, how does the monkey arrive at the Task A2 solution? There
226 are two possibilities. The first possibility is that there is a partial “unwinding”

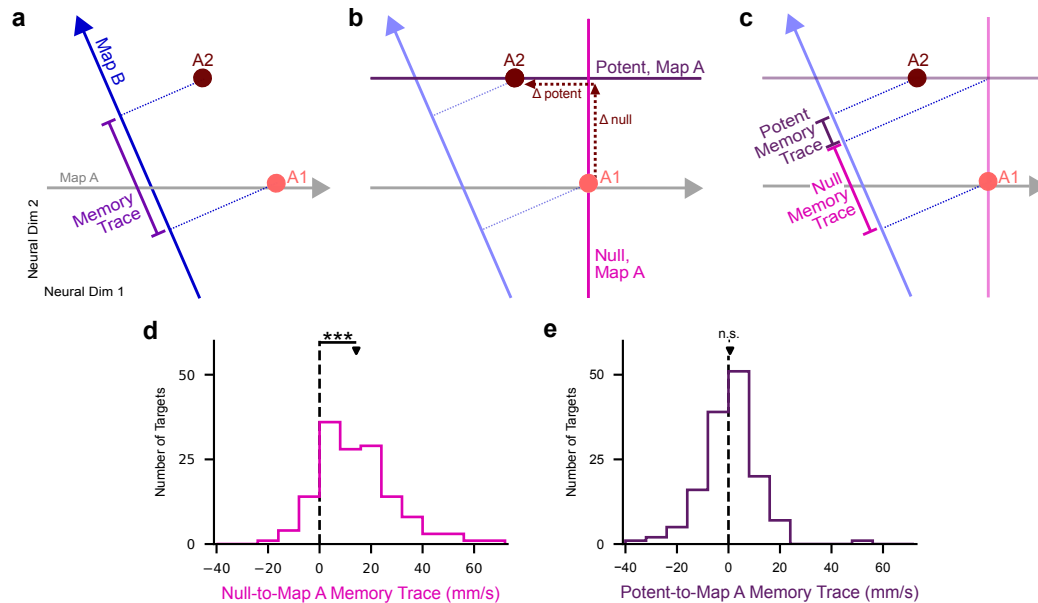


Figure 5. The memory trace is predominantly in the null space of Map A.

(a) Memory trace depicted in same space as Fig. 1d-h. Between Task A1 (light red dot) and Task A2 (dark red dot), neural activity changes. During these time periods, the cursor is controlled using Map A (grey arrow). Task A2 activity is further along Map B (blue arrow) than Task A1 activity, indicating higher progress. The memory trace is defined as difference in the projection onto Map B. (b) The change in neural activity from (a) can be decomposed into a component that is output-potent to Map A (Δ potent) and a component that is output-null to Map A (Δ null). (c) Having decomposed the change in neural activity into output-potent and output-null components, we can correspondingly decompose the memory trace into output-potent and output-null components. (d) Of the targets with a positive memory trace (142 out of 172 targets), the memory trace consistently resides in dimensions null to Map A ($P < 10^{-10}$, two-sided paired Wilcoxon signed-rank test, $n=142$ targets across all monkeys). (e) The contributions from the potent space are not significantly different from zero ($P = 0.31$, two-sided paired Wilcoxon signed-rank test, $n=142$ targets across all monkeys), meaning there is no memory trace on average in the output-potent component of neural activity.

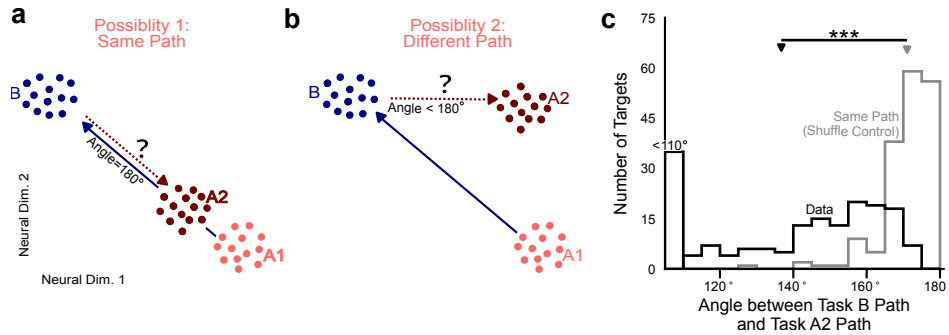


Figure 6. The path of washout does not retrace the path of learning.

We consider two possibilities for how the memory trace arises during Task A2. (a) The first possibility is that the path of washout (i.e., the path neural activity takes from the end of Task B to Task A2) retraces the path of learning (i.e., the path neural activity takes from Task A1 to the end of Task B). This would mean that washout is simply an “unwinding” of the learning experience. (b) The second possibility is that these two paths are distinct, implying that the washout is not simply “unlearning”. (c) To distinguish between these two possibilities, we measured the angle between these two paths. The angle between these paths (black histogram) was smaller than the angles that would be obtained under possibility 1 (grey histogram; see Methods; $P < 10^{-10}$, two-sided paired Wilcoxon signed-rank test, $n=172$ targets across monkeys). This implies the paths of learning and washout are distinct (possibility 2).

227 of the learning that occurred during Task B. This would suggest that the solution
228 used during Task A2 is not novel, and was employed sometime during the learning
229 experience. If this were true, we would expect that the path neural activity takes
230 from the end of Task B to the end of Task A2 (i.e., “the path of washout”, dark red
231 arrow in Fig. 6a) would retrace the path that neural activity takes from the end of
232 Task A1 to the end of Task B (i.e., “the path of learning”, blue arrow in Fig. 6a).
233 The other possibility is that the path of washout is distinct from the path of learning
234 (Fig. 6b). This would imply that the solution the monkey uses during Task A2 is
235 novel, suggesting that the relearning of Map A is distinct from simply “forgetting
236 Map B”. To differentiate between these possibilities, we calculated the angle formed
237 between the path of learning and the path of washout (see Methods). We found that
238 the path of washout is distinct from the path of learning (Fig. 6c).

