Supplementary Materials for

New neural activity patterns emerge with long-term learning

Emily R. Oby[1,2], Matthew D. Golub[2,3,4,5], Jay A. Hennig[2,6,7], Alan D. Degenhart[1,2], Elizabeth C. Tyler-Kabara[1,8,9,10], Byron M. Yu[2,5,11]†, Steven M. Chase[2,11]†, Aaron P. Batista[1,2]†*

Correspondence to: aaron.batista@pitt.edu

This file includes:

      Materials and Methods

      Figs. S1 to S7

      Table S1

      Caption for Movie S1

Other Supplementary Materials for this manuscript include the following:

      Movie S1

MATERIALS AND METHODS

All animal procedures were approved by the University of Pittsburgh Institutional Animal Care and Use Committee in accordance with the guidelines of the US Department of Agriculture, the International Association for the Assessment and Accreditation of Laboratory Animal Care, and the National Institutes of Health.

Electrophysiology and behavioral monitoring

We recorded neural activity from the arm region of the primary motor cortex (M1) in two male rhesus monkeys (*Macaca mulatta*, age, monkey L: 9 years; monkey N: 7 years) using 96-channel microelectrode arrays (Blackrock Microsystems, Salt Lake City, UT, USA). We recorded neural units as threshold crossings, where a threshold crossing was detected when the depolarizing phase of the voltage signal crossed a threshold of -3 times the root-mean-square voltage. We estimated the root-mean-square voltage of the signal on each electrode while the monkeys sat calmly in a darkened room. For multi-day experiments, we set the threshold on day one and did not adjust it again. Across multi-day experiments, we recorded 89.9±1.9 neural units (all quantities are listed as mean±SD unless otherwise noted) for monkey L, and 93.2±1.8 units for monkey N; within a given OMP experiment the number of recorded neural units was constant throughout. The number of threshold crossings were counted in non-overlapping 45-ms bins. The data were recorded between 18 and 41 months after array implantation for monkey L and from 2.5 to 8 months after array implantation for monkey N.

Both arms were loosely restrained throughout all experiments. While the monkeys could move their forearms slightly and there was no restriction on wrist and hand movement, the arm movements during the experiments were minimal.

Behavioral task

Throughout each experiment animals performed an eight-target center-out task under brain-computer interface control. Each trial began with the cursor appearing at the center of the workspace. A peripheral target 125 mm away appeared simultaneously. Target locations were uniformly spaced 45 degrees apart around a circle. Targets were presented in a random (monkey N) or a pseudo-random (monkey L) order. Pseudo-random target presentation consisted of 8-trial blocks, in which target location was randomly chosen without replacement, and the next block of trials began only after all eight targets were successfully acquired. The cursor was frozen at the center of the workspace for the first 300 ms of the trial. After this "freeze period" the cursor began moving under neural control. Details about the neural control are given in the *Intuitive BCI mappings* and *Perturbed BCI mappings* sections below. Target acquisition occurred when the cursor entered the target window (circle, radius = 14 mm, monkey L; 35 mm, monkey N). If the target was successfully acquired within 7.5 s, the trial was deemed a success, the monkey was given a water reward, and the next trial began 200 ms later. If the target was not acquired within the 7.5 s allotted time, there was a 1.7 s timeout before the next trial.

Identifying the intrinsic manifold

The Intrinsic Manifold is defined as the low-dimensional space that describes the neural activity patterns generated by the animal prior to learning. To estimate it, we used the data collected on day 1, while we calibrated the intuitive BCI mapping (as described below *Intuitive BCI mappings*). We used factor analysis to describe each high-dimensional spike count vector in terms of a low-dimensional set of factors. Under factor analysis, factors are described by a standard normal distribution:

$$z_t \sim N(0, I) \quad (1)$$

where I is the identity matrix. The neural activity is related to those factors by a normal distribution:

$$u_t | z_t \sim N(\Lambda z_t, \psi) \quad (2)$$

where $u_t \in \mathbb{R}^{q \times 1}$ is the z-scored spike count vector across q simultaneously-recorded neural units at timestep t. Each neural unit was z-scored separately. $z_t \in \mathbb{R}^{p \times 1}$ is the set of p factors at timestep t. $\Lambda \in \mathbb{R}^{q x p}$ and $\psi \in \Sigma_+^q$ (a diagonal $q \times q$ covariance matrix) were estimated using the expectation-maximization algorithm (*37*). The intrinsic manifold is defined as the column space of Λ, where each factor corresponds to a column of Λ. By z-scoring we ensure that animals are not required to produce firing rates that are outside of the observed physiological range of each neural unit in order to learn how to control the cursor while using the perturbed BCI mappings.

We set the dimensionality of the intrinsic manifold to be 10 for all experiments. This was the dimensionality used in Sadtler et al. (*8*), and is consistent with the dimensionality of the neural activity in the current study. Here, we computed the dimensionality as the number of dimensions needed to explain 95% of the shared variance (*38*). This method differs slightly from that used in Sadtler et al. (*8*). We used the method described in (*38*) because it returns a more reliable estimate of dimensionality, even though it tends to underestimate the "true" dimensionality. The underestimation occurs because it seeks only to explain 95% of the shared variance. Using this method, the multi-day experiments had a mean dimensionality of 8.4 (range: 5-11) across both monkeys (computed using 3500-5500 time points for each session). (For each monkey individually, dimensionalities were monkey N: mean = 9.6, range: $8-11$; monkey L: mean = 7.4, range: 5 - 9). This distribution is consistent with the use of 10 dimensions in this study.

Given the neural activity u_t, we extract the factors

$$\hat{z}_t = E[z_t | u_t] = \Lambda^T (\Lambda \Lambda^T + \psi)^{-1} u_t = \beta u_t \quad (3)$$

where $\beta = \Lambda^T (\Lambda \Lambda^T + \psi)^{-1}$. A separate factor analysis model was fit for each experiment based on the neural activity recorded during the calibration trials (as described below *Intuitive BCI mappings*).

Intuitive BCI mappings

BCI mappings translated the neural activity into 2D cursor velocities v_t using a Kalman filter (8). A Kalman filter consists of a state model, which describes how cursor velocity changes over time, and an observation model, which describes how observed neural activity relates to cursor velocity. The state model is:

$$v_t|v_{t-1} \sim N(Av_{t-1}, Q) \qquad (4)$$

where $A = I$ and $Q = 2(m^2/s^2) \times I$. Q controls the smoothness of cursor velocities over time, and was chosen based on the setting that yielded the best closed-loop performance. In a standard Kalman filter, the high-dimensional neural activity is directly translated into cursor velocities. Here, we first estimated the lower-dimensional factors using Eqn. 3, and then translated the factors to cursor velocities. The observation model is:

$$\hat{z}_t|v_t \sim N(Cv_t + d, R) \qquad (5)$$

where $C \in \mathbb{R}^{10 \times 2}$, $d \in \mathbb{R}^{10}$, and $R \in \Sigma_+^{10}$. These parameters were fit using maximum likelihood to relate the factors and intended cursor velocities (defined below) from the calibration trials at the start of each experiment. Taken together, the intuitive BCI mapping can be written as

$$\hat{v}_t = M_1 \hat{v}_{t-1} + M_2 u_t + m_0 \quad (6)$$

$$M_1 = A - KCA$$

$$M_2 = K\beta$$

$$m_0 = -Kd$$

where $\hat{v}_t$ is the velocity used to move the cursor on the screen at time t and K is the converged Kalman gain. M_1 temporally smooths the velocities, M_2 describes the high-dimensional neural activity in terms of the low-dimensional factors and relates the factors to the cursor velocity, and m_0 is a constant offset.

