Neuron
Article

Abrupt Transitions between Prefrontal Neural Ensemble States Accompany Behavioral Transitions during Rule Learning

Daniel Durstewitz,[1,4,*] Nicole M. Vittoz,[2,4] Stan B. Floresco,[3] and Jeremy K. Seamans[2,*]
[1]RG Computational Neuroscience, Central Institute of Mental Health and Interdisciplinary Center for Neurosciences, University of Heidelberg, J 5, 68159 Mannheim, Germany
[2]Brain Research Centre, Psychiatry, Faculty of Medicine, University of British Columbia, Vancouver, BC V6T 2B5, Canada
[3]Psychology Department, University of British Columbia, Vancouver, BC V6T 2B5, Canada
[4]These authors contributed equally to this work
*Correspondence: daniel.durstewitz@zi-mannheim.de (D.D.), seamans@interchange.ubc.ca (J.K.S.)
DOI 10.1016/j.neuron.2010.03.029

SUMMARY

One of the most intriguing aspects of adaptive behavior involves the inference of regularities and rules in ever-changing environments. Rules are often deduced through evidence-based learning which relies on the prefrontal cortex (PFC). This is a highly dynamic process, evolving trial by trial and therefore may not be adequately captured by averaging single-unit responses over numerous repetitions. Here, we employed advanced statistical techniques to visualize the trajectories of ensembles of simultaneously recorded medial PFC neurons on a trial-by-trial basis as rats deduced a novel rule in a set-shifting task. Neural populations formed clearly distinct and lasting representations of familiar and novel rules by entering unique network states. During rule acquisition, the recorded ensembles often exhibited abrupt transitions, rather than evolving continuously, in tight temporal relation to behavioral performance shifts. These results support the idea that rule learning is an evidence-based decision process, perhaps accompanied by moments of sudden insight.

INTRODUCTION

We are constantly faced with changes in our daily lives and are often forced to realize that behavioral strategies that were once appropriate may now be highly disadvantageous. In these situations, novel response strategies must be developed and old ones abandoned. Often, however, environmental situations are complex with many unknowns such that underlying regularities must be probed through trial and error or evidence-driven deductive search. Most species are capable of this type of trial-and-error learning, which may either proceed in a slow incremental fashion or be accompanied by an "a-ha" moment in which the problem is solved all at once through "sudden insight" (Aziz-Zadeh et al., 2009; Bowden et al., 2005; Gallistel et al., 2004).

One task that classically assesses trial-and-error learning in humans is the Wisconsin card-sorting test (WCST). This task intermittently requires subjects to deduce a new sorting strategy based on feedback from the experimenter. The ability to switch to a different sorting rule (an extra-dimensional shift) is critically dependent on an intact prefrontal cortex (PFC) (Milner, 1963). Furthermore, the dorsolateral and medial PFC (mPFC) are consistently activated on this and a variety of other trial-and-error learning tasks (Boettiger and D'Esposito, 2005; Landmann et al., 2007; Nakahara et al., 2002; Sandkühler and Bhattacharya, 2008). Electrophysiology studies in animals employing analogous behavioral tasks have also confirmed the involvement of the PFC at the single-cell level. Neurons within the mPFC exhibit robust error-related correlates (Mansouri et al., 2006, 2007) and cells in both the dorsolateral PFC and mPFC also show task-specific activity as monkeys learn novel touch sequences or mappings through trial and error or instructional cues (Genovesio et al., 2005; Procyk et al., 2000). Moreover, subgroups of neurons in both PFC regions show greater activity on rule-search trials than on trials in which the monkey is simply repeating a known sequence. In rats, mPFC activity changes as they switch task strategies even if both strategies require identical behaviors (Rich and Shapiro, 2009).

While activity in the mPFC is tightly correlated with various aspects of the trial-and-error learning process, it is unclear how the neural dynamics unfold in time as the animal progresses trial-by-trial from a familiar to a novel rule. Is it the case that the network moves along a more-or-less linear, incremental course as the animal gradually abandons an old strategy, slowly begins to relearn, and finally responds according to the new rule? Or, as suggested by Gallistel et al. (2004), is there a moment at which an abrupt transition occurs, perhaps accompanied by a "sudden insight," when sufficient evidence for a new environmental contingency has accrued? These questions are difficult to answer with the traditional method of averaging single-unit responses over tens to hundreds of trials. Instead, it is necessary to track neural behavior on a trial-to-trial basis. With the advance of multiple single-unit recording techniques and appropriate multivariate statistical methods (Chapin and Nicolelis, 1999;

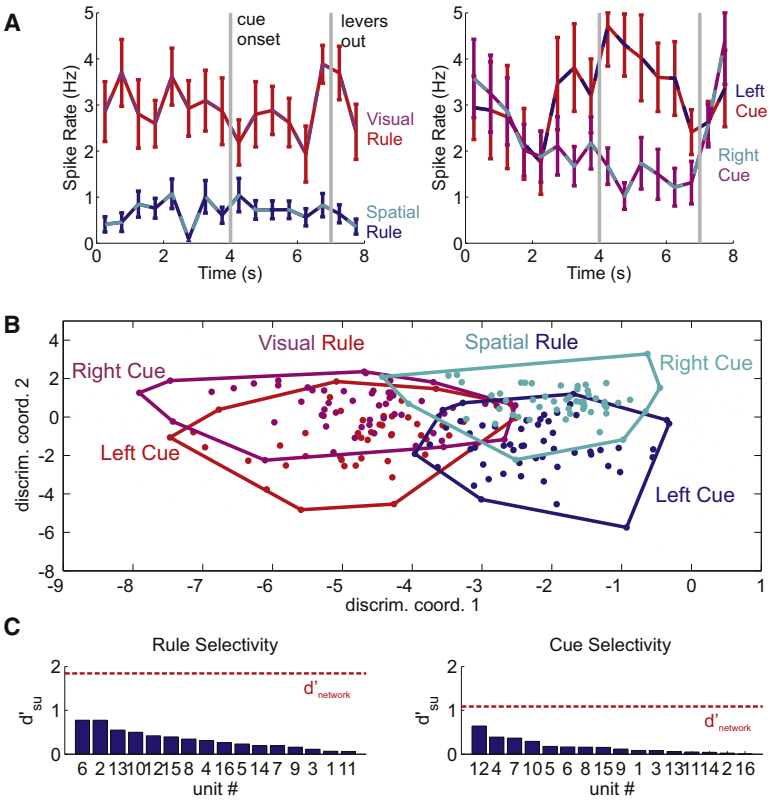

Figure 1. Single-Unit and Population Discrimination among Task Rules and Cues

(A) Examples of single units which show significant discrimination among the two task rules (left) or between the two cue lights (right). Examples are the units with the highest rule and cue selectivity indices, respectively, from the data set shown in (B) and (C). Note that while cue selectivity is transient, rule selectivity for this cell is maintained throughout the whole 8 s analysis period. Error bars = SEM.

(B) Projection of the multiple single-unit activity (MSUA) (total of 16 units) onto the two most discriminating dimensions using DA (see text). Dots in bluish (spatial rule) and reddish (visual rule) colors mark activity vectors belonging to one of the two rule sets, while darker colors (red, blue) mark activity vectors corresponding to left cue light and lighter colors (magenta, cyan) those corresponding to right cue light trials. The polygons show the *convex hull* of the cue and rule representations, i.e., their *largest* extent in the data.

(C) Selectivity indices for all single units, compared to the selectivity index of the whole network (red line), for discrimination among rule sets (left) and among cues (right).

Deadwyler and Hampson, 1997; Brown et al., 2004) the extraction of network-level information from single trials has become more feasible (Churchland et al., 2007; Jones et al., 2007; Lapish et al., 2008; Yu et al., 2009). Here, we used this approach to visualize the dynamics of an ensemble of simultaneously recorded mPFC neurons as rats shifted away from a familiar cue-based response strategy in an operant chamber and began to acquire an egocentric response strategy by gathering evidence through trial and error. This task is considered a rodent analog of the WCST and is severely disrupted by inactivation of the rat mPFC (Floresco et al., 2008; Ragozzino et al., 1999).

