



HAL
open science

Behavioral shifts and action valuation in the anterior cingulate cortex.

René Quilodran, Marie Rothé, Emmanuel Procyk

► **To cite this version:**

René Quilodran, Marie Rothé, Emmanuel Procyk. Behavioral shifts and action valuation in the anterior cingulate cortex.. *Neuron*, 2008, 57 (2), pp.314-25. 10.1016/j.neuron.2007.11.031 . inserm-00906686

HAL Id: inserm-00906686

<https://inserm.hal.science/inserm-00906686>

Submitted on 20 Nov 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Behavioral shifts and action valuation in the anterior cingulate cortex

Quilodran René ^{1,2}, Rothé Marie ^{1,2} and Procyk Emmanuel ^{1,2,3}

¹ Inserm, U846, Stem Cell and Brain Research Institute, 69500 Bron, France;

² Université de Lyon, Lyon 1, UMR-S 846, 69003 Lyon, France ;

³ Centre National de la Recherche Scientifique, DR7, Lyon, France

SUMMARY

Rapid optimization of behavior requires decisions about when to explore and when to exploit discovered resources. The mechanisms that lead to fast adaptations and their interaction with action valuation are a central issue. We show here that the anterior cingulate cortex (ACC) encodes multiple feedbacks devoted to exploration and its immediate termination. In a task that alternates exploration and exploitation periods, the ACC monitored negative and positive outcomes relevant for different adaptations. In particular it produced signals specific of the first reward i.e. the end of exploration. Those signals disappeared in exploitation periods but immediately transferred to the initiation of trials; a transfer comparable to learning phenomena observed for dopaminergic neurons. Importantly, these were also observed for high gamma oscillations of local field potentials shown to correlate with brain imaging signal. Thus mechanisms of action valuation and monitoring of events/actions are combined for rapid behavioral regulation.

INTRODUCTION

Exploring is of primary importance for adapting to challenging situations, constructing novel internal maps, and developing the value of each choice in the context of new environments. Defining when exploration should be terminated is crucial as well. This refers to the ability of efficient adaptive systems to optimize performance and regulate the shift between exploring for rewards and exploiting known resources. The speed of adaptation is a key issue (Kawato and Samejima, 2007). Solutions emerge with models that alternate between flexible or controlled behavioral adaptations and poorly flexible long term habituations (Daw et al., 2005a; Daw et al., 2005b; Doya, 2002; Sutton and Barto, 1998). One main difference between the two controls on behavior is the strong or weak impact of a particular outcome on the evaluation of the action that produced it (Daw et al., 2005b). An outcome can trigger immediate adaptation as

observed in many conditional protocols, or be used as just a piece of evidence weighted against recent reward history (Kennerley et al., 2006). Information issued from outcomes can either be used as positive and negative assessments of action value, and/or as triggers to change or repeat a response. Where in the brain and how information about outcomes participates in fast action valuation and in shifting between behavioral modes is a central question.

Aston-Jones and Cohen recently proposed that a set of frontal areas (namely the anterior cingulate cortex (ACC) and the orbitofrontal cortex (OFC)) might have a key role in producing signals used to control the balance between exploratory and exploitative behavior (Aston-Jones and Cohen, 2005). The proposition is based on the hypothesis that OFC and ACC are critical players for processing rewards and costs respectively. However the role of these structures is still in debate. We focus on ACC functions, but see Rushworth et al for a recent comparative review (Rushworth et al., 2007). The ACC has been reproducibly reported to activate during adaptive behavior and to shift activity between exploratory and exploitative behaviors (Procyk et al., 2000; Walton et al., 2004). The debates over ACC functions have often focused on its involvement and specificity in detection and evaluation of errors and/or conflict in processing. Holroyd and Coles proposed that ACC monitors behavioral errors based on negative prediction errors. The negative evaluation of performance would take place either at the level of motor response or at the level of external feedback indicating errors (Amiez et al., 2005; Holroyd and Coles, 2002). Some neurons in ACC code for a unidirectional error signal, a signal that could be dependant on dopaminergic afferences. This refers to the role of the dopaminergic system in reinforcement learning mechanisms by which prediction errors are used as teaching signals for neural plasticity (Schultz, 2006). Although numerous works have supported a role for ACC in error monitoring, other authors emphasize reports of ACC activation independent of error commission (Botvinick et al., 2004). Indeed recent work