Each experiment began with 80 trials to calibrate an intuitive BCI mapping. Calibration involved a mixture of passive observation of center-out cursor movements and closed-loop BCI cursor control, during which we gradually increased the animal's level of control over the cursor, until all assistance was removed, and he had full control. Calibration began with 16 trials (2 to each target) of the monkey observing automatic center-out cursor movements straight to the target at a constant speed (0.15 m/s). Here, intended cursor velocity at each timestep was taken to be the observed cursor velocity (0.15 m/s) in the center-to-target direction. For the next 8 trials, the animal controlled the cursor using a mapping calibrated using the data from the 16 observation trials, but the cursor was restricted to move only along the center-to-target direction (velocity components perpendicular to the center-to-target direction were set to 0). The next 8 trials used a mapping calibrated from the previous 8 trials, and perpendicular velocity

components were scaled by a factor of 0.2. The next 8 trials used a mapping calibrated from the previous 16 trials, and perpendicular velocity components were scaled by a factor of 0.4. We repeated this procedure for a total of 40 to 80 trials until the animal was given complete control of the cursor (perpendicular scale factor = 1). All calibrations performed within this procedure defined intended cursor velocities to be straight to the target, with speeds taken from the cursor movements that were displayed to the animal. Animals demonstrated proficient cursor control using the intuitive BCI mapping from the very first intuitive trial of each experiment, as evidenced by success rates (100%), fast acquisition times (L: 678.7 ± 42.9 ms; N: 522.3 ± 23.8 ms), and relatively straight cursor trajectories (e.g., Fig. 1C). The data from these 80 calibration trials were used to determine the intrinsic manifold.

Perturbed BCI mappings

Perturbed BCI mappings alter the relationship between recorded neural activity and cursor velocities to induce learning. The OMP mappings are of the same general form as the intuitive mapping (Eqn. 6). OMP mappings modify an intuitive BCI mapping by permuting the elements of u_t (that is, the ~90 neural channels) before passing it into factor analysis. This changes the relationship between the neural activity and factors, but preserves the relationship between the factors and cursor velocity. The impact of this manipulation is that it encourages the monkey to change the way in which neurons co-vary in order to restore proficient cursor control. (In contrast, to generate a WMP mapping, the ten factors were permuted before passing them to the Kalman filter. This manipulation changes the relationship between the factors and cursor velocity, but preserves the relationship between the neural activity and the factors (*8*).) An OMP mapping can be described in terms of the high-dimensional neural activity by:

$$\hat{v}_t = M_1 \hat{v}_{t-1} + M_2^{OMP} u_t + m_0, \qquad (7)$$

where $M_2^{OMP} = M_2 \eta_{OM}$ and η_{OM} is $q \times q$ permutation matrix defining the outside-manifold perturbation (where q indexes the neural channels).

The OMP mappings we chose were not orthogonal to the manifold, as this would result in essentially zero cursor movement, at least initially. Rather, each mapping was selected to be difficult enough to require substantial learning but not so difficult that the animal would give up (*8*).

Behaviorally, these perturbations had complex effects on cursor movements, which cannot be expressed as pure visuomotor rotations or gains. Before learning, the effects of a typical perturbation can be approximately summarized by a combination of per-target velocity rotations and speed scalings. Because the perturbations were implemented in high-dimensional space, these rotations and scalings need not be consistent across movement directions and speeds (as they would be in the case of a pure visuomotor rotation or gain). Perturbations often affected movement speeds more profoundly along one movement direction than along the perpendicular direction (as seen in Fig. 2D-F).

Incremental training paradigm

      We developed an incremental training paradigm to facilitate learning. The incremental training paradigm consisted of giving the monkey a series of incremental mappings which were successively further from the intrinsic manifold and successively closer to the full OMP mapping, to guide the monkey through neural space to the new patterns of neural activity requested by the OMP mapping (Fig. S1A). The incremental mappings were a weighted combination of the intuitive mapping and the OMP mapping, such that

$$M_2^{incremental} = \left(1 - \frac{a}{5}\right) * M_2 + \left(\frac{a}{5}\right) * M_2^{OMP} \qquad (8)$$

where $a = 1, 2, 3, 4$ and 4.5. Incremental step 1 is closer to the intrinsic manifold and should be easier to learn to control and incremental step 4 is closer to the OMP mapping and is expected to be more difficult to learn to control. The inclusion of the additional half step (4.5) was made to increase the monkeys' willingness to continue to attempt the task. The criteria for incrementing the mapping are described below.

Task flow

      Each experiment began with a calibration block during which we identified the intrinsic manifold and determined the parameters of the intuitive mapping. The monkeys used the intuitive mapping for 206 ± 56 trials (monkey L) or 193 ± 3.68 trials (monkey N). We then ran one of four possible OMP experiment types (Fig. 1B): multi-day, incremental training (monkey L: n=9; monkey N: n=6); multi-day, no incremental training (monkey L: n=6; monkey N: n=5); single day, incremental training (monkey L: n=12; monkey N: n=15); single day, no incremental training (monkey L: n=13; monkey N: n=8). Single day WMP experiments are also included (monkey L: n=11; monkey N: n=11).

      We conducted a total of 51 experiments with monkey L. The 13 single day OMP, no incremental training experiments and 11 single day WMP experiments from monkey L have been previously reported (*8*). We conducted a total of 45 experiments with monkey N. Single day WMP experiments from monkey N have been analyzed in previous studies (*9, 15*). Below, we detail the task for each experiment type. The number of each type of experiment is summarized in Table S1.

Multi-day OMP, incremental training

      This section corresponds to the experiments described in Figs. 1-4 and Fig. S1C, D, G (lower right). The procedures for multi-day incremental training experiments differed somewhat for the two monkeys. For monkey N, on day 1 the intuitive mapping block was followed by that experiment's full OMP mapping (434 ± 81 trials). Day 2 began with 40 trials of the full OMP mapping before beginning the incremental training with the incremental step 1 mapping. (These

data were collected to test whether OMP learning simply required a night's sleep. We did not observe this to be the case. These data are not included in the analyses.) The mapping was incremented when the monkey reached a success rate of 80% or greater over the prior 40 trials and target acquisition times appeared to asymptote according to visual inspection by the experimenter (E. Oby). On subsequent days, the monkey began with an incremental mapping (usually step 3) and incremental training continued in the same way. After arriving at the full OMP mapping, the experiment continued for another 2 to 3 days (2.8±1.2 days). At the end of the final day, the intuitive mapping from day 1 was re-introduced for a washout block (356±130 trials).

For monkey L, day 1 of a multi-day, incremental training experiment proceeded in the same manner as a single day, incremental training experiment (as described below) without a washout block. On subsequent days, the monkey began with an incremental mapping (usually step 3) and incremental training continued. That is, the full OMP mapping was not re-presented at the start of day 2, as it was for monkey N. There were not objective criteria for incrementing to the next mapping for monkey L; the mapping was incremented when he showed good performance as judged by the experimenter. This procedure was adopted in part because monkey L was more likely to give up when the task was difficult, and we felt that a fixed rule might actually disrupt learning. After reaching the full OMP mapping, the experiment continued for another 4.5±3.1 days. At the end of the final day, the intuitive mapping from day 1 was re-introduced for a washout block (425±283 trials).

Multi-day OMP, no incremental training

This section corresponds to Fig. S1E, F, G (upper right). In this type of experiment, on day 1 the intuitive mapping block was followed by a full OMP mapping block for 250 to 815 trials (L: 563±226 trials; N: 440±90 trials). The first day ended with no washout block. The same OMP mapping was presented for 3 to 7 days (L: 4.5±1.5 days; N: 5±0 days). On subsequent days the monkey saw only that experiment's full OMP mapping, and he had at least 400 trials (L: 738±257 trials; N: 443±139 trials) of exposure to it each day. At the end of the final day of each multi-day OMP, no incremental training experiment the intuitive mapping from day 1 was re-introduced for a washout block (699±205 trials).