RESULTS

Neural Population Representation of Rules and Cues

Thirteen animals were first trained to perform a simple visual cue-based discrimination task (termed "visual rule" in the following; Figure S1A). After reaching criterion, on one day, unknown to the animal, the reward contingencies were changed after 20 trials on the visual rule such that only responses to the left or the right lever were rewarded, regardless of the cue light location (termed "spatial rule" in the following). To study the neural representation of these two rules, we first focused on those last 20 trials of the visual rule right before the shift to the spatial rule, and the last 20 trials of the spatial rule after it had become established. The behavioral error rates during these trials were low (<12% on average), and they will be subsequently referred to as the visual rule and spatial rule "steady states," respectively. Individual spike trains of all isolated units were convolved with Gaussian kernels and converted into a time series of instantaneous firing rates for each unit i as a function of time bin t, $r_i(t)$, with bin width Δt = 500 ms (we checked a range of binning and kernel parameters; see Experimental Procedures). All of the N simultaneously recorded cells from each individual rat were then combined into population vectors $\mathbf{r}(t)$ evolving as a function of time bin. Rate vectors were aligned to the start of each trial.

There were a number of individual cells significantly discriminating in their average firing activity between the two behavioral rules (>30% according to conservative criteria; see Supplemental Experimental Procedures available online; Figure 1A, left panel) or between the two cue lights (>8%; Figure 1A, right panel) denoting the correct response during the visual rule. Interestingly, rule-selective cells often maintained (in ~1/3 of the cases) their differential firing rates throughout the examined time window (8 s), including 4 s of *precue* time, while cue-selective cells limited their selective rate changes to comparatively short periods around the presentation of the cues (a phenomenon also confirmed at the population level; Figure S2). The selectivity of single units for a particular rule or cue was evaluated during the 3 s postcue period indicated in Figure 1A using the sensitivity parameter $d'_i = |\langle r_{i(1)}\rangle - \langle r_{i(2)}\rangle| / \mathrm{sqrt}(\sigma_{i(1)}^2 + \sigma_{i(2)}^2)$ as introduced in signal detection theory. In words, these values take the (absolute) difference between the mean firing rates (denoted by $\langle\cdot\rangle$) associated with two sets of trials, divided by the square root of the sum of their variances. The sets $r_{i(1)}$ and $r_{i(2)}$ correspond either to steady-state trials with left or right cue light presentations (for cue selectivity) or to the steady-state trials on which the visual or spatial rule was the correct strategy (for rule selectivity), respectively. The larger the mean difference between the two sets of firing rates and the smaller their variance, the greater is the sensitivity of that unit for the difference between rules/ cues. (In probabilistic terms, a value of $d' > 2$ implies a misclassification rate of less than 8% based on normality

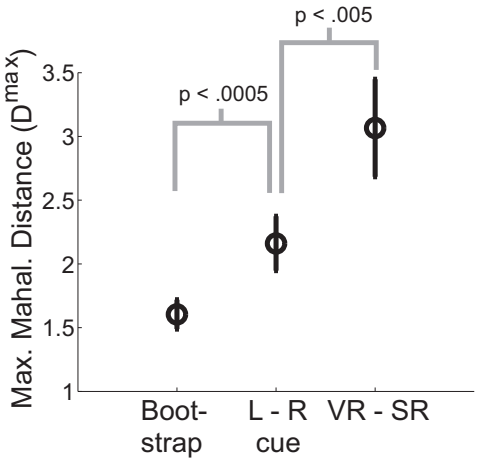

Figure 2. Summary Statistics for Rule and Cue Discrimination
Summary statistics for discrimination among rules and cues as determined from the maximum (along trajectories) Mahalanobis distances among the respective sets of points, averaged across all 13 data sets. L = left, R = right, VR = visual rule, SR = spatial rule. Error bars = SEM. See also Figure S2.

assumptions.) Distributions of these selectivity indices are shown in Figure 1C for discrimination among rules (left panel) and cues (right panel) for one of the 13 data sets collected.

To visualize the discrimination among rules and cues at the population level, Fisher's discriminant criterion (e.g., Krzanowski, 2000) was used to derive a two-dimensional plane from the *N*-dimensional vector time series **r**(*t*) (Figure 1B). Discriminant analysis (DA) is a standard statistical procedure similar to principal component analysis, with the crucial difference that the axes of the reduced space are determined such that the differences between the group means are maximized while the within-group scatter is minimized (i.e., according to a criterion which is a multidimensional extension of *d′* above, the Mahalanobis distance among groups). Hence, the two axes shown in Figure 1B are the ones which most clearly bring out differences in PFC network activity between the different conditions. Each of the new axes represents a linear combination of the firing rates of the recorded units. As exemplified in Figure 1B and confirmed statistically further below, in most of the data sets ($\geq$9/13) the two rules (bluish and reddish clusters) could clearly be separated, and often the two cues (darker versus lighter colors) as well. To compare the selectivity of single units to the discriminative power of the network as a whole, we also calculated *d′* for each recorded network on the most discriminating DA axis with regards to separation among rules (dashed line in Figure 1C, left) or cues (Figure 1C, right). This analysis demonstrates that several units combine their discriminative power to give rise to the network level performance. Thus, the two rules and cues are associated with four separable and coherent clusters of firing rate patterns across the recorded population.

To evaluate the statistical significance of these observations across all 13 data sets obtained, for each data set we computed the Mahalanobis distances (the criterion maximized by DA; see Supplemental Information and Figure S2A for the rationale behind this approach; Krzanowski, 2000) (a) between the *N*-dimensional vectors associated with the two steady-state rule sets, (b) between the sets of population vectors associated with the two different cues across both rules, and (c) as a control, between two sets of vectors where each of the correctly performed steady-state trials was randomly assigned to one of the two cue conditions. Hence, for each of the 13 data sets we obtained three numbers for the conditions a–c defined above, and we tested the specific hypotheses a > b, a > c, and b > c through paired t tests. As shown in Figure 2, across all 13 data sets the Mahalanobis distances between population vectors associated with the two cues were significantly larger ($t(12) > 4.83$, $p < 3 \times 10^{-4}$) than when trials were assigned randomly (b > c). Furthermore, the differences among the two rule sets were significantly larger than those between either cue sets ($a > b$; $t(12) > 3.4$, $p < 3 \times 10^{-3}$) or between the random control sets ($a > c$; $t(12) > 4.87$, $p < 2\text{x}10^{-4}$), confirming our visual inspection of the data.

Transition Trials: Correlation between Neural States and Behavioral Choices

Next, we examined the neural population dynamics specifically during the transition period from the well-learned visual rule to the novel and conflicting spatial rule. It should be noted that there were no extrinsic indications other than the changes in response-reward contingencies that would inform the animal about the change in rules. Figure 3A replots the example from Figure 1B using two bivariate normal distributions with parameters estimated from the data to represent the two rule steady states. In addition it plots the (smoothed) trajectory which represents the temporal evolution of neural population activity within this two-dimensional plane as it moves from one to the other rule steady state (color-coded according to trial number). To relate the neural transition dynamics to behavioral performance, we focused specifically on those transition trials where the visual and spatial rules were in conflict, e.g., when the rat had to press the left lever according to the spatial rule, but the cue light would have indicated a right lever response under the visual rule. Across all of these trials for each animal, the level of agreement between the behavioral and the neural choices was computed. Denoting the difference between the Mahalanobis distances to the visual and spatial rules on trial *i* by $S_i = D_i^{(visual)} - D_i^{(spatial)}$, neural choices were defined as the sign of the z-transformed values S_i. Thus, a "−1" would indicate a neural preference for the visual rule and a "+1" for the spatial rule. Figure 3B (blue bars) plots for all data sets the percentage of trials where the neural and behavioral choices agree (with 50% being chance level). Across all data sets ($t(12) > 3.58$, $p < 0.005$), and for 7 out of the 13 data sets evaluated individually (based on the binomial distribution with $p < 0.05$), the level of agreement was significantly above chance level, despite the fact that during a large proportion of the transition period the animal may not be aware of the correct strategy and may thus be expected to respond rather randomly. Moreover, as demonstrated in Figure 3C, the percentage of correctly predicted behavioral choices increased steadily as the trajectory approached one of the two rule steady states, i.e., the accuracy of prediction is directly related to the degree to which the neural dynamics indicates one or the other choice.