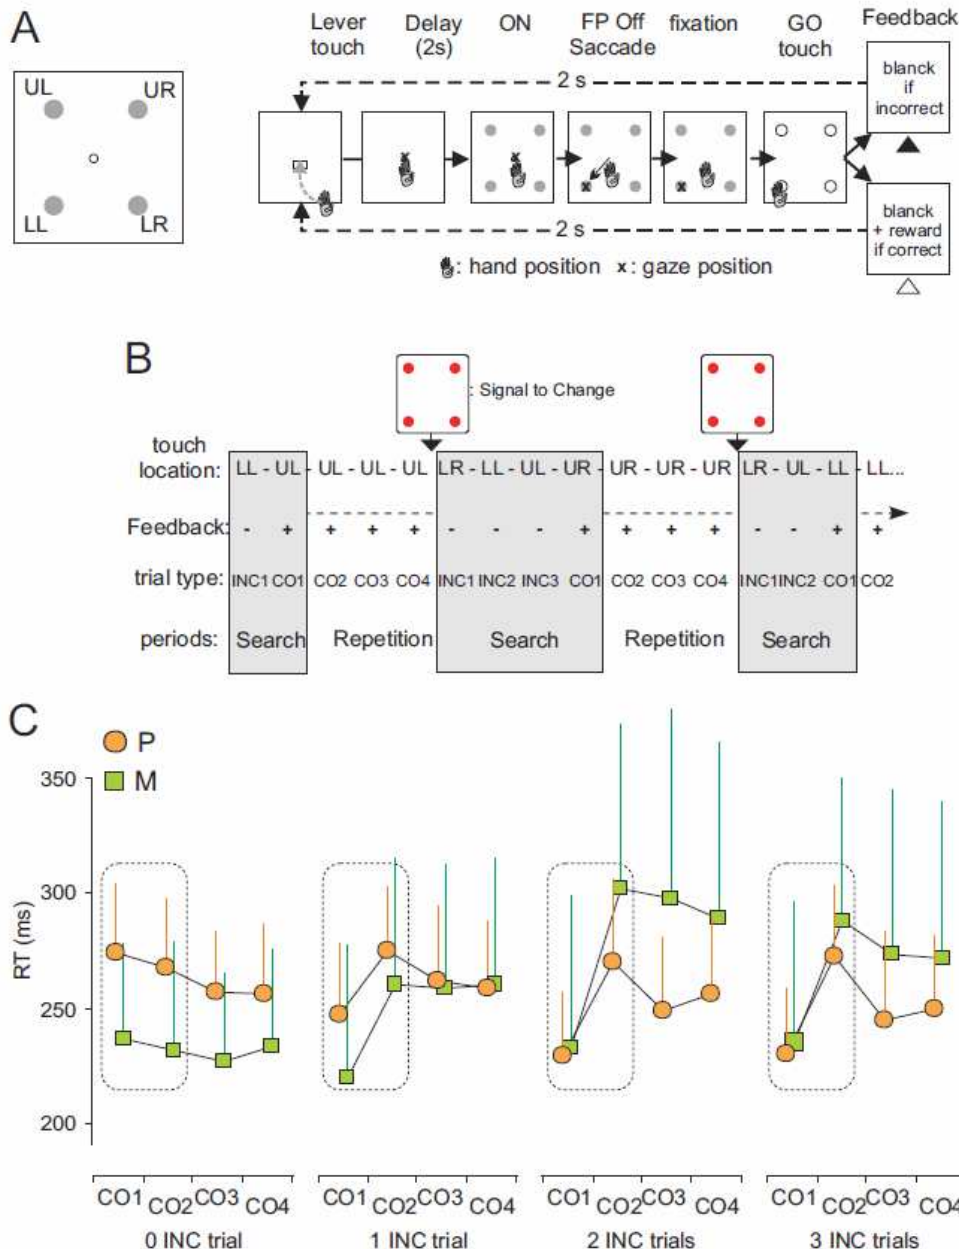


Figure 1. Behavioral Task and performance. The animal had to search by trial and error for the correct target.