239 Discussion

240 We studied how the brain can retain a memory of a newly learned task without
241 compromising the performance of familiar tasks. A BCI enables us to make new
242 progress on this longstanding question. This is because using a BCI allows us assess
243 the extent to which the same neural population activity patterns were suitable for
244 multiple different maps (that is, relationships between neural activity and cursor
245 movements), including a map not actively being used. We found that, after a learning
246 experience, neural activity remained appropriate for the learned map even when the
247 animal was using a different (familiar) map. The memory of the learned map was
248 primarily in dimensions in neural space which were output-null to the familiar map.

249 In this way, neural activity simultaneously supported action through the familiar
250 map and still maintained a memory of the recently learned map.

251 It could have been that motor memories were stored in a manner that is not de-
252 tectable when another action is being performed (Herzfeld et al., 2014, Jeon et al.,
253 2022), nor be present in the same neural activity when it is driving behavior. For
254 example, in our experiments, the motor memory could have been stored (perhaps
255 outside of M1) such that the memory is only detectable in M1 the appropriate behav-
256 ior is being performed. Instead, we found that memories can be stored in a manner
257 that makes them evident in M1 even during the execution of other actions.

258 Motor memory consolidation is the process by which memories become more ro-
259 bust to interference (Krakauer et al., 2005). This process takes at least several hours
260 (Shadmehr and Holcomb, 1997), and may require M1 (Muellbacher et al., 2002,
261 Kawai et al., 2015, Rubin et al., 2022). How might the brain bridge from the short-
262 timescale retention of a memory trace that we studied here to the longer-timescale
263 consolidation of a motor memory (Shadmehr and Holcomb, 1997, Gulati et al., 2017)?
264 Our results focused on the short-term inception of a motor memory, within an hour
265 or so of the learned experience. Three possibilities would be consistent with our
266 results. First, a long-term consolidated memory might resemble the memory trace
267 we observed here. Second, it might be that the memory trace we observed is only
268 a short-term phenomenon in M1, dissipating after consolidation. Thus, the memory
269 trace evident in M1 could constitute a short-term storage for the memory before
270 it is offloaded to another brain area during consolidation. Finally, it could be that
271 with further practice with both maps over many days, the neural activity changes to
272 provide even better performance through both BCI maps. In this way, the monkey
273 could effortlessly switch between the two tasks without a drop in performance using
274 the same population activity patterns. That is, with further practice the memory
275 trace could evolve (Nader and Hardt, 2009, Gershman et al., 2017) to lead to even
276 greater coexistence between the two behaviors (Ajemian et al., 2013, Gallego et al.,
277 2018).

278 What is the utility of maintaining a memory trace in neural population activity?
279 A memory trace could enable proficient performance to be reached more quickly
280 upon re-exposure to the learned task. This phenomenon, known as savings, has
281 been frequently observed in motor learning behavior and is often taken as evidence
282 of memory formation (Krakauer et al., 2005, Herzfeld et al., 2014). Our results
283 propose a neural population mechanism for savings. Namely, if Map B were to be
284 re-introduced following performance of Task A2, neural activity would already be
285 situated in population activity space in a manner that would yield better initial
286 performance while using Map B than before learning. While this mechanism can
287 lead to savings due to starting from a better position, our results do not speak to
288 whether there would also be an increased rate of relearning, i.e., a greater reduction
289 in error per trial after the first trial.

290 The memory trace we found in M1 represents one scheme whereby the brain can
291 store multiple memories without interference. We found that the firing of many
292 neurons contribute to the memory trace. This coding scheme marks an interesting

293 contrast to how some memories are formed in the hippocampus, where a sparse subset
294 of neurons encode the memory (Josselyn and Tonegawa, 2020). We observed that the
295 memory trace was primarily due to changes in neural activity orthogonal (i.e., output-
296 null) to the familiar task. Notably, the utilization of different subsets of neurons to
297 encode memories is a special case of orthogonal representations in population activity
298 space (Alme et al., 2014). These lines of evidence together indicate that the brain
299 needs to incorporate new memories into subspaces orthogonal to existing memories
300 in order to avoid interference (Ajemian et al., 2013, Tang et al., 2020, Gava et al.,
301 2021, Libby and Buschman, 2021, Nieh et al., 2021, Xie et al., 2022). Avoiding
302 interference may be harder in the spinal cord, where there are fewer neurons than
303 in cortex or the hippocampus. As fewer neurons likely leads to a more constrained
304 encoding space, a “negotiated equilibrium” between multiple learned behaviors may
305 be required (Wolpaw, 2018).

306 By demonstrating the presence of a memory trace, we ruled out the possibility
307 that changes in neural activity between Task A1 and Task A2 were due solely to
308 representational drift, a change in neural activity manner agnostic to the learning
309 experience. However, representational drift has been observed throughout the brain
310 (Druckmann and Chklovskii, 2012, Rule et al., 2019, Mau et al., 2020, Deitch et al.,
311 2021, Schoonover et al., 2021), and could be occurring alongside the memory trace
312 that we observe. Representational drift differs from the formation of a memory trace
313 in that the changes in neural activity due to representational drift do not directly
314 serve the purpose of memory, but instead are driven by other factors, not under
315 experiment control.

316 Sun et al. (2022) also observed systematic changes in neural activity related to
317 the learning experience. In their study, learning an arm-reaching task in a curl force
318 field induced a uniform shift in preparatory neural activity, which persisted after
319 the force field was removed. The authors conjecture that this shift indexes motor
320 memories (Sheahan et al., 2016). It remains to be seen whether uniform shifts during
321 preparatory activity lead to the reorganization of activity in M1 that constitutes a
322 memory trace, or if these findings support two separate processes.

323 Human and animal learners distinguish themselves from current artificial learning
324 systems in that they can learn to perform a large number of different behaviors. It
325 is a notoriously challenging problem for artificial agents to learn new tasks without
326 overwriting the ability to perform previously learned tasks, an effect termed “catastrophic forgetting”
327 (Masse et al., 2018, Parisi et al., 2019, Yang et al., 2019). Our
328 findings suggest that artificial learning systems could overcome catastrophic forget-
329 ting by implementing some of the same learning principles employed by biological
330 learning systems (Duncker et al., 2020, Hennig et al., 2021b). A sufficiently high
331 dimensional activity space may be important not only in the brain, but also for
332 artificial agents, for learning multiple tasks without interference.