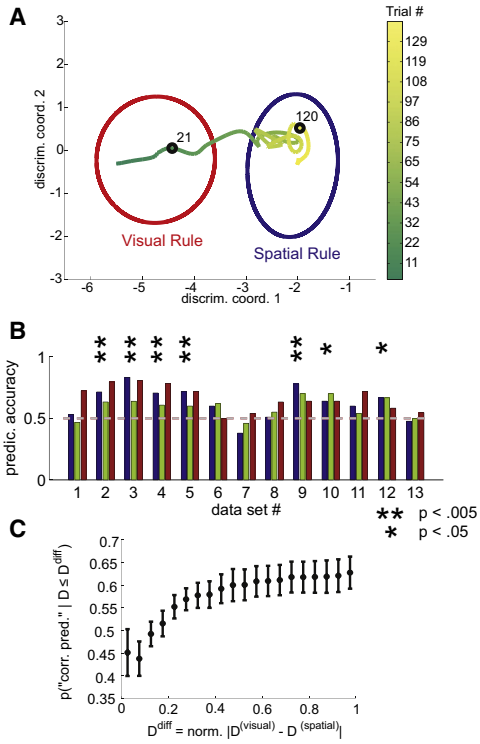

Figure 3. Transition between the Visual and the Spatial Rule

(A) The same MSUA space projection and data as in Figure 1B are shown with the neural trajectory (smoothed with a Gaussian kernel with SD = 3 trials) connecting the two rule steady states during the transition phase. The two rule steady states are represented by 1 SD contours of bivariate normal densities fitted to the distributions in Figure 1B. The trajectory is color-coded according to trials (see legend), and the first (#21) and last (#120) transition trials (right after or before the steady states) are indicated.

(B) Prediction of behavioral choices during the transition phase. Relative number of correct predictions of the behavioral choices based on the neural rule preference for each of the 13 data sets (blue bars), evaluated on those trials where visual and spatial rules were in conflict (chance level = 0.5). The seven data sets marked by stars exceeded chance performance individually when evaluated through the binomial distribution. Green bars indicate behavioral prediction accuracy evaluated only from a 3 s period preceding cue onset by 1 s (i.e., between-trials periods), while brown bars are the same for the 3 s period starting with cue onset (i.e., within-trial periods). Prediction accuracy is significantly better for within-trial periods ($t(12) > 2.18$, $p < 0.05$; where also 9/13 individual comparisons became significant), yet behavioral choices can also be predicted with beyond chance accuracy from between-trial periods ($t(12) > 4$, $p < 0.001$, for averages, and for 4/13 individual data sets).

(C) Prediction accuracy as a function of the distance to one of the two rule steady states. The abscissa gives the normalized relative proximity to one of the two rule states (the higher D^{diff}, the closer the neural trajectory is to one and the farther from the other of the two rule states). The ordinate shows prediction performance cumulated for relative distance values up to the ones given on the abscissa. Error bars = SEM.

Sudden Transitions of the Neural Ensemble Dynamics during the Rule Switch

We next addressed quantitatively the behavior of the neural trajectories during the transition period. When examining the neural choice criterion as defined above on a trial-by-trial basis, i.e., the difference between visual and spatial rule distances, we noticed that in many data sets there were quite steep transitions within the trial time series of this measure (Figure 4A). Sometimes the neural dynamic seemed to jump directly from the visual to the spatial rule (Figure 4A, left), while on other occasions there appeared to be an intermediate state of "indecisiveness" intervening between the two rule states (Figure 4A, right). To quantify these apparent transitions between states, we first fitted hidden Markov models (HMM) to the trial time series of the neural distance differences S_i. An HMM assumes that there is an underlying sequence z_k of states hidden to the observer which emit the observable variables with probability $p(S|z_k)$. Given such a model and the data, the Viterbi algorithm (e.g., Bishop, 2006) finds the most likely sequence of the states z_k as indicated by the color coding in Figure 4A.

A measure of the abruptness of the transitions at the points identified by the Viterbi algorithm was now defined as the number of trials it takes the conditional probability $p(z_i = k|S_i)$ for the new state k given neural choice S_i on trial i to rise from 0.1 × to 0.9 × its range. The analysis was focused on the last of these points which mark the transitions to the spatial rule steady state. Since the probability curves $p(z_i = k|S_i)$ can be quite bumpy, logistic functions were fitted to them using the least-mean-squared-error criterion (with slope as the only free parameter) to obtain statistically more reliable estimates (Figure 4B). The number of transition trials, T_{range}, was defined as the x axis range corresponding to the 10%–90% interval of the y axis range of these curves (see Figure 4B). As shown in Figure 4C, according to this criterion about 25 trials *on average* mark these state transitions. However, the *median* is just about 1.8, since there are cases contributing to the average for which a transition point was hard to discern and which therefore in principle can yield infinitely large ranges. Thus, according to this criterion, for at least half of the data sets the transitions occurred very fast.

To evaluate the statistical significance of this observation, we constructed bootstrap sequences by randomly repositioning the last transition point within some range (see Experimental Procedures) and reassigning states to the trials according to these shifts (see Figure S3). The conditional probabilities $p(z_i = k|S_i)$ were then recalculated for each bootstrap data set based on these state reassignments. If there were no clear and distinct transitions among states within the S_i trial time series but rather a gradual shifting of the neural position across trials, the transition points would be expected to be more or less randomly located, and the bootstraps should not be much different from the original time series (as illustrated in Figure S3B). However, across all 13 data sets and their matched bootstrap means there was a highly significant difference between the number of state-transition trials within the original and the bootstrap data (Figure 4C; Wilcoxon sign-rank test, $T = 0$, $p < 3 \times 10^{-4}$; $t(12) > 3.8$, $p < 0.003$). Moreover, nonparametrically comparing transition times within the original data to those within the corresponding bootstrap sequences, for 5/13 of the individual transition points significance ($p < 0.05$) could be established (compared to less than 1 by chance). Thus, the statistical analysis of these state transitions as identified by HMMs suggests the existence of distinct points in the time series that exceed the usual fluctuations and drifts exhibited elsewhere along the series.

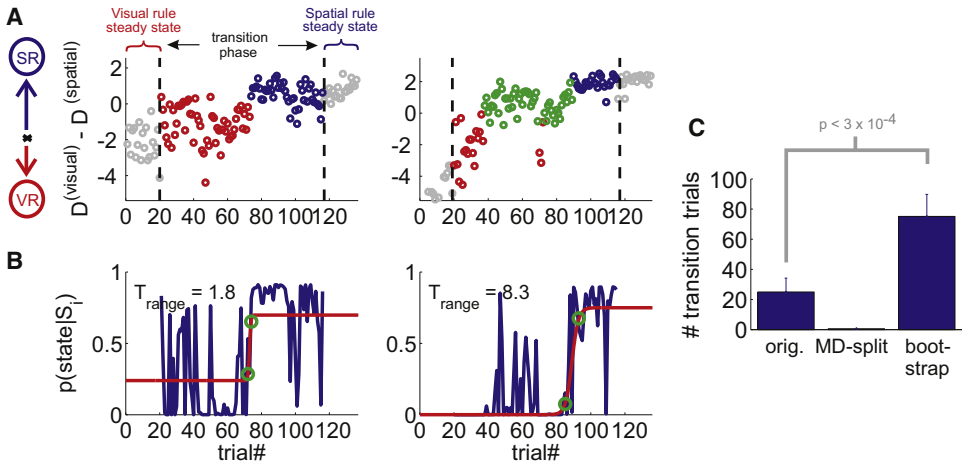

Figure 4. Sudden Transitions among Neural Rule Sets

(A) Difference between the distances to the visual and the spatial steady state rule sets (S_i, see text) as a function of trial number for two of the data sets. The transition phase is enclosed by the two vertical dashed lines. Data points are color-coded according to the states identified by a hidden Markov model. The steady state trials were excluded from the HMM fitting and are therefore indicated in gray. Left, example with two states. Note that a quite steep transition occurs around trial 74. Right, example of a data set which exhibited a third intermediate state.

(B) Conditional probability of the final state, p(state|S_i), associated with the spatial rule given the current position along the neural trajectory as quantified by S_i. These probabilities were fit with logistic functions (red traces) as indicated. The number of trials it took this probability function to traverse from 10% to 90% of its range (as delimited by the green circles) was taken as the "steepness of transition" statistic, T_{range} (equivalently, the slope of the logistic function could have been used).

(C) Average number of trials (T_{range}) it took the 13 recorded networks to transit from one state to the next (left bar), for those networks only which fell below the median (MD, center bar), and for the bootstrap sets for which the transition points were randomly relocated (right bar).

Error bars = SEM. See also Figures S3 and S4.