Single-day OMP, incremental training

This section corresponds to Fig. S1G (lower left). In this type of experiment, the intuitive mapping block was followed by a 40 to 160 trial block with that day's full OMP mapping. Then incremental training began. The monkey saw each incremental mapping for the same number of trials regardless of performance: Incremental step 1 for 40 trials, incremental step 2 for 40 trials, incremental step 3 for 80 trials, incremental step 4 for 80 trials, incremental step 4.5 for 120 or 160 trials, and then the full OMP mapping for at least 160 trials. These block sizes were chosen to balance the number of trials it took for the monkey to improve performance and to minimize the number of rewards the monkey received when the task required less learning (i.e., when the

incremental mapping was closer to the intuitive mapping). Following the full OMP mapping block, the intuitive mapping was re-introduced for a washout block (342±99 trials).

Single-day OMP, no incremental training

This section corresponds to Fig. 1D (blue) and Fig. S1G (upper left). In this type of experiment, the intuitive mapping block was followed by a block with that day's full OMP mapping for 116 to 600 trials (L: 365±85 trials; N: 575±70 trials). Following the OMP mapping block, the intuitive mapping was re-introduced for a washout block (206±83 trials).

Single-day WMP

This section corresponds to Fig. 1D (red). In this type of experiment, the intuitive mapping block was followed by a WMP mapping block for 189 to 800 trials (L: 331±98 trials; N: 616±58 trials). Following the WMP mapping block, the intuitive mapping was re-introduced for a washout block (L: 167±60 trials; N: 334±108 trials).

Analysis

Amount of learning

For each experiment, we computed the amount of learning (*AoL*) during OMP control, which provides a measure of the extent to which learning recovered intuitive levels of performance subsequent to the introduction of the BCI perturbation. The metric is based on reward rate (i.e. the number of successful trials per unit time when the cursor is under the monkey's control) and accounts for the reward rate with the intuitive mapping, as well as the potential learning possible given the difficulty of the perturbed mapping. Reward rate (*RR*) was computed for a sliding window of 40 consecutive trials. Then, we quantified the amount of learning (Fig. 1) as

$$AoL_T = \frac{RR_T^{OMP} - RR_1^{OMP}}{RR^{intuitive} - RR_1^{OMP}} \qquad (9)$$

where $RR^{intuitive}$ is the reward rate calculated for the 40 trials with the intuitive mapping just before the perturbation is introduced, and RR_T^{OMP} is the reward rate with the OMP mapping calculated for a sliding window of 40 trials beginning at trial T after the OMP is introduced. A value of 1 reflects performance that was fully restored to intuitive levels of control, and a value of 0 reflects no performance improvement. This is a different learning metric than was used in Sadtler et al. (*8*), and was chosen because it better captures the learning across all targets when the monkeys had a particularly low success rate to a subset of the targets, which was common for the OMP mappings studied here.

For single day experiments, $AoL = max_T(AoL_T)$. For multi-day experiments, we found AoL for each T on each day. The AoL values reported in the histograms in Fig. 1D are the maximum AoL across all days. Thus, it reflects the best 40 consecutive trials on the best day, to

showcase the learning. The results were qualitatively similar when we considered the best 40 trials from the first day the monkey saw the full OMP mapping after incremental training or if we considered the best 40 trials from the last day the monkey saw that OMP mapping. Because our effects were robust to these choices, this shows that the learning is not trivially dependent on the particular trials we chose.

Neural analyses

Because our goal here was to characterize changes in neural activity that facilitate learning, we focused on the trials that showed the greatest improvement in behavioral performance. For our neural analyses, we analyzed the successful trials from the 40 consecutive perturbation trials that showed the best AoL on each day. Failed trials interspersed within that block of 40 trials were not analyzed because it is difficult to determine whether the animal was actively engaged in the task during failed trials.

Defining new patterns of neural activity with a speed limit

To look for the emergence of new neural activity patterns, we defined a "speed limit". To define the speed limit (Fig. 2), we first projected each intuitive neural activity pattern through the OMP mapping

$$\hat{v}_t = M_2^{OMP} u_t + m_0 \quad (10)$$

Here we are interested in the impact of the neural activity patterns on the cursor kinematics and so do not include the previous timestep's velocity (i.e., $M_1 \hat{v}_{t-1}$ in Eqn. 6). We computed the 95% convex hull of the resulting velocities. This convex hull is a 2D representation of the intuitive neural repertoire (*9*) in the behaviorally-relevant space. The convex hull is the smallest polygon enclosing all the 2D points that also encloses all possible line segments between any two points within the polygon. To be robust to outliers, the outermost points were successively dropped until 5% of all the points had been excluded. Points nearest to the boundary of the hull were dropped in order from largest to smallest Mahalanobis distance from the centroid of all points in the 2D projection. Mahalanobis distances were computed relative to the covariance across all points in 2D.

We refer to the convex hull as a *speed limit* because it represents the maximum speed in each direction that neural activity patterns within the intuitive neural repertoire would have produced under the OMP mapping. By definition, cursor speeds exceeding the speed limit could not have been achieved by the previously-observed neural activity patterns. Therefore, we classified any neural activity patterns leading to cursor speeds exceeding the speed limit as new. This is a conservative definition of new patterns, in the sense that it does not capture patterns of neural activity outside the ~90D neural repertoire that map to speeds *within* the speed limit. The percentage of new patterns (e.g., Fig. 2G) is the ratio of new neural activity patterns to total

neural activity patterns generated by the monkey during that 40-trial block (successful trials only).

In all figures except Fig. S7, we have used the 95% convex hull to define the speed limit and to identify the emergence of new neural activity patterns. To ensure that our results did not depend on the percentage of points contained within the convex hull, we repeated our analyses for 98% and 100% convex hulls (Fig. S7). The results were consistent across these definitions of the speed limit.

Cursor progress

To quantify the behavioral consequence of each neural activity pattern, we measured the component of cursor velocity in the direction of the target at each timestep. We termed this metric *progress*

$$P_t = \begin{bmatrix} \cos \theta_t \\ \sin \theta_t \end{bmatrix} * (M_2^{OMP} u_t + m_0) \qquad (11)$$

which is the projection of the single-timestep cursor velocity onto a unit vector in the cursor-target direction θ_t. This metric is depicted graphically in Fig. 3A. While our *AoL* metric (i.e., reward rate) provides an overall metric of learning, the change in progress provides a metric of the impact of each neural activity pattern on performance. Changes in progress provide a metric of per-target learning.

Outside-manifold contribution to progress

In Figures 4 and S6 we sought to assess whether the new patterns of neural activity that we observed were still subject to the constraints of the intrinsic manifold. Each population activity pattern can be decomposed into an inside- and outside-manifold component. Correspondingly, the observed progress at each timestep is the sum of the inside-manifold contribution to progress and the outside-manifold contribution to progress, as depicted Fig. S6B. We computed the progress that can be attributed to getting outside the intrinsic manifold according to the equation:

$$P_t^{outsideManifold} = P_t^{total} - P_t^{insideManifold} \qquad (12)$$

where P_t^{total} is calculated from Eqn. 11. To calculate $P_t^{insideManifold}$, first we found the orthogonal projection of each neural population activity pattern into the manifold

$$orthogonal\ projection\ into\ manifold = U^T u_t \qquad (13)$$

where u_t is the z-scored spike count vector. Because the spike counts have been z-scored, the mean of u_t is 0. U contains the orthonormalized columns of Λ and is found using the singular value decomposition, such that $\Lambda = USV^T$. The column space of U defines the intrinsic

manifold. Then we found the representation of the orthogonal projection in the high-dimensional space

$$\tilde{u}_t = UU^T u_t \qquad (14)$$

where $\tilde{u}_t$ is the inside-manifold component of the population activity pattern. Finally, we computed the progress from this inside-manifold component of the spike count vector.

$$P_t^{insideManifold} = \begin{bmatrix} \cos\theta_t \\ \sin\theta_t \end{bmatrix} * (M_2^{OMP}\tilde{u}_t + m_0) \qquad (15)$$

Fig. S6C shows the $P_t^{insideManifold}$ and $P_t^{outsideManifold}$ for each of the new neural activity patterns expressed on the last day of OMP1. Some neural activity patterns produce large inside-manifold contributions to progress and small outside-manifold contributions to progress (Fig. S6C, red). Other neural activity patterns produce small inside-manifold contributions to progress and large outside-manifold contributions to progress (Fig. S6C, blue). To further show that learning is not constrained by the intrinsic manifold, in Fig. S6D we calculated the change in the mean $P_t^{insideManifold}$ and mean $P_t^{outsideManifold}$ across all neural activity patterns produced for each target relative to day 1. For this calculation we include all the neural activity patterns expressed for each target, not only the new neural activity patterns. For many targets learning occurs by increasing the $P_t^{outsideManifold}$, thus demonstrating that learning can be driven by producing new neural activity patterns outside the intrinsic manifold. Note that even for some targets that show good learning, the inside-manifold or outside-manifold component can be unhelpful, but then be overcome by helpful progress in the other component (Fig. S6E).