333 Acknowledgments

334 We thank Caleb Kemere for helpful discussions. This work was supported by the Na-
 335 tional Science Foundation Graduate Research Fellowship DGE1745016 and DGE2140739
 336 (to D.M.L.), the Richard King Mellon Presidential Fellowship (to J.A.H.), the Carnegie
 337 Prize Fellowship in Mind and Brain Sciences (to J.A.H.), NIH R01 HD071686 (to
 338 A.P.B., B.M.Y. and S.M.C.), NSF NCS BCS1533672 (to S.M.C., B.M.Y. and A.P.B.),
 339 NSF NCS DRL 2124066 and 2123911 (to B.M.Y., S.M.C. and A.P.B.), NSF CA-
 340 REER award IOS1553252 (to S.M.C.), NIH CRCNS R01 NS105318 (to B.M.Y.
 341 and A.P.B.), NSF NCS BCS1734916 (to B.M.Y.), NIH CRCNS R01 MH118929 (to
 342 B.M.Y.), NIH R01 EB026953 (to B.M.Y.), Simons Foundation 543065 (to B.M.Y.),
 343 NIH R01 NS120579 (to A.P.B.) and NIH R01 HD090125 (to A.P.B.).

344 Author Contributions

	D.M.L.	J.A.H. ⁺	E.R.O. ⁺	M.D.G.	P.T.S.	K.M.Q.	S.I.R.	E.C.T.-K.	A.P.B.*	B.M.Y.*	S.M.C.*
designed analyses	x	x	x						x	x	x
performed analyses	x										
performed animal surgeries			x				x	x			
designed experiments			x	x	x	x			x	x	x
performed experiments			x		x						
wrote/revised manuscript	x	x	x						x	x	x
commented on and approved manuscript	x	x	x	x	x	x	x	x	x	x	x

+ , * denote equal contribution

345 Competing Interests

346 The authors declare no competing interests.

References

- Robert Ajemian, Alessandro D’Ausilio, Helene Moorman, and Emilio Bizzi. A theory for how sensorimotor skills are learned and retained in noisy and nonstationary neural circuits. *Proceedings of the National Academy of Sciences*, 110(52):E5078–E5087, 2013.
- Charlotte B Alme, Chenglin Miao, Karel Jezek, Alessandro Treves, Edvard I Moser, and May-Britt Moser. Place cells in the hippocampus: eleven maps for eleven rooms. *Proceedings of the National Academy of Sciences*, 111(52):18428–18435, 2014.
- Fritzie Arce, Itai Novick, Yael Mandelblat-Cerf, and Eilon Vaadia. Neuronal correlates of memory formation in motor cortex after adaptation to force field. *Journal of Neuroscience*, 30(27):9189–9198, 2010.

- Anil Cherian, Hugo L Fernandes, and Lee E Miller. Primary motor cortical discharge during force field adaptation reflects muscle-like dynamics. *Journal of neurophysiology*, 110(3):768–783, 2013.
- Benjamin R Cowley, Adam C Snyder, Katerina Acar, Ryan C Williamson, Byron M Yu, and Matthew A Smith. Slow drift of neural activity as a signature of impulsivity in macaque visual and prefrontal cortex. *Neuron*, 108(3):551–567, 2020.
- Daniel Deitch, Alon Rubin, and Yaniv Ziv. Representational drift in the mouse visual cortex. *Current Biology*, 31(19):4327–4339, 2021.
- Shaul Druckmann and Dmitri B Chklovskii. Neuronal circuits underlying persistent representations despite time varying activity. *Current Biology*, 22(22):2095–2103, 2012.
- Lea Duncker, Laura Driscoll, Krishna V Shenoy, Maneesh Sahani, and David Sussillo. Organizing recurrent network dynamics by task-computation to enable continual learning. *Advances in neural information processing systems*, 33:14387–14397, 2020.
- Juan A Gallego, Matthew G Perich, Stephanie N Naufel, Christian Ethier, Sara A Solla, and Lee E Miller. Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature communications*, 9(1):1–13, 2018.
- Karunesh Ganguly and Jose M Carmena. Emergence of a stable cortical map for neuroprosthetic control. *PLoS biology*, 7(7):e1000153, 2009.
- Giuseppe P Gava, Stephen B McHugh, Laura Lefèvre, Vítor Lopes-dos Santos, Stéphanie Trouche, Mohamady El-Gaby, Simon R Schultz, and David Dupret. Integrating new memories into the hippocampal network activity space. *Nature neuroscience*, 24(3):326–330, 2021.
- Samuel J Gershman, Marie-H Monfils, Kenneth A Norman, and Yael Niv. The computational nature of memory modification. *Elife*, 6:e23763, 2017.
- Matthew D Golub, Byron M Yu, and Steven M Chase. Internal models for interpreting neural population activity during sensorimotor control. *Elife*, 4:e10015, 2015.
- Matthew D Golub, Steven M Chase, Aaron P Batista, and Byron M Yu. Brain-computer interfaces for dissecting cognitive processes underlying sensorimotor control. *Current opinion in neurobiology*, 37:53–58, 2016.
- Matthew D Golub, Patrick T Sadtler, Emily R Oby, Kristin M Quick, Stephen I Ryu, Elizabeth C Tyler-Kabara, Aaron P Batista, Steven M Chase, and Byron M Yu. Learning by neural reassociation. *Nature neuroscience*, 21(4):607–616, 2018.