Another way to address this phenomenon statistically is change point analysis which tries to identify and statistically test discrete points along the time series where a change in mean (or some other parameter) occurred (Kirch 2007, 2008; Huskova and Kirch, 2008). As common in change point analysis, we based the location and test statistics for such points on so-called CUSUM plots which simply give the cumulated sum of differences to the mean along the time series. Thus, given the neural time series S_i as defined above, $CUSUM(S_i) = \sum_{j \le i}(S_j - \langle S \rangle)$ is obtained where $\langle S \rangle$ is the mean of the whole time series. By cumulating along the series of observed values, rather than working on this series directly, this approach reduces variance and hence allows a more reliable identification of change points and their statistical properties. If values S_i are first consistently below the mean $\langle S \rangle$ up to some point $c \geq i$ and consistently above it for $c < i$, one would obtain a steadily decreasing curve up to point c and a steady increase thereafter. Figure 5A illustrates in black the CUSUM curves obtained from the neural time series presented in Figure 4A. The one (left) or two (right) change points, respectively, can be identified as minima in these curves.

To test the statistical significance of such discrete points associated with a change in mean, especially their departure from what would be expected from drift or gradual changes, phase-randomized bootstrap data were constructed (Kantz and Schreiber, 2004). Such bootstraps preserve both the *distribution* of the original time series values S_i as well as their *autocorrelations* (or, equivalently, the power spectrum), and hence—on average—contain the same amount of drift. Original time series were furthermore detrended by linear regression prior to these bootstrap comparisons (see Experimental Procedures). The gray areas in Figure 5A demonstrate the range spanned by about 1000 of such phase-randomized bootstrap time series. Due to their construction (see Supplemental Information), these bootstraps can be very similar to the original time series. Note, however, that the original time series in black are still situated at the lower extremes of the bootstrap range.

For quantitative comparison of original and bootstrap time series, the maximum deflection of the CUSUM graphs from zero (T_{CP}) was used as a test statistic (i.e., the absolute values of the curves' extrema). As was the case for the two examples shown, 9 out of the 13 time series significantly exceeded the bootstrap range (nonparametric $p < 0.05$), and in addition on average there was a highly significant difference ($t(12) > 4.49$; $p < 4 \times 10^{-4}$) between the original and bootstrap statistic across all 13 data sets (Figure 5B). Phase-randomized bootstraps in addition allow to establish confidence limits for the change point location (see Experimental Procedures): in 5/13 data sets the 95% confidence interval spanned no more than six trials, attesting again the sudden nature of the jumps. (Further bootstrap tests are reported in the Supplemental Information.)

Relation of Neural Change Points to Behavioral Performance

Figure 3 confirmed a generally good agreement between behavioral choices and the current state along the neural trajectory. However, given that changes in the neural dynamics appear

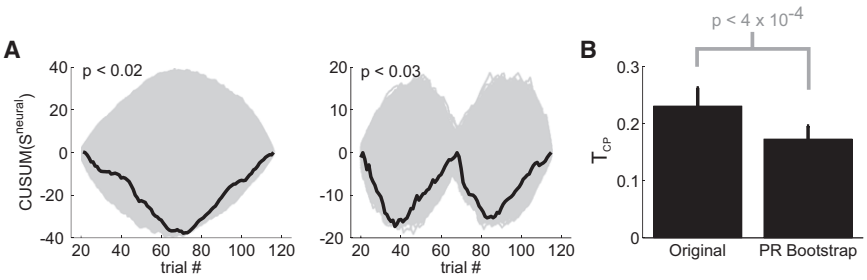

Figure 5. Bootstrap Testing of Change Points in Neural Time Series

(A) Cumulative sum of differences to the mean (CUSUM) plots for the same two data sets as in Figure 4A are shown for the neural choice criterion $S_i = D_i^{(visual)} - D_i^{(spatial)}$ (black curves). The range of ~1000 phase-randomized bootstrap time series is indicated by the gray areas. These bootstraps exhibit drift at the same temporal scales as the original time series, yet could start both below and above the time series mean, hence both upwards and downwards curves are present. However, the change point statistics used are based on the maximum deflection of the curves from zero, and hence the direction of the change in mean is irrelevant: A curve mirrored at the x axis would have exactly the same value as the original curve. The original time series in black are the same as in Figure 6A below (but look slightly different because they were detrended prior to the bootstrap comparisons).

(B) Comparison of change point statistics T_{CP} between the original and bootstrap time series averaged across all 13 data sets (error bars = SEM). PR = phase-randomized.

rather abrupt, the question arises whether more specifically these discrete points in the neural time series correspond to turning points in the behavioral performance. To address this, we first ran the same statistical procedures as used for the neural time series on the behavioral performance curves, and then related CUSUM-based change point indicators obtained from the behavioral curves to the neural ones. Sigmoid function fits to the behavioral performance curves revealed a 10%–90% rise time (T_{range}) with a median of less than 2.8 trials (see Figure S1B for examples). Moreover, significant differences to phase-randomized bootstraps were obtained both overall ($t(12) > 12.4$; $p < 2 \times 10^{-8}$) and in 8/13 cases individually according to the CUSUM-based deviation statistic (T_{CP}) applied to the de-trended behavioral time series (see Experimental Procedures and Supplemental Information). In 7/13 cases the behavioral change point could furthermore be pinned down to a narrow range of just six trials with 95% confidence. Thus, rather abrupt changes were as common in the behavioral as in the neural time series.

Figure 6A illustrates the correspondence between neural and behavioral CUSUM curves for the same two examples used in Figures 4A and 5A. As these graphs demonstrate, the behavioral CUSUM curves were in general quite well aligned with the neural ones (average correlation of >0.79, SEM < 0.08, excluding two cases as explained below), i.e., ups and downs in the behavioral performance were associated with corresponding ups and downs in the neural dynamics (Figure S5 plots these graphs for all 13 data sets from the present study). Importantly, the discrete points of change identified from the neural time series closely matched those extracted from the behavioral performance curves. The relationship between neural and behavioral change points is summarized for all 13 data sets in Figure 6B. As can be appreciated from this figure, there were only two outlying data sets (marked by stars) which did not exhibit any reasonable association between neural and behavioral change points. However, these two were the ones with the weakest (in fact nonsignificant) evidence among all data sets for a change point in the neural dynamics, and in general exhibited comparatively little change (whether gradual or sudden) during the transition trials (see Figure S5). Thus, the lack of correlation with behavioral change points for these two cases may be caused by the absence of strong and reliable neural change points, leading to a more or less random relationship between neural and behavioral time series. Eliminating these two cases therefore from the analysis, a highly significant correlation between neural and behavioral change point locations was revealed as shown. Thus, the behavioral shift from a familiar to a novel rule is accompanied by a sudden transition in the neural dynamic which tightly correlates with the behavioral decision process.

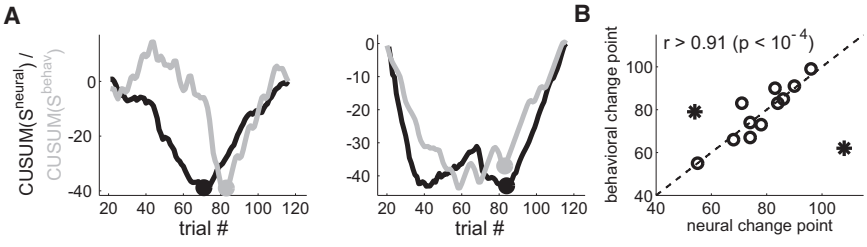

Figure 6. Correspondence of Change Points in the Neural and Behavioral Time Series

(A) The circles show the change points on the neural (black) and behavioral (gray) CUSUM curves as identified by a second derivative criterion (see Experimental Procedures). Behavioral CUSUM curves were scaled to have the same minimum as the neural ones to expose their close relationship.

(B) Relation between change points identified from the neural and corresponding behavioral time series for all 13 data sets. Two outliers (indicated by asterisks) are apparent from the two data sets with the weakest evidence for a change point (see text). One of these animals was also the one with the worst behavioral performance. Excluding these two data sets, a tight and highly significant correlation between neural and behavioral change points is revealed. Note that the straight line is a unity slope curve (corresponding to equivalence of the neural and behavioral change points) and not a regression line.

See also Figures S1B, S5, and S7.