SI References:

37.   Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B Methodol.* **39**: 1–38.

38.   Williamson RC, et al. (2016) Scaling properties of dimensionality reduction for neural populations and network models. *PLoS Comp Bio.* **12**: e1005141.

A

OMP mapping
Incremental mappings
Unit 1 firing rate (FR)
Intuitive mapping
Intrinsic Manifold
Unit 2 FR
Unit 3 FR

B

Multi-day, incremental training

Monkey N
Monkey L

Multi-day, no incremental training

Single-day, incremental training

Single-day, no incremental training

Time ⟶

Legend
Intuitive mapping
Step 1
Step 2
Step 3
Step 4
Step 4.5
OMP mapping

C Multi-day OMP, incremental training: monkey L
n=9
Amount of Learning
Day

D Multi-day OMP, incremental training: monkey N
n=6
Amount of Learning
OMP 1
Day

E Multi-day OMP, no incremental training: monkey L
n=6
Amount of Learning
Day

F Multi-day OMP, no incremental training: monkey N
n=5
Amount of Learning
Day

G

Single-day OMP

Multi-day OMP

No incremental training
n=21
% of experiments

n=11
% of experiments

With incremental training
n=27
% of experiments
Amount of learning

n=15
% of experiments
Amount of learning

12

Fig. S1. Incremental training facilitates multi-day learning. (**A**) An incremental training approach was used to guide monkeys during learning of the OMP mappings. Incremental mappings spanned the neural space between the intuitive mapping and the OMP mapping. (**B**) Training scheme used for single- and multi-day experiments with incremental training and no incremental training. Each block indicates a day. (**C**) Learning curves for all multi-day OMP, incremental training experiments for monkey L. The amount of learning (circles) is shown only for days on which the full OMP was presented to the animal. Each line indicates a unique OMP mapping. (**D**) All multi-day OMP, incremental training learning curves for monkey N. Same conventions as in C. OMP 1 corresponds to the experiment highlighted in other figures. (**E**) All multi-day OMP, no incremental training learning curves for monkey L. (**F**) All multi-day OMP, no incremental training learning curves for monkey N. (**G**) The amount of learning for multi-day, incremental training OMP experiments (green) was greater than for all other training strategies we used (vs. single-day, incremental training OMP experiments (cyan; t-test, p=0.0019); vs. multi-day, no incremental training OMP experiments (light green; t-test, p=0.052); vs. single-day no incremental training (blue; t-test, $p<10^{-4}$).
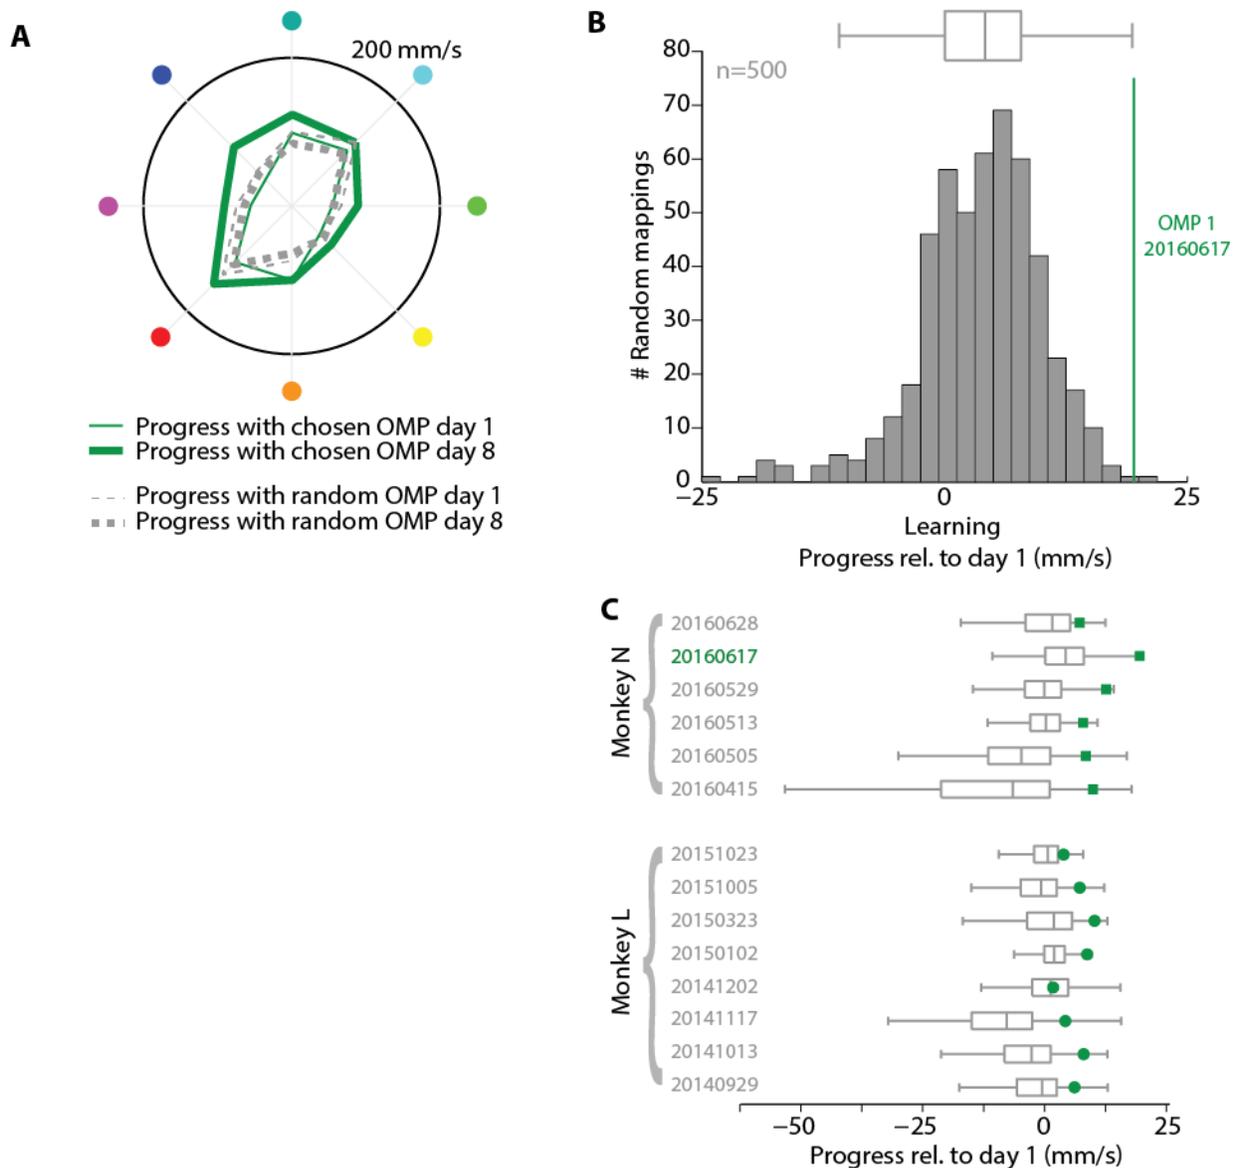
Fig. S2. Changes in neural activity are specific to the OMP mapping presented to the monkey. (**A**) Mean progress (cf. Fig. 3) observed through the OMP 1 mapping increased from day 1 to day 8 (green). For comparison, we mapped the learned neural activity patterns through 500 random OMP mappings and measured progress. The gray dashed lines show the average progress through these random mappings. Progress does not improve through these random mappings over time. This shows that neural activity changed in a manner that is appropriate for the particular OMP mapping the monkey was learning. (**B**) The change in progress between day 1 and the last day (averaged over all targets) is plotted for the OMP mapping the animal experienced (green) and each of the 500 random OMP mappings (gray). (**C**) For each multi-day OMP, incremental training experiment we compared the observed learning (change in progress relative to day 1 averaged across targets, green dot) to the apparent learning distribution through 500 random mappings (box and whisker plot, where the center line indicates the median, the edges of the box indicate the 25th and 75th percentiles, and the whiskers extended to the smallest