- Michael Graziano. The organization of behavioral repertoire in motor cortex. *Annu. Rev. Neurosci.*, 29:105–134, 2006.
- Tanuj Gulati, Ling Guo, Dhakshin S Ramanathan, Anitha Bodepudi, and Karunesh Ganguly. Neural reactivations during sleep determine network credit assignment. *Nature neuroscience*, 20(9):1277–1284, 2017.
- Jay A Hennig, Matthew D Golub, Peter J Lund, Patrick T Sadtler, Emily R Oby, Kristin M Quick, Stephen I Ryu, Elizabeth C Tyler-Kabara, Aaron P Batista, Byron M Yu, et al. Constraints on neural redundancy. *Elife*, 7:e36774, 2018.
- Jay A Hennig, Emily R Oby, Matthew D Golub, Lindsay A Bahureksa, Patrick T Sadtler, Kristin M Quick, Stephen I Ryu, Elizabeth C Tyler-Kabara, Aaron P Batista, Steven M Chase, et al. Learning is shaped by abrupt changes in neural engagement. *Nature Neuroscience*, 24(5):727–736, 2021a.
- Jay A Hennig, Emily R Oby, Darby M Losey, Aaron P Batista, Byron M Yu, and Steven M Chase. How learning unfolds in the brain: toward an optimization view. *Neuron*, 109(23):3720–3735, 2021b.
- David J Herzfeld, Pavan A Vaswani, Mollie K Marko, and Reza Shadmehr. A memory of errors in sensorimotor learning. *Science*, 345(6202):1349–1353, 2014.
- Eun Jung Hwang, Paul M Bailey, and Richard A Andersen. Volitional control of neural activity relies on the natural motor repertoire. *Current Biology*, 23(5):353–361, 2013.
- Beata Jarosiewicz, Steven M Chase, George W Fraser, Meel Velliste, Robert E Kass, and Andrew B Schwartz. Functional network reorganization during learning in a brain-computer interface paradigm. *Proceedings of the National Academy of Sciences*, 105(49):19486–19491, 2008.
- Brian B Jeon, Thomas Fuchs, Steven M Chase, and Sandra J Kuhlman. Existing function in primary visual cortex is not perturbed by new skill acquisition of a non-matched sensory task. *Nature communications*, 13(1):1–16, 2022.
- Sheena A Josselyn and Susumu Tonegawa. Memory engrams: Recalling the past and imagining the future. *Science*, 367(6473), 2020.
- Matthew T Kaufman, Mark M Churchland, Stephen I Ryu, and Krishna V Shenoy. Cortical activity in the null space: permitting preparation without movement. *Nature neuroscience*, 17(3):440–448, 2014.
- Risa Kawai, Timothy Markman, Rajesh Poddar, Raymond Ko, Antoniu L Fantana, Ashesh K Dhawale, Adam R Kampff, and Bence P Ölveczky. Motor cortex is required for learning but not for executing a motor skill. *Neuron*, 86(3):800–812, 2015.

- Aaron C Koralek, Xin Jin, John D Long II, Rui M Costa, and Jose M Carmena. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature*, 483(7389):331–335, 2012.
- John W Krakauer, Claude Ghez, and M Felice Ghilardi. Adaptation to visuomotor transformations: consolidation, interference, and forgetting. *Journal of Neuroscience*, 25(2):473–478, 2005.
- Chiang-Shan Ray Li, Camillo Padoa-Schioppa, and Emilio Bizzi. Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron*, 30(2):593–607, 2001.
- Alexandra Libby and Timothy J Buschman. Rotational dynamics reduce interference between sensory and memory representations. *Nature neuroscience*, 24(5):715–726, 2021.
- Nicolas Y Masse, Gregory D Grant, and David J Freedman. Alleviating catastrophic forgetting using context-dependent gating and synaptic stabilization. *Proceedings of the National Academy of Sciences*, 115(44):E10467–E10475, 2018.
- William Mau, Michael E Hasselmo, and Denise J Cai. The brain in motion: How ensemble fluidity drives memory-updating and flexibility. *Elife*, 9:e63550, 2020.
- Wolf Muellbacher, Ulf Ziemann, Joerg Wissel, Nguyet Dang, Markus Kofler, Stefano Facchini, Babak Boroojerdi, Werner Poewe, and Mark Hallett. Early consolidation in human primary motor cortex. *Nature*, 415(6872):640–644, 2002.
- Karim Nader and Oliver Hardt. A single standard for memory: the case for reconsolidation. *Nature Reviews Neuroscience*, 10(3):224–234, 2009.
- Edward H Nieh, Manuel Schottdorf, Nicolas W Freeman, Ryan J Low, Sam Lewallen, Sue Ann Koay, Lucas Pinto, Jeffrey L Gauthier, Carlos D Brody, and David W Tank. Geometry of abstract learned knowledge in the hippocampus. *Nature*, pages 1–5, 2021.
- Emily R Oby, Matthew D Golub, Jay A Hennig, Alan D Degenhart, Elizabeth C Tyler-Kabara, Byron M Yu, Steven M Chase, and Aaron P Batista. New neural activity patterns emerge with long-term learning. *Proceedings of the National Academy of Sciences*, 116(30):15210–15215, 2019.
- Amy L Orsborn and Bijan Pesaran. Parsing learning in networks using brain–machine interfaces. *Current opinion in neurobiology*, 46:76–83, 2017.
- German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.

- Matthew G Perich and Lee E Miller. Altered tuning in primary motor cortex does not account for behavioral adaptation during force field learning. *Experimental brain research*, 235(9):2689–2704, 2017.
- Matthew R Roesch and Carl R Olson. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science*, 304(5668):307–310, 2004.
- Daniel B Rubin, Tommy Hosman, Jessica N Kelemen, Anastasia Kapitonava, Francis R Willett, Brian F Coughlin, Eric Halgren, Eyal Y Kimchi, Ziv M Williams, John D Simeral, et al. Learned motor patterns are replayed in human motor cortex during sleep. *Journal of Neuroscience*, 42(25):5007–5020, 2022.
- Michael E Rule, Timothy O’Leary, and Christopher D Harvey. Causes and consequences of representational drift. *Current opinion in neurobiology*, 58:141–147, 2019.
- Patrick T Sadtler, Kristin M Quick, Matthew D Golub, Steven M Chase, Stephen I Ryu, Elizabeth C Tyler-Kabara, Byron M Yu, and Aaron P Batista. Neural constraints on learning. *Nature*, 512(7515):423–426, 2014.
- Carl E Schoonover, Sarah N Ohashi, Richard Axel, and Andrew JP Fink. Representational drift in primary olfactory cortex. *Nature*, 594(7864):541–546, 2021.
- Reza Shadmehr and Henry H Holcomb. Neural correlates of motor memory consolidation. *Science*, 277(5327):821–825, 1997.
- Reza Shadmehr and Ferdinando A Mussa-Ivaldi. Adaptive representation of dynamics during learning of a motor task. *Journal of neuroscience*, 14(5):3208–3224, 1994.
- Hannah R Sheahan, David W Franklin, and Daniel M Wolpert. Motor planning, not execution, separates motor memories. *Neuron*, 92(4):773–779, 2016.
- Xulu Sun, Daniel J O’Shea, Matthew D Golub, Eric M Trautmann, Saurabh Vyas, Stephen I Ryu, and Krishna V Shenoy. Cortical preparatory activity indexes learned motor memories. *Nature*, pages 1–6, 2022.
- Cheng Tang, Roger Herikstad, Aishwarya Parthasarathy, Camilo Libedinsky, and Shih-Cheng Yen. Minimally dependent activity subspaces for working memory and motor preparation in the lateral prefrontal cortex. *Elife*, 9:e58154, 2020.
- Jonathan R Wolpaw. The negotiated equilibrium model of spinal cord function. *The Journal of physiology*, 596(16):3469–3491, 2018.
- Yang Xie, Peiyao Hu, Junru Li, Jingwen Chen, Weibin Song, Xiao-Jing Wang, Tianming Yang, Stanislas Dehaene, Shiming Tang, Bin Min, et al. Geometry of sequence working memory in macaque prefrontal cortex. *Science*, 375(6581):632–639, 2022.