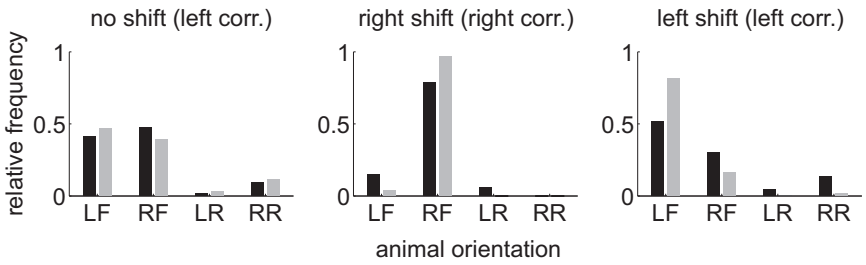

Figure 7. Changes in Body Posture Associated with Changes in the Behavioral Rule around the Change Point Trial

The graphs show the distribution of head orientations to the left and right front versus rear quadrant of the Skinner box before (black bars) and after (gray bars) the neural change point for three representative examples with no shift (left), significant right shift (center), and significant left shift (right). LF = left front, RF = right front, LR = left rear, RR = right rear. See also Figure S6.

Finally, we examined the relation between neural change points and a range of other behavioral indicators (see Experimental Procedures). Given the recently reported importance of body position as a means for encoding working memory and controlling PFC activity (Cowen and McNaughton, 2007; Euston and McNaughton, 2006), the distribution of head orientations of the rat toward each of the four quadrants of the operant box (front left/ right, rear left/ right) was charted each second for the ten trials just preceding and those just following the neural change point. As shown in Figure 7, some animals exhibited a shift in their distributions toward the quadrant associated with the correct response under the spatial rule. Using the Kolmogorov-Smirnov test statistic and permutation bootstraps, this shift in distributions was significant for 5/11 animals for which a neural change point could be reliably identified (see above; see Figure S6 for all pre-/post-change point distributions). Thus, some of the animals may employ a body-based strategy to encode the current rule (as in Cowen and McNaughton, 2007). However, there was no significant difference in the CUSUM-based maximum deviation statistic between the five animals which exhibited such a significant change in distribution, and those that did not ($t(9) < 0.2$; $p > 0.85$). This indicates that the presence of neural change points is not affected by a shift in average head orientation around the change point trial.

DISCUSSION

The ability of animals and humans to infer and apply new rules in order to maximize reward relies critically on the frontal lobes. If PFC networks are compromised, individuals show clear impairments when required either to switch between concepts or to form new concepts or strategies (Milner, 1963; Rich and Shapiro, 2007; Stuss et al., 2000). In the current study we examined how PFC neuronal ensembles switch from encoding a familiar rule to a completely novel rule that could only be deduced through trial and error. By closely tracking the state of the neural network on a trial-to-trial basis while it passed from one rule representation to the other, we noticed that in most cases the transitions between neural states were quite sudden rather than exhibiting a slow gradual change.

There were at least three pieces of statistical evidence supporting this conclusion: First, we fit the posterior probabilities $p(z_i = k|S_i)$ of the neural system being in the spatial rule steady state k given time series observation S_i with a sigmoid function allowed to vary in slope (see Figures 4B and S3). The best fit was defined as the one yielding the smallest residual error from a range of slopes reaching from completely flat (slope of zero) to completely steep (slope approaching infinity). From a gradual transition one would expect a shallow slope to yield the smallest residual error, yet in about half of the time series the final state probability rose from 10% to 90% of its range in less than just two trials. Moreover, the (change) points at which the fits were centered were not arbitrary since placing them randomly along the time series resulted in significantly wider trial-spans (Figure 4C). Second, slow fluctuations and drift in a time series should be well captured by a so-called autoregressive moving-average (ARMA) process, i.e., a linear dynamical process where the current value in a time series is given by a weighted sum over previous values and Gaussian noise (see Supplemental Information). In contrast, a sudden change in mean constitutes a strong nonlinearity and nonstationarity which a pure ARMA process cannot produce (but see Experimental Procedures and Supplemental Information for details). Our phase-randomized bootstraps embody the hypothesis that the fluctuations which remain after detrending are consistent with an ARMA process of any order plus a monotonic time-independent transform (Kantz and Schreiber, 2004). However, this null hypothesis was dismissed overall and in 9/13 cases individually using the CUSUM-based maximum deviation statistic (see Figure 5). Third, in about half of the (detrended) neural and behavioral time series the change point could be located with high (95%) confidence within a narrow interval ($\leq$6 trials), once again attesting that the change must have occurred rather rapidly within just a few trials (see also Figure S4). In all three aspects, similar results were obtained for the behavioral time series and, moreover, the points at which abrupt neural changes happened agreed closely with the change points in behavioral performance (cf. Figure 6B).

Hence, at both the neural and behavioral levels, at least half of the studied time series exhibit significant evidence for sudden transitions. There may be several reasons why tight confidence limits or steep slopes could *not* be established in the remaining cases: This could either be due to a too low signal to noise ratio (insufficient statistical power), e.g., an unlucky draw of the set of recorded neurons (a failure to sample from those exhibiting step changes), or there may be some instances where the transition is indeed more gradual. In general these observations fit very well with a detailed statistical analysis of learning behavior by Gallistel and colleagues (2004), who quantitatively examined learning curves in different species and from a variety of conditioning paradigms, arriving at the same conclusions (see below). The present study may therefore provide the first indication of the *neural* mechanisms underlying this behavioral phenomenon.

Role of the PFC in Rule Deduction and Representation

Inactivation or dysfunction of the mPFC selectively increases perseverative responding, indicative of an inability to respond according to a new rule (Block et al., 2007; Boulougouris et al., 2007; Floresco et al., 2008; Ragozzino et al., 1999). The set-shifting paradigm employed here was chosen because previous work had established that mPFC inactivation specifically impairs performance during this type of rule switch (Floresco et al., 2008), and because it has the distinct advantage that the second egocentric rule is simple enough to be acquired within one recording session, such that the activity of the same neurons could be followed throughout the complete rule transition. Studies in both humans and animals have furthermore shown that damage to the PFC impedes flexible responding only on the first, but not on subsequent shifts between rules (Rich and Shapiro, 2007; Stuss et al., 2000). Therefore, we focused on the deduction of a completely novel rule rather than the shifting between familiar rules or multiple reversals as in most previous studies (Asaad et al., 1998; Everling and DeSouza, 2005; Genovesio et al., 2005; Pasupathy and Miller, 2005; Wallis et al., 2001).

The current findings demonstrate that different rules are represented by distinct patterns of network activity in the PFC which are much more prominent than the distinction between patterns associated with the two visual cues (Figures 1B and 2). Furthermore, these rule representations persisted throughout the analyzed time windows far beyond the actual trial periods (Figures 1A and S2), and were predictive of behavioral choices even before a trial had started (Figure 3B). Surprisingly, the rule representations appeared fairly stable even for some time after a rule had been acquired, despite the fact the PFC is not critical for the maintenance of performance under a familiar rule (Rich and Shapiro, 2007; Floresco et al., 2008). Thus, these observations suggest that the PFC continues to maintain an online representation of the current behavioral strategy even if this information is currently not relevant for responding. The reason for this may be that in order to detect violations of a rule, either due to performance errors or to external alterations in reward contingencies, the PFC would need to maintain a rule-specific state as a point of reference. Thus, although the distinct rule representations may not be critical per se for accurate responding according to well-learned rules, they may well be critical for detecting rule violations and efficient updating of behavioral strategies in the event of a rule shift. This interpretation is consistent with the idea that PFC networks generally monitor behavior in order to resolve response conflicts and detect errors (Ito et al., 2003; Lapish et al., 2008; Narayanan and Laubach, 2008; Shima and Tanji, 1998).

Neural and Behavioral Transition Dynamics during Rule Search

As argued by Gallistel et al. (2004), the appearance of gradual learning curves may simply result from averaging across sessions or animals. In contrast, the learning curves for individual animals on a variety of tasks often seem abrupt and more consistent with an evidence-based decision process where the animal suddenly alters its choice criterion in line with the accumulated evidence that speaks for or against one "hypothesis" (Gallistel et al., 2004; Papachristos and Gallistel 2006; cf. Figure S1B). This formulation of learning fits well with the present observations. Two further observations in our task are consistent with the idea of sudden neural transitions as a reflection of evidence-based decision making: First, we observed that the most rule-selective neurons from all data sets appeared to switch their activity state *between* rather than within trials (Figure S4), i.e., *after* the animal received a new bit of evidence in the form of reward feedback. Second, sudden neural switches were also found in other task contexts related to learning (Figure S7): behavioral performance shifts as rats apparently explored different strategies were observed as well on the visual-rule only training sessions or on a subsequent day of spatial rule training, and were also correlated with abrupt shifts in the network dynamics (Figure S7; although they tended to be less frequent and extreme than those observed on the first shift day in Figures 4–6). In contrast, sessions for which the behavioral performance was comparatively more stable exhibited much less evidence for the occurrence of neural change points (Figures S7B and S7C).