and largest of the data points). For all but one experiment, the chosen mapping showed more learning than would have occurred for 75% of the random mappings.
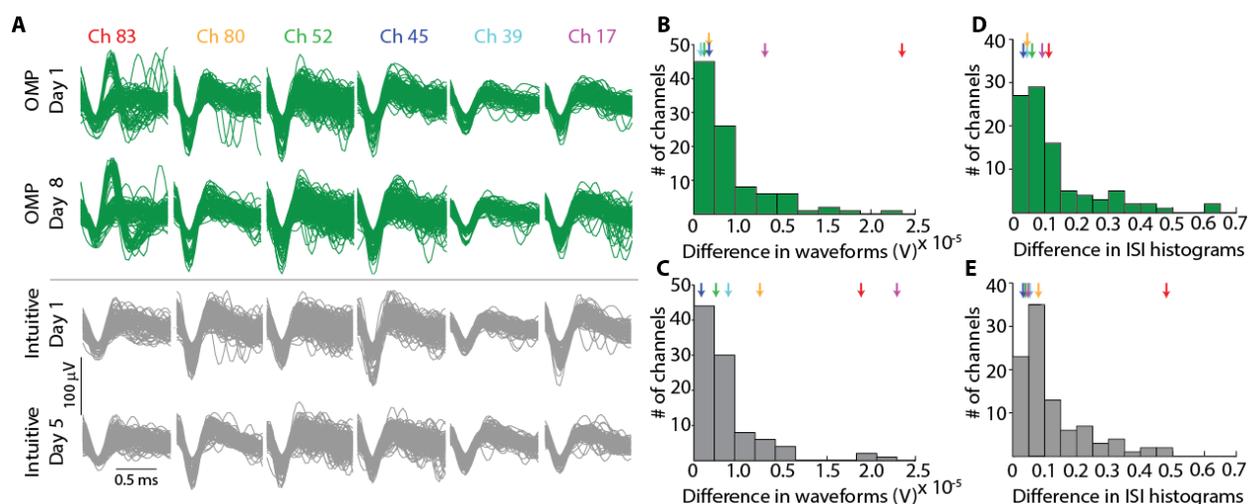
Fig. S3. Recordings are stable on the timescale of several days. Here, stability is assessed during an example multi-day OMP learning experiment and a multi-day intuitive mapping experiment. (**A**) Waveforms of the threshold crossings used for decoding from six selected channels on day 1 and the last day (day 8) of OMP1 for Monkey N (green) and day 1 and the last day (day 5) of a multi-day intuitive experiment for Monkey N (gray). The waveform shape for each channel remains stable across days. (**B**) Waveform differences for OMP1. We quantified the difference in waveform shape by calculating the distance between the mean waveform on day 1 and that on the last day for each channel. Arrows correspond to the selected channels in panel A. (**C**) Waveform differences for a multi-day intuitive experiment. Arrows correspond to the selected channels in panel A. The waveform shape changes are not statistically different between the multi-day OMP learning and multi-day intuitive scenarios (K-S test, p=0.97). (**D**) ISI differences for OMP1. The difference in ISI histograms was calculated as the K-S statistic between the distribution of ISIs on day 1 and the last day. Arrows correspond to the selected channels in panel A. (**E**) ISI differences for a multi-day intuitive experiment. Arrows correspond to the selected channels in panel A. ISI changes are not statistically different between the multi-day OMP learning and multi-day intuitive scenarios (K-S test, p=0.99).

A Multi-day OMP: high learning pressure

OMP day 1 — 6% new, Δ prog = 2.3
OMP day 6 — 10% new, Δ prog = 11.3
OMP day 7 — 22% new, Δ prog = 18.9
OMP day 8 — 28% new, Δ prog = 20.2
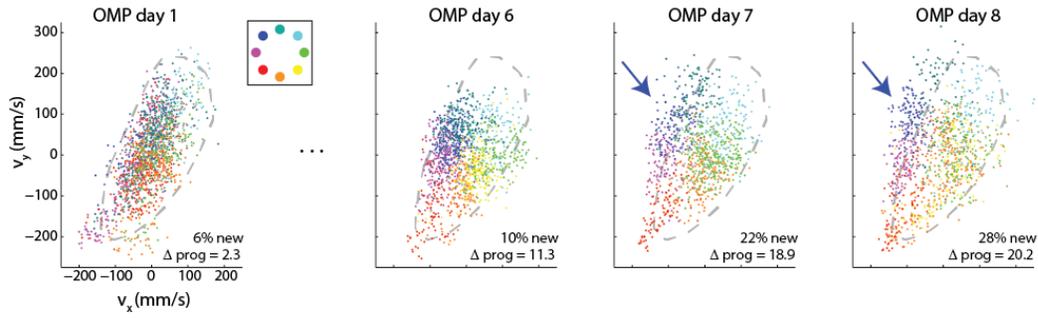
B Multi-day intuitive mapping: "low learning pressure"

Intuitive day 1 — 5% new, Δ prog = xx
Intuitive day 2 — 8% new, Δ prog = -1.1
Intuitive day 3 — 11% new, Δ prog = -1.1
Intuitive day 4 — 10% new, Δ prog = -3.8
Intuitive day 5 — 13% new, Δ prog = -3.4
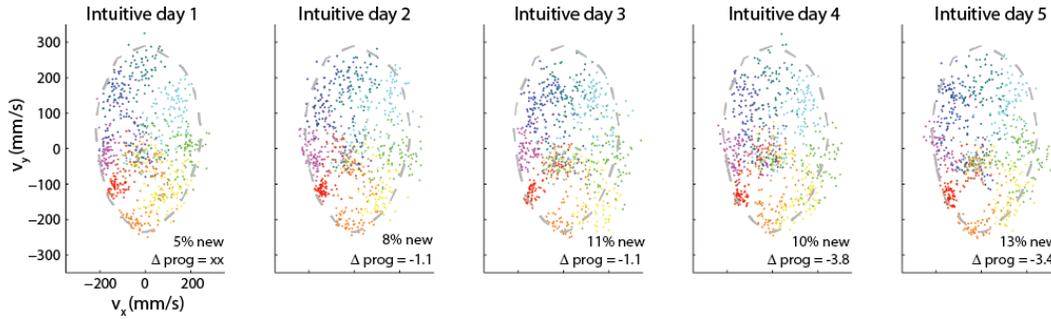
C Multi-day intuitive neural activity through random OMP: "no learning pressure"

Intuitive day 1 — 4% new, Δ prog = 2.26
Intuitive day 2 — 10% new, Δ prog = 2.37
Intuitive day 3 — 10% new, Δ prog = 6.23
Intuitive day 4 — 9% new, Δ prog = 7.38
Intuitive day 5 — 13% new, Δ prog = 3.97