Guangyu Robert Yang, Madhura R Joglekar, H Francis Song, William T Newsome, and Xiao-Jing Wang. Task representations in neural networks trained to perform many cognitive tasks. *Nature neuroscience*, 22(2):297–306, 2019.

347 Methods

348 **Experimental procedures.** Experimental methods are detailed in our previous
349 work (Sadtler et al., 2014, Golub et al., 2018). Briefly, we recorded neural activity
350 from three male Rhesus macaques (*Maccaca mulatta*, ages 7, 7 and 8 for monkeys
351 J, N and L, respectively) using 96 electrode arrays (Blackrock Microsystems) im-
352 planted in the proximal arm region of the primary motor cortex. All animal care
353 and handling procedures conformed to the NIH Guidelines for the Care And Use of
354 Laboratory Animals and were approved by the University of Pittsburgh’s Institu-
355 tional Animal Care and Use Committee.

356
357 The monkeys performed an eight-target center-out BCI task. In the BCI, a mon-
358 key guided a computer cursor by modulating its neural activity. The recorded neural
359 activity was translated into movements of the computer cursor according to a BCI
360 map (see *Translating neural activity to cursor movement*). Each session was split
361 into three task periods, “Task A1”, “Task B”, and “Task A2”. The three task peri-
362 ods followed the same experimental paradigm, differing only in the BCI map. During
363 Task A1 the monkey used Map A, which was selected to be intuitive for the monkey
364 to use from the outset. The monkey controlled the cursor during Task A1 for 318.8
365 ± 95.4 (mean \pm s.d.) trials. Uncued to the monkey, we then switched to Map B for
366 the second period of the experiment (Task B). The monkey had never seen before
367 Map B and was selected in order to initially be difficult for the monkey to use to
368 control the cursor. The monkey was given 696.7 ± 219.4 (mean \pm s.d.) trials to
369 learn to control the cursor with Map B. Finally, again uncued, Map A was reinstated
370 (Task A2). The Task A2 period lasted the remainder of the experiment for $318.2 \pm$
371 153.9 (mean \pm s.d.) trials.

372
373 *Trial flow.* At the start of each trial, the cursor appeared at the center of the mon-
374 key’s workspace. Target locations were selected pseudo-randomly from a set of eight
375 uniformly spaced locations around a circle (radius, Monkey J: 150 mm; Monkeys L
376 and N: 125 mm). The target appeared on the screen at the beginning of the trial.
377 For the first 300 ms, the cursor’s velocity was fixed at zero. After this, the velocity of
378 the cursor was controlled by the monkey through the BCI map corresponding to the
379 task period of the experiment. If the monkey was able to acquire the target within
380 7.5s after the start of the trial, a water reward was dispersed. If the monkey failed
381 to acquire the target within the allotted time, there was a 1.5s timeout prior to the
382 start of the next trial.

383
384 *Identifying latent dimensions of neural activity.* Experiments began with a cal-
385 ibration period in order to define Map A. Monkey J’s calibration employed either
386 passive cursor observation or closed-loop BCI control using the previous day’s BCI
387 map. For monkeys L and N, we used a calibration procedure that gradually stepped
388 from passive observation to closed-loop control. We then applied factor analysis (see
389 below) to identify the 10D linear subspace (the “intrinsic manifold”) that captured

390 the dimensions of greatest shared variance in the neural population. Ten dimensions
 391 was selected using cross-validation, as described in prior work (Sadtler et al., 2014).

392 Spike counts (i.e. threshold crossings) were taken in nonoverlapping 45 ms time
 393 windows. We denote the spike counts at timestep t as $\mathbf{u}_t \in \mathbb{R}^{q \times 1}$, where q is the
 394 number of neural units. Factor analysis describes this high-dimensional population
 395 activity, \mathbf{u}_t , in terms of a low-dimensional set of factors, $\mathbf{z}_t \in \mathbb{R}^{10 \times 1}$. Latent factors,
 396 \mathbf{z}_t , are distributed as:

$$\mathbf{z}_t \sim N(\mathbf{0}, I) \quad (1)$$

397 where I is the identity matrix. Spike counts, \mathbf{u}_t , are related to the factors by:

$$\mathbf{u}_t | \mathbf{z}_t \sim N(L\mathbf{z}_t + \boldsymbol{\mu}, \Psi) \quad (2)$$

398 where parameters $L \in \mathbb{R}^{q \times 10}$ (termed the loading matrix), $\boldsymbol{\mu} \in \mathbb{R}^{q \times 1}$, and $\Psi \in \mathbb{R}^{q \times q}$
 399 (a diagonal matrix of variances independent to each neuron) are estimated using the
 400 expectation-maximization algorithm. The latent factor activity, \mathbf{z}_t , at timestep t is
 401 estimated as the posterior expectation given the spike counts as:

$$\mathbf{z}_t = L^T(LL^T + \Psi)^{-1}(\mathbf{u}_t - \boldsymbol{\mu}) \quad (3)$$

402 For all analyses, we orthonormalized \mathbf{z}_t so that it had units of spike counts per
 403 timestep to facilitate the interpretability of the factor activity. As the majority of
 404 the shared variance of the neural population is captured in these latent dimensions,
 405 and neural activity cannot be readily produced outside this low-dimensional subspace
 406 during short-term learning (Sadtler et al., 2014, Oby et al., 2019), we focus our anal-
 407 yses on this factor activity, referred to as “population activity patterns” throughout.