From a dynamical systems perspective, one potential explanation for the observed transition dynamics is that, during the set shift, the inconsistency of the reward feedback with the previously relevant but now incorrect rule leads to the destabilization of one "quasiattracting" state (cf. Durstewitz and Deco, 2008) and the emergence of another corresponding to the new rule set. Such "phase transitions" (bifurcations) are indeed often accompanied by a sudden change in the system's state. Potential neural mechanisms driving this destabilization of previous attractor states and leading to the emergence of new ones include a change in PFC dopamine tone or D2 receptor activation (Durstewitz and Seamans, 2008; Floresco et al., 2006), the accumulation of evidence by climbing activity observed in brain regions that may feed this signal into the PFC (Kim and Shadlen, 1999; Schall, 2001; Durstewitz, 2003) or synaptic changes driven by reward feedback (Schultz, 2006).

In conclusion, in accordance with the interpretation of Gallistel et al. (2004) of abrupt transitions in the learning curves as reflections of altered decision processes, neural transition points may mark shifts between behavioral strategies, i.e., apparent changes in the choice criterion of an animal. In the present problem solving context where the animal had to infer a new rule by accumulating evidence through trial and error, such sudden neural and behavioral transitions may perhaps correspond to moments of "sudden insight" (Epstein et al., 1984).

EXPERIMENTAL PROCEDURES

Animals and Behavior

Twenty-one male Long Evans rats (Charles River) were trained for two sessions (30–40 min/day) on a visual discrimination task, in which they were required to press one of the two levers that had a stimulus cue light illuminated above it to receive reward. The following day, after completing 20 trials on the same visual cue discrimination problem, rats were required to shift to an egocentric spatial strategy in order to receive reward (Figure S1A). In this case, only responses on one of the levers provided reward, regardless of which cue light was illuminated (120 trials, 20 s intertrial interval). The side on which the rewarded lever appeared was assigned as the opposite of a given rat's side bias determined from a prior 20-trial free-choice session.

Surgery and Electrophysiological Recordings

After initial habituation to the task elements, animals underwent stereotaxic surgery for unilateral implantation of the recording electrodes into the PFC. The center of the 32-wire array was placed at coordinates +3.0 AP, 0.8 ML, −3.0 DV in mm from Bregma, with an angle of 4 degrees toward midline and the array was lowered slowly into place and cemented with dental acrylic (Jet dry, Henry-Schein-Asch-Arcona, Melville, NY). All procedures were performed in accordance with the animal care and ethics guidelines of the Canadian Council on Animal Care as well as the UBC Committee on Animal Care.

Electrophysiological recordings were made with a 32-channel Digital Lynx system and Cheetah data collection software in a custom operant chamber (MedAssociates, St. Albans, VT). Video (online frame-capture COHU camera) and behavioral event markers (MedPC) were collected with synchronized timestamps. Individual cells were discriminated using SpikeSort (Neuralynx; Figure S1D), and timestamps were imported into Matlab for analysis. One-third of animals were excluded from analysis for one of two reasons: (1) recordings were of a very poor quality due to poor grounding or (2) histology revealed poor placement of the majority of electrodes. The current data were collected using arrays of single wires, but similar sudden shifts in neural dynamics were observed in two control data sets with tetrode recordings obtained after this study had been finished (see also Supplemental Information for further discussion of this issue).

Data Analysis

To obtain statistically reliable estimates of local spike densities from single trials (e.g., Hastie et al., 2009) for each isolated cell i as a function of time bin t, $r_i(t)$, all spike trains were convolved with Gaussian kernels and binned at 500 ms (a bit shorter than the inverse of the average neural firing rates, <2 Hz), where binning was locked to the appearance of the response levers on each trial. All major results were confirmed for different binnings (200, 500, 1000 ms), different widths of the smoothing kernel (σ = 0.01, 0.25, 0.6 × bin width), different kernel functions (Gaussian, gamma function), max-normalization of firing rates, and using sqrt-transforms to stabilize the variance. Neurons with average firing rates below 0.1 Hz were excluded from further analysis. Testing single neurons for significant "responsiveness" is described in the Supplemental Information.

For population (state space) analysis, population vectors $\mathbf{r}(t) = [r_1(t) \ldots r_N(t)]$ were formed, with N the number of single units isolated from a given recording session. For notational convenience, let L_k, k = 1..4, be the sets of time bins (across steady-state trials as defined in the main text) associated with the four conditions "left-cue/visual-rule" (k = 1), "right-cue/visual-rule" (k = 2), "left-cue/spatial-rule" (k = 3), "right-cue/spatial-rule" (k = 4). A selectivity index for each unit i with respect to the type of cue or type of rule was obtained by grouping the firing rates into two classes corresponding to the two types of cues or rules, respectively, and computing

$$d'_i = \frac{|\langle\{r_i(t)|t\in A\}\rangle - \langle\{r_i(t)|t\in B\}\rangle|}{\sqrt{\sigma^2_{i,t\in A} + \sigma^2_{i,t\in B}}}$$

where $\langle\cdot\rangle$ denotes the mean, A = $L_1 \cup L_3$ and B = $L_2 \cup L_4$ for comparison among cues, A = $L_1 \cup L_2$ and B = $L_3 \cup L_4$ for comparison among rules, and the set of indices L_k were collected from the 3 s periods between cue onset and appearance of the levers on all correctly performed "steady-state trials" for both cues and rules.

Fisher's discriminant criterion (Krzanowski, 2000, Hastie et al., 2009) was used to derive a two-dimensional population representation that best visualized the differences between the steady-state rule and cue activity patterns. The discriminant coordinate representation was obtained from all correctly performed steady state trials and the same 3 s periods as used for sensitivity index calculation, i.e., with data points grouped according to the four sets L_k, k = 1..4, as defined above. To obtain one-dimensional network level selectivity indices for cues and rules (comparable on same footage to the single-unit indices), discriminant analysis was performed with just the two cues (i.e., sets $L_1 \cup L_3$ versus $L_2 \cup L_4$) or the two rules (i.e., sets $L_1 \cup L_2$ versus $L_3 \cup L_4$) as grouping variables, and d'_{network} was computed from the variation on the first discriminant axis.

To check for the statistical significance of group separation, Mahalanobis distances (D) between groups were calculated. Although Euclidean distances often gave similar results, Mahalanobis distances are preferable since they take the spatial spread of the data into account (Figure S2A). To ensure that group comparisons were not affected by the number of data points that went into the Mahalanobis distance estimates, the same number m of time bins were drawn from each of the four sets L_k defined above (where $m = \min_k |L_k|$). Drawings were repeated 1000 times and distance estimates averaged to nevertheless make full use of all time bins recorded. Mahalanobis distances were computed along trajectories, i.e., time bin by time bin within the 3 s periods described above, and the maximum distance ($D^{\max}$) within these periods was chosen as the statistic for comparing conditions (cf. Mazor and Laurent, 2005). The four numbers $D^{\max}_{1-2}$, $D^{\max}_{3-4}$, $D^{\max}_{1-3}$, and $D^{\max}_{2-4}$ were calculated with covariance matrices pooled across the two conditions compared, where the indices refer to the four sets L_k, and the final estimates were obtained as $D^{\max}_{(cue)} = (D^{\max}_{1-2} + D^{\max}_{3-4})/2$ and $D^{\max}_{(rule)} = (D^{\max}_{1-3} + D^{\max}_{2-4})/2$. Covariance matrix estimates were regularized (see Hastie et al., 2009) as further explained in the Supplemental Information. For each of the 13 original data sets, 1000 permutation bootstraps were drawn and averaged where the assignment of correct steady state trials to cue conditions was randomized, i.e., where time bins were randomly shuffled between L_1 and L_2 and between L_3 and L_4, respectively, and $D^{\max}_{(cue)}$ was then recalculated based on these shuffled sets.

To derive statistically robust distance estimates for single trials during the transition phase, all time bins from a 7 s period starting 4 s before cue onset were combined (instead of calculating D time bin by time bin as above). For calculating distances to the two rule steady states this is reasonable given that cells often exhibited rule-selective activity throughout the whole 7 s window (cf. Figures 1A, 3B, and S2). Distances were then computed to the two different rule sets with corresponding cue lights, i.e., to either L_1 and L_3, or to L_2 and L_4, respectively, but covariance matrices were pooled across both cue lights for each rule (i.e., $L_1 \cup L_2$ or $L_3 \cup L_4$) and the set of time bins corresponding to the trial. This was done to reduce error variance in comparing the means yet to achieve robust estimates of the covariance matrices. (Similar results were obtained when pooling across all four sets L_k, yielding a linear instead of a quadratic decision boundary between rule steady states or when the covariance matrix was just estimated from the steady states, i.e., not pooled with the single trials, and held fixed for all single-trial comparisons.)