D

17
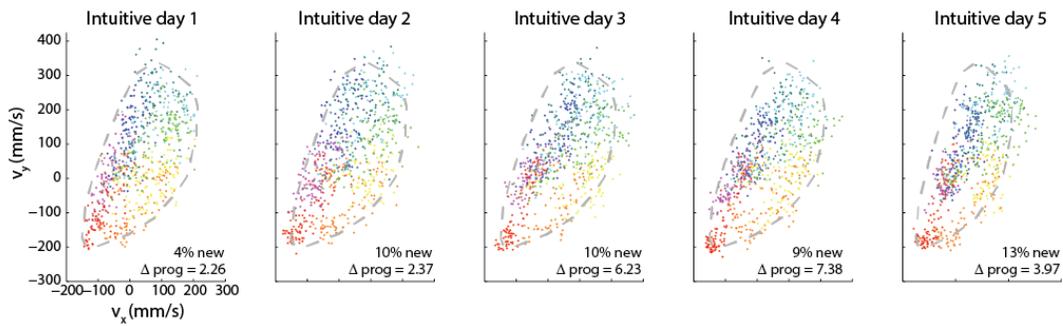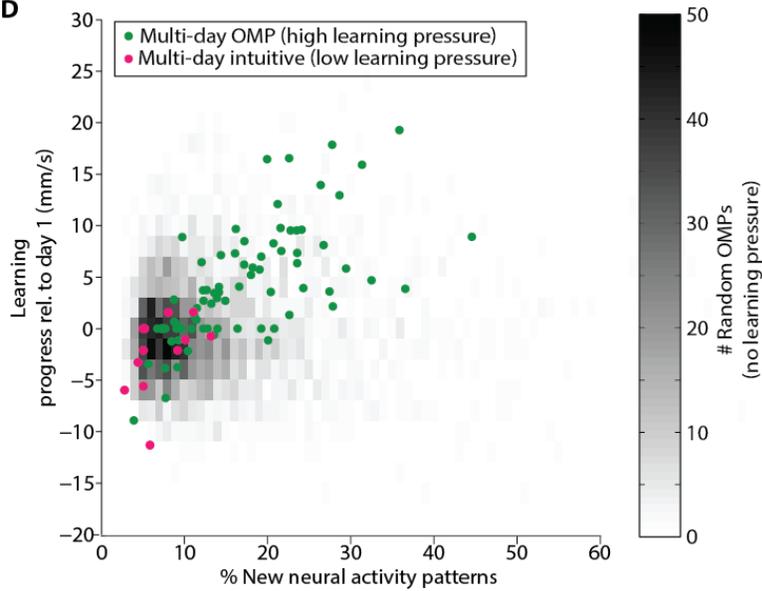
Fig. S4. New neural activity patterns preferentially emerge during OMP learning. (**A-C**) "New" neural activity patterns correspond to velocities that exceed the speed limit. As in Fig. 2, the "speed limit" is defined as the 95% convex hull of velocities generated from the intuitive neural repertoire projected through the OMP mapping (gray dashed line). The velocities generated while using the OMP mapping are colored according to the instructed target location. Each dot is the velocity resulting from one neural activity pattern (45 ms bin). (**A**) Multi-day OMP learning for OMP 1. This experiment is the same as in Fig. 2, but here we include all days for which the monkey used the full OMP mapping. On day 1, the velocities during OMP control mostly fell within the speed limit defined from the intuitive neural repertoire. On day 6, most of the velocities are still within the speed limit. However, the velocities have become more appropriate for successful target acquisition, as can be seen by the points for each target (color code given by inset). For example, when the monkey is instructed to move the cursor to the upper-left target (blue), he generates neural activity patterns that cause velocities in that direction. On days 7 and 8, velocities begin to consistently exceed the speed limit in the direction that is most beneficial for behavior (blue arrows). Thus, during OMP learning, new neural activity patterns emerge in a manner that increases the progress of the cursor toward the target. % new indicates the percentage of points that fall outside the speed limit. Δ prog refers to the progress relative to day 1. (**B**) Multi-day intuitive mapping. This is considered a "low learning pressure" scenario because while cursor control performance is high from the beginning of the experiment, there might still be motivation for the monkey to increase his reward rate by generating faster speeds. The speed limit (gray dashed outline) is identified on day 1. Each panel shows the velocities generated during subsequent days with the same intuitive mapping relative to the speed limit. Across all days there are only small changes in progress (which happen to all be negative, i.e., behavior has gotten worse relative to day 1) and few new neural activity patterns (that is, velocities which exceed the speed limit). The velocities that do exceed the speed limit do not occur in an appropriately target-directed manner, as was seen during OMP learning in row A. (**C**) For a "no learning pressure" scenario, we passed the neural activity from the multi-day intuitive experiment through 500 random OMP mappings. There is no learning pressure because the monkey did not receive any online feedback about these random OMP mappings. The speed limit (gray dashed line) and velocities corresponds to one representative random OMP mapping which the monkey never saw. In this scenario, it is not the emergence of new neural activity patterns that leads to apparent behavioral improvement. Rather, there is a slight shift of the velocities relative to the speed limit that leads to an increase in speeds in some directions and a decrease in speeds in other directions. (**D**) New neural activity patterns preferentially emerge during OMP learning. Increases in progress are plotted as a function of the percentage of new neural activity patterns for the three conditions of high- (green), low- (magenta), and no-learning pressure (density histogram across 500 random OMP mappings). There is significant correlation between the formation of new neural activity patterns and increases in progress for the high-learning condition (Pearson correlation coefficient r=0.67, p=$1*10^{-11}$).
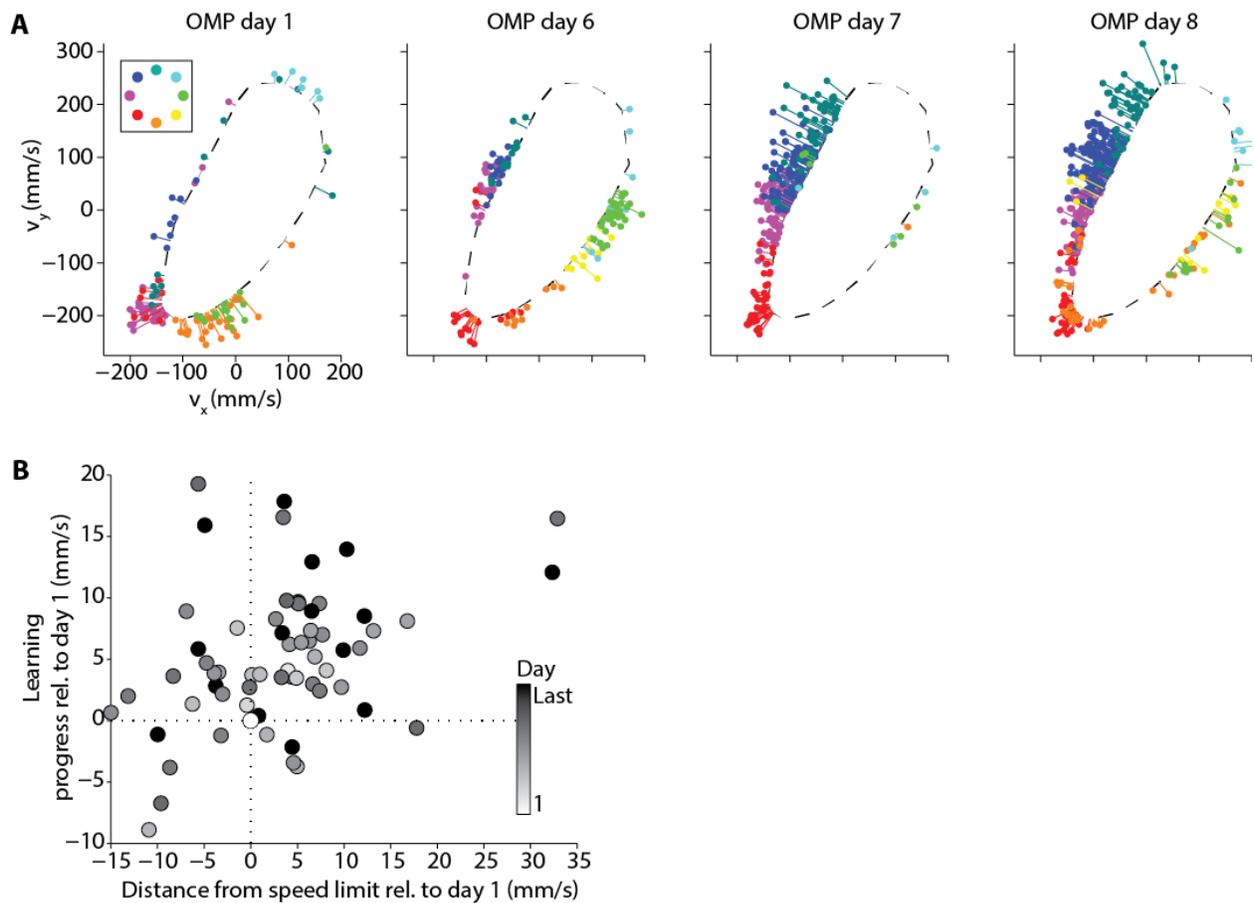
Fig. S5. New neural activity patterns move progressively farther away from the speed limit with learning. In Fig. 2, we defined the speed limit as the 95% convex hull of the velocities generated when the intuitive neural repertoire is mapped through the OMP. We defined new neural activity patterns as patterns that generate velocities which exceed the speed limit. With learning, we observe more new neural activity patterns. One possible interpretation of this result is that the velocities that exceed the speed limit arise from visiting those 5% of points outside the speed limit progressively more often. If this were case, the percentage of "new" neural activity patterns identified by the speed limit metric might go up, but the distance of those points from the speed limit would remain unchanged. Here, we show that this possibility is not consistent with the data. (**A**) Instead, the points move progressively farther from the speed limit (gray dashed) during learning (i.e., the cursor velocities continue to get faster and more target-directed during learning). (**B**) With learning over days, neural activity patterns are further from the speed limit, and they contribute to greater progress. For the multi-day OMP experiments, there is a significant correlation between the extent of the points beyond the speed limit and increases in progress (Pearson correlation coefficient r=0.40, p=0.0002). Each symbol is the average over all 8 targets from one day of one OMP experiment relative to day 1. Symbols are shaded to indicate the day within a given multi-day experiment.
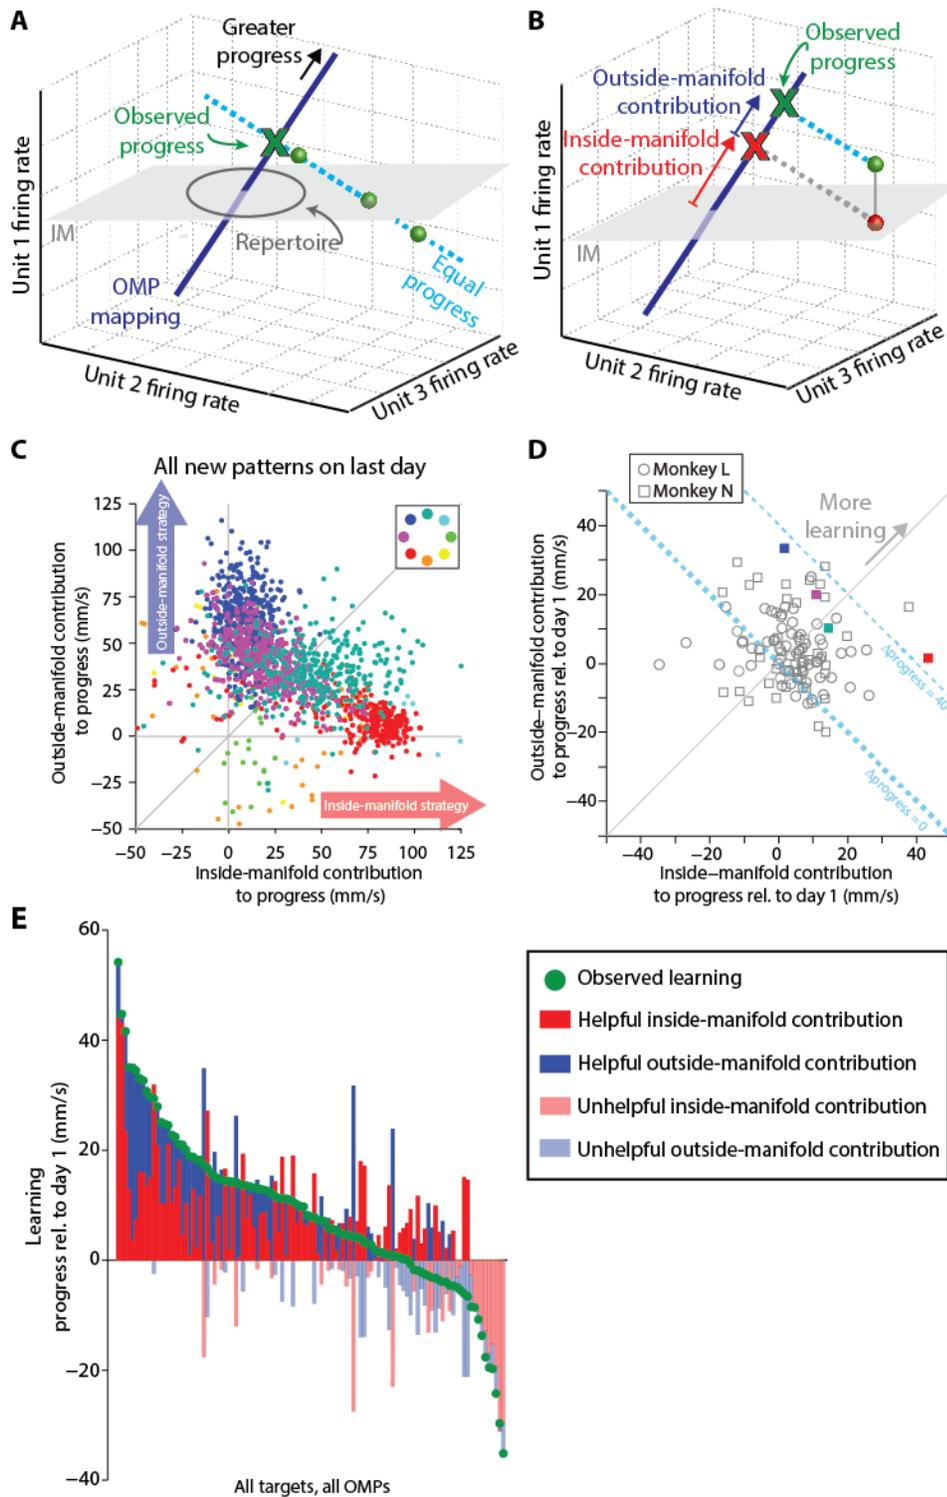
Fig. S6. New patterns of neural activity can be inside or outside of the intrinsic manifold. (**A**) A particular value of cursor progress can be produced by many different neural activity patterns (dashed cyan line), including both patterns inside and outside of the manifold. (**B**) Each neural