408 *Translating neural activity to cursor movement.* At each 45 ms timestep t , neural
 409 activity drove the computer cursor according to the BCI map for that task period.
 410 Specifically, the cursor velocity was determined using a Kalman filter:

$$\mathbf{v}_t = M_1\mathbf{v}_{t-1} + M_2\mathbf{z}_t + \mathbf{m}_0 \quad (4)$$

412 The parameters $M_1 \in \mathbb{R}^{2 \times 2}$, $M_2 \in \mathbb{R}^{2 \times 10}$ and $\mathbf{m}_0 \in \mathbb{R}^{2 \times 1}$ are determined during the
 413 calibration period (see Sadtler et al. (2014) for details), and $\mathbf{v}_t \in \mathbb{R}^{2 \times 1}$ comprises
 414 the horizontal and vertical cursor velocities. The two BCI maps differ only in the
 415 M_2 term. For Map A, $M_2 = M_2^{(A)}$, which is found during the calibration session.
 416 For Map B, $M_2 = M_2^{(B)}$ was a permutation applied to the columns of $M_2^{(A)}$, equiv-
 417 alent to permuting the elements of \mathbf{z}_t before applying equation 4. This means that
 418 Map B remained within the intrinsic manifold (a “within-manifold perturbation”).
 419 Thus Map B changed the relationship between the factor activity and cursor velocity.

420 **Data Analysis** The data analyzed in this study was part of a larger study that
 421 included both within-manifold perturbations (WMPs) and outside-manifold pertur-
 422 bations (OMPs) (Sadtler et al., 2014). As we have previously found that WMPs
 423 show stronger learning than OMPs, we only considered sessions that used WMPs.
 424

425 Data from the Task A1 and Task B periods of these WMP sessions were analyzed in
426 prior work (Golub et al., 2018, Hennig et al., 2018, 2021a). Here we focused on neural
427 activity recorded during Task A2, which has not been previously studied. To ensure
428 an adequate amount of Task A2 data to analyze per session, we only considered
429 sessions that included at least 100 Task A2 trials. This yielded a total of 43 sessions
430 (Monkey J, 22 sessions, 362.6 ± 170.2 Task A2 trials; Monkey N, 12 sessions, 333.3
431 ± 107.3 Task A2 trials; Monkey L, 9 sessions, 171.0 ± 49.7 Task A2 trials; all values
432 mean \pm s.d.).

433
434 *Selecting experiments and trials for analysis.* Some targets showed more learning
435 than others. As the focus of this work is on the memory of a learned task, we analyzed
436 targets that showed the most learning. We defined *learning* as how well the monkey
437 performed with Map B after learning, relative to how well it would have performed
438 with Map B if it continued producing the same neural activity as it did during Task
439 A1 (i.e., if there was no learning). Thus, we defined learning as the difference in
440 the average Map B progress (see *Quantifying the memory trace* for how progress is
441 computed) of the last 10 trials to a given target during Task B compared to average
442 Map B progress to that same target during Task A1. For each monkey, the 50%
443 of targets with the most learning were designated as “well-learned.” Well-learned
444 targets had an average amount of learning of 26.61 ± 13.07 mm/s, compared to 0.69
445 ± 8.79 mm/s for the other targets. Fig. 2a, Fig. 2b, Fig. 4a, Fig. 4c, Extended
446 Data Fig. 3 and Extended Data Fig. 4 include all targets. All other analyses focus
447 on the well-learned targets.

448 As our central question focuses on neural activity during proficient Task A2 per-
449 formance, we restricted analyses of Task A2 to after behavior had stabilized. To do
450 this, we excluded the first 50 trials of Task A2 from each session (see Fig. 4). Unless
451 stated otherwise, the remaining Task A2 trials are referred to as Task A2 throughout
452 the manuscript. Additionally, we only analyzed successful trials, as it is otherwise
453 difficult to determine whether the monkey was engaged in the task.

454 On each trial, we discarded the first 90 ms (2 timesteps during the freeze period) as
455 the activity in M1 would not yet reflect the target due to sensory processing delays
456 (Golub et al., 2015). Additionally, because we report trial-averaged and target-
457 averaged quantities, we wanted to ensure neural activity came from instances in
458 which the monkey needed to push the cursor in the same direction. Thus, we only
459 analyzed timesteps in which the angle between the cursor and the target was no
460 greater than 22.5° away from the target direction for that trial. Performing our
461 analyses without this exclusion criterion did not change our results qualitatively.

462 Even after learning to use Map B, the monkeys generally exhibited lower perfor-
463 mance with Map B than Map A (see Fig. 1c). Thus, Task B trials tended to be
464 longer than the Task A1 and A2 trials. To compare the Task B trials to the Task A1
465 and A2 trials, we only utilized the first 25 timesteps from each trial. This number
466 was selected because it is approximately equal to the average Task A1 acquisition
467 time across all monkeys.

468

469 *Testing the reversion hypothesis.* To measure tuning changes between task periods
470 (Fig. 2), we fit cosine tuning curves for each neural unit using ordinary least squares
471 regression:

$$\lambda(\theta) = r_0 + (r_{\max} - r_0) \cos(\theta - \theta_{pd})$$

472 where $\lambda(\theta)$ is the estimated firing rate for a given cursor-target direction θ . The
473 parameters θ_{pd} , r_0 and r_{\max} can be interpreted as the preferred direction, the average
474 firing rate, and the tuning amplitude of the unit, respectively. For each neural unit,
475 we fit a separate tuning curve for each task period of the experiment.

476 We compared the preferred direction θ_{pd} for each neural unit between Tasks A1
477 and A2 by computing the average absolute change in preferred direction (Fig. 2b).
478 To calculate the control distribution, for each neural unit, we randomly permuted
479 the task labels for each timestep during Task A1 and Task A2. The difference in
480 preferred direction between Task A1 and A2 was then recalculated using these new
481 task labels.

482 To visualize how neural activity changes in the 10D population space, we applied
483 linear discriminant analysis to \mathbf{z}_t , taken in 45ms timesteps, in to order to find the
484 2D plane that best separates the activity from the three task periods (Fig. 2c). We
485 applied a QR decomposition in order to orthonormalize the basis vectors found by
486 LDA, then projected the neural activity onto this orthonormal basis.