To evaluate the predictive power of the neural state along the trajectory for behavioral choices, we focused on the conflict trials on which the visual rule and the spatial rule would demand different behavioral responses. Let us denote by $\{a_i\}$ the binary time series of behavioral responses on these conflict trials (with a_i = 1 if the response was correct on trial i, and a_i = 0 otherwise), and by S_i the corresponding neural time series as defined in the main text. The neural predictor was now defined as $P_i = sign(S_i - \langle S_i\rangle/\sigma_S)$, and the level of agreement between behavioral and neural choices was taken to be $|\{i|P_i = 2a_i - 1\}|/|\{i\}|$ where $\{i\}$ denotes the set of all conflict trials and $|.|$ denotes the cardinality of a set. (Text book definitions of linear and quadratic discriminant functions [e.g., Hastie et al., 2009], however, gave similar results.)

Different approaches were used to detect and quantify transitions in the neural and behavioral states during the transition trials. First, a hidden Markov model (HMM) was fitted to the neural time series S_i (including *all* transition trials) using the MVN-HMM toolbox (Kirshner, S., MVNHMM Toolbox, http://www.stat.purdue.edu/~skirshne/MVNHMM). Fits with both two or three hidden states were attempted based simply on the visual appearance of the data. Since the probability curves $p(z_i = k|S_i)$ (see main text for definition) had quite high variance, logistic functions of the form $g(p) = a + \frac{b-a}{1+\exp(s|p-\theta|)}$ were a natural choice to model the transitions in these probabilities without smoothing out the transition point itself. Parameters a and b were set equal to $\langle p(z_i = k|S_i)\rangle$ for the first and the last 10–15 trials, respectively, θ was fixed to the switch point location, and hence only the slope s was subject to least-squared-error curve fitting (starting from different initial conditions including s = 0). The test statistic was now defined as $T_{range} = g^{-1}(a + 0.9[b-a]) - g^{-1}(a + 0.1[b-a])$. Bootstraps were constructed for each time series $\{S_i\}$ by repositioning the last transition point corresponding to the spatial rule transition randomly within a range anchored by the end

points of the sequence minus 9 trials, or the neighboring transition points, if present. The probabilities $p(z_i = k|S_i)$ were then recomputed based on the resulting new sequence. For behavioral time series, sigmoids were fitted directly to behavioral performance curves without prior state identification (see Supplemental Information for additional information).

For change point (CP) analysis, the following CUSUM-based (as defined in the main text) change point location and test statistics were used (see Kirch, 2007, 2008; Huskova and Kirch, 2008): CPs were identified as the points c for which $c := \underset{i}{\operatorname{argmax}} \left| \left(\frac{M}{i(M-i)}\right)^{\gamma} \sum_{j \le i} (S_j - \langle S \rangle) \right|$, while $T_{CP} = \max_i \left| \sum_{j \le i} (S_j - \langle S \rangle)/M \right|$ was used as the test statistic, with $\langle S \rangle$ the mean of the time series, M the total number of trials, and $\gamma = 1/2$ (however, other test statistics including the CP locator above itself, and $T_n^{(3)}(q)$ with q_1 and q_2 as defined in Kirch (2007), were checked as well, yielding similar results). For all bootstrap comparisons sequences were always constrained to have the same length for original and matched bootstrap time series (although formally T_{CP} is independent of sequence length). The first part of the series containing the first CP was removed (up to five trials after the first CP) in those cases where there were two CPs. All time series were z transformed prior to computing test statistics. For comparison with phase-randomized bootstrap data and for all *nonpaired* comparisons, time series were first detrended by linear regression left and right from the identified change point to remove this type of nonstationarity. That is, a linear model was fitted by least-squares separately to the series S_i, $i = 1..c$, and S_i, $i = c+1..M$, and the series corrected to remove the slope. For neural CP analysis, steady state trials were removed as these constituted the reference sets for S_i calculation (hence potentially "breaks" may occur in the curves when moving into or out of these sets). For the comparison between neural and behavioral CPs, we also used another CP locator based on maxima in the second derivative of the CUSUM(S_i) curve and the consistency of a subsequent upwards trend, as explained in the Supplemental Information. Overall, however, the CPs identified by this criterion and the one defined further above were in tight agreement ($r > 0.91$, $p < 1.3 \times 10^{-5}$) and hence gave very similar results.

For phase randomization, put briefly, time series are first transferred into the frequency domain where the amplitudes are kept and the phases are scrambled, and then transferred back into the time domain. For each data set, 999 phase-randomized bootstraps were created using the "surrogates" routine from the TISEAN-3.0.1 package (Hegger et al., 1999; Schreiber and Schmitz, 2000; www.mpipks-dresden.mpg.de/~tisean/). For obtaining bootstrap confidence limits, phase randomization was performed on the residuals (cf. Huskova and Kirch, 2008), i.e., the means right and left from the change point were subtracted off, 999 bootstraps were obtained, means were added on again after phase scrambling, and change point location was determined for each of these bootstraps giving estimated quantiles of the distribution under the H_1. Since for all statistical comparisons except for the behavioral distributions described next we had directed hypotheses, one-tailed tests were employed most of the time.

For comparing distributions of head orientations before and after the (second) change point, the Kolmogorov-Smirnov test statistic $T_{KS} = \max|F_{pre} - F_{post}|$ was used, where F_{pre} and F_{post} are the empirical cumulative distribution functions prior to and after the change point, respectively. Since these distributions are defined over categorical (or ordinal) variables (head orientation), permutation bootstraps were used to check for significance. For the 1000 permutation bootstraps, samples of sizes N_1 and N_2 were randomly drawn (without replacement) from the combined set of pre-CP and post-CP values (i.e., the union of F_{pre} and F_{post}), where N_1 and N_2 refer to the cardinalities of the pre-CP and post-CP sets, respectively. Three types of other behavioral events, "grooming," "rearing," and "exploration," were scored as well during each second of the experiment, but they occurred so frequently (almost every trial [grooming and rearing] or about every third trial [exploration]) that any apparent relationship with the change point location could just be due to chance.

SUPPLEMENTAL INFORMATION

Supplemental Information include Supplemental Experimental Procedures, supplemental text, and seven figures and can be found with this article online at doi:10.1016/j.neuron.2010.03.029.

ACKNOWLEDGMENTS

This work was funded by grants from the Deutsche Forschungsgemeinschaft to D.D. (Du 354/5-1 and Du 354/6-1), the Canadian Institute for Health Research to J.K.S., the Natural Sciences and Engineering Research Council to S.B.F., and NARSAD to J.K.S. and N.M.V. J.K.S. and S.B.F. are Michael Smith Foundation for Health Research Senior Scholars. Many thanks to Emili Balaguer-Ballester for thorough reading of this manuscript and to Liya Ma and James Hyman for collection of control data.

Accepted: March 22, 2010
Published: May 12, 2010

REFERENCES

Asaad, W.F., Rainer, G., and Miller, E.K. (1998). Neural activity in the primate prefrontal cortex during associative learning. Neuron *21*, 1399–1407.

Aziz-Zadeh, L., Kaplan, J.T., and Iacoboni, M. (2009). "Aha!": The neural correlates of verbal insight solutions. Hum. Brain Mapp. *30*, 908–916.

Bishop, C.M. (2006). Pattern Recognition and Machine Learning (New York: Springer).

Block, A.E., Dhanji, H., Thompson-Tardif, S.F., and Floresco, S.B. (2007). Thalamic-prefrontal cortical-ventral striatal circuitry mediates dissociable components of strategy set shifting. Cereb. Cortex *17*, 1625–1636.

Boettiger, C.A., and D'Esposito, M. (2005). Frontal networks for learning and executing arbitrary stimulus-response associations. J. Neurosci. *25*, 2723–2732.

Boulougouris, V., Dalley, J.W., and Robbins, T.W. (2007). Effects of orbitofrontal, infralimbic and prelimbic cortical lesions on serial spatial reversal learning in the rat. Behav. Brain Res. *179*, 219–228.

Bowden, E.M., Jung-Beeman, M., Fleck, J., and Kounios, J. (2005). New approaches to demystifying insight. Trends Cogn. Sci. *9*, 322–328.