activity pattern (green dot) can be decomposed into an inside-manifold component (red dot) and outside-manifold component. Correspondingly, observed progress (green X) is the sum of inside-manifold contribution to progress (red arrow) and outside-manifold contribution to progress (blue arrow). (**C**) Animals learn using both inside-manifold and outside-manifold strategies for a given OMP mapping. Each dot shows a new (i.e., outside-repertoire) neural activity pattern observed on the last day of OMP 1. For each new neural activity pattern, the inside-manifold component is plotted against the outside-manifold component. Dots are colored according to the instructed target when that pattern was generated by the monkey (inset). Dots above the diagonal have largely outside-manifold contributions to progress. Dots below the diagonal have largely inside-manifold contributions to progress. We found that some targets (e.g., red) are learned using a largely inside-manifold strategy, whereas other targets (e.g., blue) are learned using a largely outside-manifold strategy. Most of the targets (e.g., teal) are learned using a mixed strategy. (**D**) Learning (i.e., change in progress relative to day 1) can arise from both outside-manifold and inside-manifold strategies for different targets. Each symbol shows the extent of inside- and outside-manifold contributions to learning for one target from one experiment. Each point is the average of all of the same-colored dots plotted in the space of panel C. Here all population activity patterns are included in the average, not just the new patterns. Colored symbols correspond to the targets represented in panel C. Equal Δ progress lines are plotted to emphasize that both inside- and outside-manifold strategies were used by the animals to achieve similar amounts of learning. (**E**) Animals learn using both inside-manifold and outside-manifold strategies. Each bar shows one target from one multi-day OMP experiment. This figure differs from Fig. 4C in that here all targets are shown, including those that show negative learning. The overall learning, defined as change in progress from day 1 to the last day, is represented by a green dot. The inside-manifold contributions to that learning are shown in red. The outside-manifold contributions are shown in blue. The overall learning is equal to the sum of the inside-manifold contribution (whether it be positive, i.e., helpful, or negative, i.e., unhelpful) and the outside-manifold contribution (whether it be positive or negative). Thus, it is possible for a bar to exceed the corresponding green dot if the inside-manifold and outside-manifold contributions take on different signs.