487 To quantify the changes in population activity between Task A1 and Task A2, we
488 calculated the Mahalanobis distance on a per-target basis between the population
489 activity means across \mathbf{z}_t , taken in 45ms timesteps, for each task period (Fig. 2d).
490 This distance was computed in the 10D space, using the covariance of the Task A1
491 neural activity for that target. To calculate the control distribution, for each target,
492 we randomly permuted the task labels for each timestep during Task A1 and Task
493 A2. The Mahalanobis distance between the mean activity for each target was recal-
494 culated using the new task labels.

495
496 *Defining the memory trace.* Progress quantifies the appropriateness of a particular
497 population activity pattern for a particular BCI map, i.e., the extent to which that
498 population activity pattern drives the cursor towards the target, and is computed
499 as follows. First, we determine the neural push of this activity pattern, \mathbf{z}_t , through
500 a particular map, M_2 , as $M_2\mathbf{z}_t$. In equation 4, \mathbf{m}_0 and M_1 do not rely on the
501 instantaneous neural activity, and so we do not consider the contributions from these
502 terms. Next, we compute the component of this neural push in the direction of the
503 target. More specifically, for each timestep t , we define a unit vector, $\mathbf{c}_t \in \mathbb{R}^{2 \times 1}$,
504 pointing from the current location of the cursor to the target. Thus, the progress at
505 timestep t is evaluated as:

$$p_t = \mathbf{c}_t^T M_2 \mathbf{z}_t \quad (5)$$

506 We sought to determine how much more appropriate neural activity is for Map B
507 during Task A2 than it is during Task A1. We call this change in appropriateness
508 a “memory trace” because it measures the lasting alteration of neural activity used

509 during a familiar task (Map A) after a learning experience (Map B). Specifically, we
510 define the memory trace as the difference in progress when neural activity is passed
511 through Map B during Task A2 minus that during Task A1.

512

513 *Testing how Task A2 duration affects the memory trace.* We sought to determine
514 whether the memory trace persisted over time (Fig. 4a and Fig. 4b). For each mon-
515 key, we divided the sessions into two groups based on whether the Task A2 period
516 was longer or shorter than the median length across all sessions (300 trials). This
517 resulted in 22 sessions in the long Task A2 group (14/22 sessions from Monkey J,
518 average length of 464.36 ± 119.76 Task A2 trials; 8/12 sessions from Monkey N,
519 400.00 ± 53.45 , 0/9 sessions from Monkey L; all values are mean +/- s.d.) and 21
520 sessions in the short Task A2 group. In order to focus on trials where the monkey
521 had longer exposure to Task A2, we excluded the first 200 trials when calculating
522 the memory trace, leaving at least 100 trials of Task A2 for analysis. For the short
523 Task A2 group, we excluded the first 50 trials of Task A2 as usual (see *Selecting*
524 *experiments and trials for analysis*).

525

526 *Testing how Task A2 behavior affects the memory trace.* We additionally sought to
527 determine whether the memory trace differed as a function of performance through
528 Map A (Fig. 4c and Fig. 4d). To address this, for each target we compared the
529 average progress through Map A during Task A2 to that during Task A1. Targets
530 with acquisition times during Task A2 that were at least as good as Task A1 were
531 placed in the “better behavior group”. There were 21 targets in this group, with an
532 average of 75.0 ± 57.7 ms (mean \pm s.d.) faster target acquisition in Task A2 relative
533 to Task A1. Targets which had acquisition times during Task A2 that were worse
534 than Task A1 were placed in the “worse behavior group”. There were 145 targets
535 in this group, with an average of $241.3 \text{ ms} \pm 210.2 \text{ ms}$ (mean \pm s.d.) slower target
536 acquisition in Task A2 relative to Task A1.

537

538 *Decomposing the memory trace into output-potent and output-null components.* In
539 order to determine how the memory trace can coexist without degrading behavioral
540 performance during Task A2, we wanted to determine how changes in neural activ-
541 ity between Task A1 and Task A2 relate to Map A. To address this question, we
542 decomposed neural activity into a component that is output-potent to Map A and a
543 component that is output-null to Map A (Fig. 5). This decomposition was done by
544 applying the singular value decomposition (SVD) to Map A:

545

$$M_2^{(A)} = UDV^T \quad (6)$$

546 where $U \in \mathbb{R}^{2 \times 10}$, $D \in \mathbb{R}^{10 \times 10}$, and $V \in \mathbb{R}^{10 \times 10}$. D is a diagonal matrix, whose
547 diagonal elements are the singular values of $M_2^{(A)}$. As $M_2^{(A)}$ is a matrix of rank two,
548 only the first two diagonal entries of D are non-zero. This means that the first two
549 columns of V form an orthonormal basis for the output-potent space of $M_2^{(A)}$. We
550 denote this basis as $R \in \mathbb{R}^{10 \times 2}$. The last 8 columns of V form an orthonormal basis

551 of the output-null space of $M_2^{(A)}$. We denote this basis as $N \in \mathbb{R}^{10 \times 8}$.


552 We can find the component of neural activity potent to Map A as $\mathbf{z}_t^{\text{pot}} = RR^T \mathbf{z}_t$.
553 Similarly, the null component is found as $\mathbf{z}_t^{\text{null}} = NN^T \mathbf{z}_t$. Both $\mathbf{z}_t^{\text{pot}}$ and $\mathbf{z}_t^{\text{null}}$ are
554 10×1 vectors, and have the property that $\mathbf{z}_t = \mathbf{z}_t^{\text{pot}} + \mathbf{z}_t^{\text{null}}$. We calculate the potent
555 and null component of the memory trace as before, except utilizing $\mathbf{z}_t^{\text{pot}}$ and $\mathbf{z}_t^{\text{null}}$ for
556 \mathbf{z}_t respectively in equation (4). This decomposition is utilized in Fig. 5 and Extended
557 Data Fig. 5. Note that this decomposition is performed with respect to Map A and
558 not with respect to Map B. This is because, by definition, the memory trace must
559 be in output-potent dimensions of Map B, as those are the only dimensions that
560 determine the cursor velocity through Map B.