Brown, E.N., Kass, R.E., and Mitra, P.P. (2004). Multiple neural spike train data analysis: state-of-the-art and future challenges. Nat. Neurosci. *7*, 456–461.

Chapin, J.K., and Nicolelis, M.A. (1999). Principal component analysis of neuronal ensemble activity reveals multidimensional somatosensory representations. J. Neurosci. Methods *94*, 121–140.

Churchland, M.M., Yu, B.M., Sahani, M., and Shenoy, K.V. (2007). Techniques for extracting single-trial activity patterns from large-scale neural recordings. Curr. Opin. Neurobiol. *17*, 609–618.

Cowen, S.L., and McNaughton, B.L. (2007). Selective delay activity in the medial prefrontal cortex of the rat: contribution of sensorimotor information and contingency. J. Neurophysiol. *98*, 303–316.

Deadwyler, S.A., and Hampson, R.E. (1997). The significance of neural ensemble codes during behavior and cognition. Annu. Rev. Neurosci. *20*, 217–244.

Durstewitz, D. (2003). Self-organizing neural integrator predicts interval times through climbing activity. J. Neurosci. *23*, 5342–5353.

Durstewitz, D., and Deco, G. (2008). Computational significance of transient dynamics in cortical networks. Eur. J. Neurosci. *27*, 217–227.

Durstewitz, D., and Seamans, J.K. (2008). The dual-state theory of prefrontal cortex dopamine function with relevance to catechol-o-methyltransferase genotypes and schizophrenia. Biol. Psychiatry *64*, 739–749.

Epstein, R., Kirshnit, C.E., Lanza, R.P., and Rubin, L.C. (1984). 'Insight' in the pigeon: antecedents and determinants of an intelligent performance. Nature *308*, 61–62.

Euston, D.R., and McNaughton, B.L. (2006). Apparent encoding of sequential context in rat medial prefrontal cortex is accounted for by behavioral variability. J. Neurosci. *26*, 13143–13155.

Everling, S., and DeSouza, J.F. (2005). Rule-dependent activity for prosaccades and antisaccades in the primate prefrontal cortex. J. Cogn. Neurosci. *17*, 1483–1496.

Floresco, S.B., Magyar, O., Ghods-Sharifi, S., Vexelman, C., and Tse, M.T. (2006). Multiple dopamine receptor subtypes in the medial prefrontal cortex of the rat regulate set-shifting. Neuropsychopharmacology *31*, 297–309.

Floresco, S.B., Block, A.E., and Tse, M.T. (2008). Inactivation of the medial prefrontal cortex of the rat impairs strategy set-shifting, but not reversal learning, using a novel, automated procedure. Behav. Brain Res. *190*, 85–96.

Gallistel, C.R., Fairhurst, S., and Balsam, P. (2004). The learning curve: implications of a quantitative analysis. Proc. Natl. Acad. Sci. USA *101*, 13124–13131.

Genovesio, A., Brasted, P.J., Mitz, A.R., and Wise, S.P. (2005). Prefrontal cortex activity related to abstract response strategies. Neuron *47*, 307–320.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). The Elements of Statistical Learning: Data Mining, Inference, and Prediction (New York: Springer).

Hegger, R., Kantz, H., and Schreiber, T. (1999). Practical implementation of nonlinear time series methods: The TISEAN package. Chaos *9*, 413–435.

Huskova, M., and Kirch, D. (2008). Bootstrapping confidence intervals for the change-point of time series. J. Time Ser. Anal. *29*, 947–974.

Ito, S., Stuphorn, V., Brown, J.W., and Schall, J.D. (2003). Performance monitoring by the anterior cingulate cortex during saccade countermanding. Science *302*, 120–122.

Jones, L.M., Fontanini, A., Sadacca, B.F., Miller, P., and Katz, D.B. (2007). Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. Proc. Natl. Acad. Sci. USA *104*, 18772–18777.

Kantz, H., and Schreiber, T. (2004). Nonlinear Time Series Analysis (Cambridge: Cambridge University Press).

Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci. *2*, 176–185.

Kirch, C. (2007). Block permutation principles for the change analysis of dependent data. J. Statist. Plann. Inference *137*, 2453–2472.

Kirch, C. (2008). Resampling in the frequency domain of time series to determine critical values for change-point tests. Statist. Decisions *25*, 1001–1026.

Krzanowski, W.J. (2000). Principles of Multivariate Analysis, Second Edition (New York: Oxford University Press).

Landmann, C., Dehaene, S., Pappata, S., Jobert, A., Bottlaender, M., Roumenov, D., and Le Bihan, D. (2007). Dynamics of prefrontal and cingulate activity during a reward-based logical deduction task. Cereb. Cortex *17*, 749–759.

Lapish, C.C., Durstewitz, D., Chandler, L.J., and Seamans, J.K. (2008). Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex. Proc. Natl. Acad. Sci. USA *105*, 11963–11968.

Mansouri, F.A., Matsumoto, K., and Tanaka, K. (2006). Prefrontal cell activities related to monkeys' success and failure in adapting to rule changes in a Wisconsin Card Sorting Test analog. J. Neurosci. *26*, 2745–2756.

Mansouri, F.A., Buckley, M.J., and Tanaka, K. (2007). Mnemonic function of the dorsolateral prefrontal cortex in conflict-induced behavioral adjustment. Science *318*, 987–990.

Mazor, O., and Laurent, G. (2005). Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. Neuron *48*, 661–673.

Milner, B. (1963). Effects of different brain lesions on card sorting: The role of the frontal lobes. Arch. Neurol. *9*, 100–110.

Nakahara, K., Hayashi, T., Konishi, S., and Miyashita, Y. (2002). Functional MRI of macaque monkeys performing a cognitive set-shifting task. Science *295*, 1532–1536.

Narayanan, N.S., and Laubach, M. (2008). Neuronal correlates of post-error slowing in the rat dorsomedial prefrontal cortex. J. Neurophysiol. *100*, 520–525.

Papachristos, E.B., and Gallistel, C.R. (2006). Autoshaped head poking in the mouse: a quantitative analysis of the learning curve. J. Exp. Anal. Behav. *85*, 293–308.

Pasupathy, A., and Miller, E.K. (2005). Different time courses of learning-related activity in the prefrontal cortex and striatum. Nature *433*, 873–876.

Procyk, E., Tanaka, Y.L., and Joseph, J.P. (2000). Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. Nat. Neurosci. *3*, 502–508.

Ragozzino, M.E., Detrick, S., and Kesner, R.P. (1999). Involvement of the prelimbic-infralimbic areas of the rodent prefrontal cortex in behavioral flexibility for place and response learning. J. Neurosci. *19*, 4585–4594.

Rich, E.L., and Shapiro, M.L. (2007). Prelimbic/infralimbic inactivation impairs memory for multiple task switches, but not flexible selection of familiar tasks. J. Neurosci. *27*, 4747–4755.

Rich, E.L., and Shapiro, M. (2009). Rat prefrontal cortical neurons selectively code strategy switches. J. Neurosci. *29*, 7208–7219.

Sandkühler, S., and Bhattacharya, J. (2008). Deconstructing insight: EEG correlates of insightful problem solving. PLoS ONE *3*, e1459. 10.1371/journal.pone.0001459.

Schall, J.D. (2001). Neural basis of deciding, choosing and acting. Nat. Rev. Neurosci. *2*, 33–42.

Schreiber, T., and Schmitz, A. (2000). Surrogate time series. Physica D *142*, 346.

Schultz, W. (2006). Behavioral theories and the neurophysiology of reward. Annu. Rev. Psychol. *57*, 87–115.

Shima, K., and Tanji, J. (1998). Role for cingulate motor area cells in voluntary movement selection based on reward. Science *282*, 1335–1338.

Stuss, D.T., Levine, B., Alexander, M.P., Hong, J., Palumbo, C., Hamer, L., Murphy, K.J., and Izukawa, D. (2000). Wisconsin Card Sorting Test performance in patients with focal frontal and posterior brain damage: effects of lesion location and test structure on separable cognitive processes. Neuropsychologia *38*, 388–402.

Wallis, J.D., Anderson, K.C., and Miller, E.K. (2001). Single neurons in prefrontal cortex encode abstract rules. Nature *411*, 953–956.

Yu, B.M., Cunningham, J.P., Santhanam, G., Ryu, S.I., Shenoy, K.V., and Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. J. Neurophysiol. *102*, 614–635.