A 95% convex hull
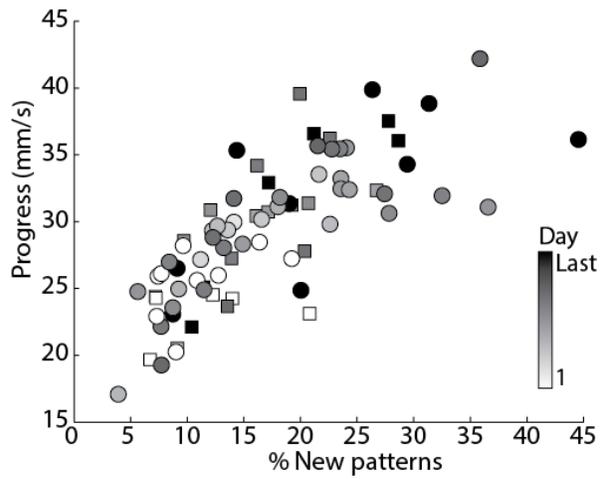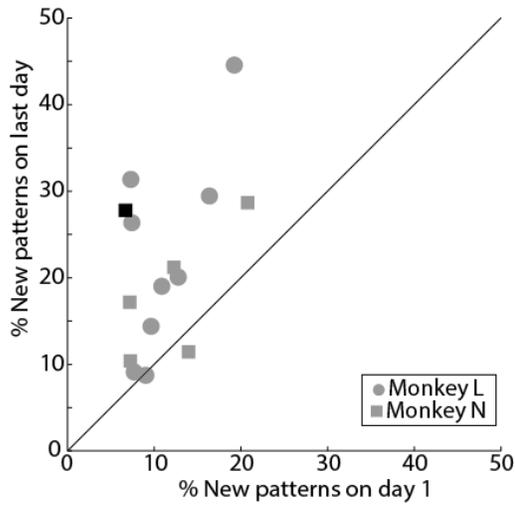
B 98% convex hull
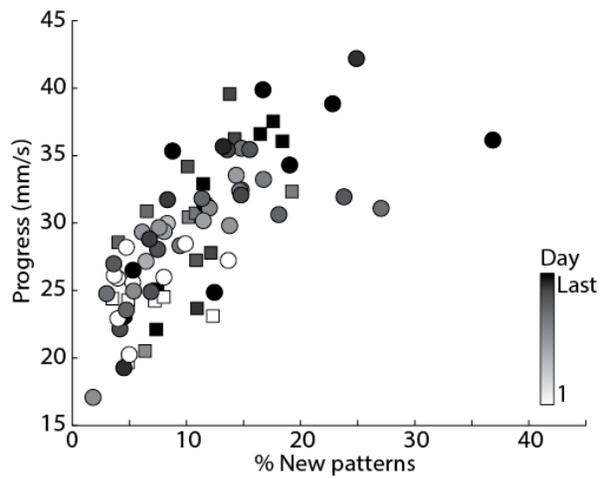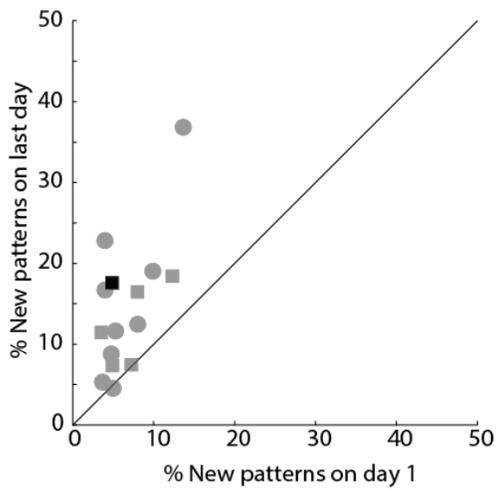
C 100% convex hull
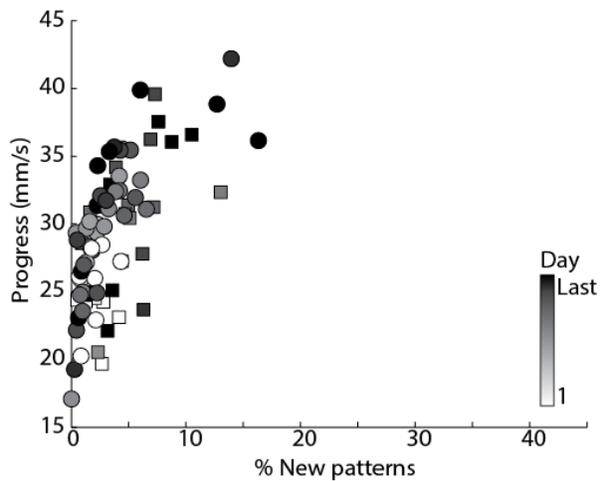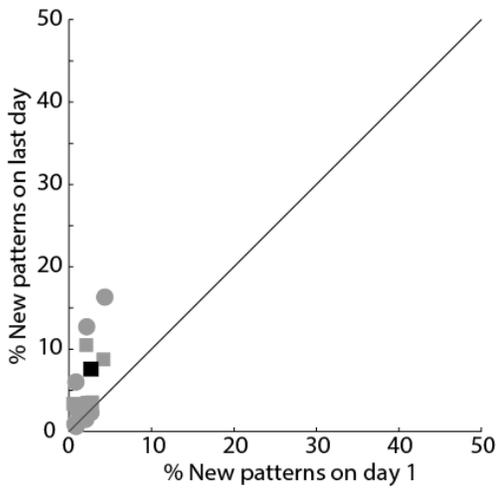
Fig. S7. Results do not depend on the specific definition of the speed limit. In all other figures, we define the speed limit as the 95% convex hull. Here we have repeated the key analyses that depend on the speed limit (Fig. 2G and Fig. 3C) for speed limits defined as the (**A**) 95% convex hull (reproduced here to facilitate comparison), (**B**) 98% convex hull and (**C**) 100% convex hull.

Monkey	Multi-day OMP, incremental training	Multi-day OMP, no incremental training	Single-day OMP, incremental training	Single-day OMP, no incremental training	Single-day WMP
L	9 (7 to 16 days)	6 (3 to 7 days)	12	13	11
N	6 (6 to 11 days)	5 (5 days)	15	8	11

Table S1. Number of experiments of each type.

Movie S1. Example BCI trials to each target with the intuitive mapping on day 1 of a multi-day, incremental training OMP experiment (OMP 1), and the OMP mapping on day 1 and after 8 days of training.