561
562 *Path of learning and washout.* To distinguish whether the path of washout retraces
563 the path of learning (Fig. 6), we first define the path of learning as the vector in 10D
564 neural activity space from the mean activity during Task A1 to the mean activity
565 during the late Task B period (see *Selecting experiments and trials for analysis*).
566 We similarly define the path of washout as the 10D vector between the mean neural
567 activity during late Task B and the mean activity during Task A2. We then com-
568 pared the paths of learning and washout by finding the the angle between these two
569 vectors. To obtain a control distribution, for each target, we randomly permuted the
570 task labels for each timestep during Task A1 and Task A2. This mimics a situation
571 in which Task A1 and Task A2 activity patterns come from the same distribution.
572 As task labels for Task B were not shuffled, the paths of learning and washout would
573 thus be equal and opposite on average under this construction. The angle between
574 the paths for each target was recalculated using the new task labels.

575
576 **Statistics.** Unless otherwise noted, to test for statistical significance, we used
577 nonparametric tests (for example, Wilcoxon signed-rank test or ranked-sum test),
578 which do not assume normality. All P-values less than 10^{-10} were reported as
579 $P < 10^{-10}$, regardless of how small the P-value was.


580
581 **Data availability** The data that support the findings of this study are available
582 from the authors upon reasonable request.

583
584 **Code availability** Python code that supports the data analyses will be made
585 publicly available upon publication.




Extended Data Fig. 1. Comparison of behavioral performance in Tasks A1 and A2

Here we plot the average acquisition time for a given target during Task A1 against its average acquisition time during Task A2. Performance in Task A2 tended to be lower than performance in Task A1, likely due to satiation or fatigue. In Fig. 4 and Extended Data Fig. 3 we demonstrate that this difference in behavior is not the cause of the memory trace. The targets that fall below the diagonal are those in which performance during Task A2 is better than during Task A1, and are the same targets that are included in the “better behavior group” in Fig. 4d.




Extended Data Fig. 2. A memory trace is also evident when measured using angular error

In Fig. 3d, we measured the memory trace using “progress”, which is defined as the velocity by which a neural activity pattern would have moved the cursor toward the target (see Methods). We could have alternatively measured the memory trace in terms of angular error instead of progress. In contrast to progress which depends on velocity magnitude and direction, angular error depends only on the velocity direction. Angular error is defined at each timestep as the angular difference between the velocity vector of the neural push and the cursor-to-target direction. As with progress, the velocity of the neural push is defined using the Task A1 (or Task A2) neural activity projected through Map B. We then compute the angular error for Task A1 minus the angular error for Task A2. We use unsigned angular error so clockwise and counterclockwise errors do not cancel each other out when averaging. Smaller angular errors are better. Thus, when angular error is smaller for Task A2 relative to Task A1, a memory trace is present ($P < 10^{-10}$, two-sided paired Wilcoxon sign-rank test, $n=88$ targets; Monkey N, $P = 1.02 \times 10^{-7}$, $n=48$; $P = 0.0027$, $n=36$ targets). The white tick mark on the horizontal axis of the middle histogram denotes the example target illustrated in Fig. 3a, Fig. 3b and Fig. 3c.




Extended Data Fig. 3. Most targets per session exhibit a memory trace

We considered whether the memory trace could be due to a global shift in neural activity (e.g., due to a neural recording instability) that leads to an increase in Map B progress for some targets at the expense of progress for targets on the opposite side of the monkey's workspace. If this were the case, we would expect that only half of the targets in each session, including targets that show little or no learning, would show a memory trace. **(a)** Instead, we found that more than half of the targets per session showed a memory trace ($P = 7.21 \times 10^{-5}$, two-sided paired Wilcoxon signed-rank test, $n=43$ sessions across monkeys). **(b)** Similarly, we found that the average memory trace across all eight targets per session is positive ($P = 4.25 \times 10^{-5}$, two-sided paired Wilcoxon signed-rank test, $n=43$ sessions across monkeys).



Extended Data Fig. 4. The amount of learning is correlated with the size of the memory trace

If the memory trace arose spuriously and not as the result of the learning experience, we would expect the size of the memory trace (measured using neural activity during Task A2) to be uncorrelated with the amount of learning (measured using neural activity during Task B). The amount of learning during the Task B period positively correlates with the magnitude of the memory trace (Monkey J $R^2 = 0.25$, $P < 10^{-10}$, one-sided F test, $n=176$ targets; Monkey N, $R^2 = 0.29$, $P = 1.23 \times 10^{-8}$, $n = 96$; Monkey L, $R^2 = 0.13$, $P = 0.0017$, $n = 72$). All targets were included in this analysis, though similar results hold when only examining well-learned targets. We considered the possibility that these results could have arisen trivially due to the memory trace and amount of learning both being calculated relative to Map B progress during Task A1. We thus reran this analysis without subtracting this quantity (that is, regressing Map B progress during Task B with Map B progress during Task A2) and arrived at similar results (Monkey J $R^2 = 0.48$, $P < 10^{-10}$; Monkey N, $R^2 = 0.45$, $P < 10^{-10}$; Monkey L, $R^2 = 0.37$, $P = 1.14 \times 10^{-8}$). This supports the notion that the memory trace is the result of the preceding learning experience. Furthermore, in Fig. 3d, Monkey L showed a smaller memory trace than Monkeys J and N. While the scatter of values for Monkey L lies within the scatter of Monkeys N and J, Monkey L showed less learning on average than Monkeys J and N. This is a possible explanation for Monkey L's smaller memory trace.



Extended Data Fig. 5. The majority of the memory trace resides in dimensions output-null to Map A

To understand which dimensions of neural activity contribute to the memory trace, we decomposed neural activity into components that are output-potent and output-null to Map A and evaluated their contribution to the memory trace (Fig. 5). Here, we breakdown Fig. 5 by target. Targets across all sessions and monkeys are ordered by the total memory trace expressed for that target (black line). The contributions by the potent and null spaces of Map A are shown in purple and magenta, respectively. As the total memory trace is the sum of the contributions from the output-potent and output-null components, it is possible for one of these components to have a negative contribution and the total memory trace to still be positive. A negative value indicates progress through Map B is smaller during Task A2 relative to Task A1 for that component. For visual clarity, we use dark shading for positive values and light shading for negative values. For a given target, there is one purple bar (light or dark) and one magenta bar (light or dark). We find the majority of the memory trace lies in resides in dimensions output-null to Map A (magenta bars tend to be larger than purple bars), as quantified in Fig. 5.