A. A trial starts with a touch on the lever and onset of central fixation spot. After a delay period with eye fixation on the central spot, all four targets switched on and the animal made a saccade toward and touched one of them. All targets switched off, and the feedback was given (no reward; negative; black arrow head – reward: positive; white arrow head). **B.** Example: In the first problem, the monkey discovered the solution (UL) in two trials (search period). After discovery the animal was allowed to repeat the response (repetition period) (-: no reward; +: reward; INC: incorrect; CO: correct). **C.** Average reaction times (RT) for target touches (for monkeys M and P; 28 sessions each). RT for CO1, CO2, CO3 and CO4 are shown for problems with different search length (i.e. number of INC trials before the first correct CO1). Note differences for CO1 vs CO2 (paired t-test on individual data; after 0 INC $p < 0,05$; after 1, 2 and 3 INC $p < 0,001$) and CO2 vs CO3 (after 0 INC $p < 0,01$; after 1 INC $p < 0,05$; after 2 and 3 INC $p < 0,001$), but not significant differences for CO3 vs CO4 in both animals;. Both animals show an effect of search length on the difference in RT from search to repetition (CO1-CO2; ANOVA; monkey M: $F(3, 108)=13.241$, $p<.00001$; monkey P: $F(3, 108)=10.693$, $p<.00001$).

in other words in a function that encompasses feedback detection and value adjustment (Rushworth et al., 2004; Walton et al., 2004). A recent experiment in monkeys described positive and negative reward prediction error signals in the ACC during learning (Matsumoto et al., 2007).

showed mid-frontal event-related potentials possibly produced by the ACC and related to correct performance or to reward gain (Gehring and Willoughby, 2002; Vidal et al., 2003). Cohen and colleagues propose that ACC reacts whenever the current task induces conflict in processing for response selection, or even for feedback detection (Botvinick et al., 2004). The conflict monitoring theory suggests that the expectancy violation due to the detection of an unexpected feedback can be conceived as a conflict. This formulation is in some aspects similar to the one invoking prediction errors. More recently, Rushworth and co-workers concluded that the ACC is critically involved in detecting both positive and negative outcomes for the purpose of action valuation,

However, experiments in humans have revealed, using various cognitive tasks, tonic and phasic ACC activations that are difficult to reconcile with a pure reinforcement learning account of ACC function. For instance ACC is activated in dual tasks and self-selection of actions (Kondo et al., 2004; Walton et al., 2004). Several authors have proposed that ACC might form part of a dedicated network involve in task maintenance or task control that is in the ensemble of processes involved in regulating task performance (Braver et al., 2003; Dosenbach et al., 2006; Johnston et al., 2007; Roelofs et al., 2006).




Figure 2. Pre-feedback population activity. Population data for 130 neurons showing statistically significant differences in pre-Feedback activity (0ms to 400ms after the touch; ANOVA $p < 0.01$) between SEARCH (INC and CO1) versus REPETITION (COR) trials. **A.** The average epoch activity in the two periods plotted for each neuron shows two populations with higher or lower activity in search than in repetition periods. **B.** The population histograms for the two populations show the average search vs repetition effect. The dashed grey line indicates feedback onset.

The question thus remains of whether ACC is purely involved in processing negative and/or positive signals, in which circumstances, and if it produces signals for adapting behavioral strategy and/or for evaluating actions through reinforcement learning mechanisms.

Here we assess the characteristics of ACC feedback-related activity during a trial and error protocol that involved different types of behavioral adaptations. The task alternates exploration and exploitation periods with sharp transitions between the two. We show that ACC activity reports and discriminates different types of feedbacks, negative and positive, that relate to diverse adaptations. A shift of activity from positive feedback in exploration to trial initiations in exploitation suggests that these activities participate in action valuation.

RESULTS

We studied ACC unit activities and local field potentials (LFPs) oscillations during a task that alternates exploration (trial-and-error) and exploitation (repetition) periods, and that manipulates outcome expectation and outcome valence. Two monkeys had to search by trial and error which of four simultaneously presented targets was associated to a reward (**Fig. 1A**). In each trial the animal had to choose a target by fixating and then touching it. Targets switched off 600ms after the touch. A reward (positive feedback) was delivered if the correct target was chosen. No reward was given in case of an incorrect choice (negative feedback). Each block of trials (or problem) contained a search period (exploration) during which the animal was searching for the rewarded target and, after its discovery, a repetition period (exploitation) during which the correct response was repeated at least three times (**Fig. 1B**). A visual signal (signal to change – SC) at the end of the repetition period indicated the beginning of a new problem. The animal started a trial by touching a target (this target is subsequently named ‘lever’). The touch induced the onset of a fixation point which marked the initiation of the trial. Subjects were required to fixate the fixation point

until targets onset, and fixate the target once selected by eye (see methods). Any break in fixation requirements resulted in trial cessation (break fixation error – negative feedback). During recordings monkeys performed optimal searches, i.e. did not repeat incorrect trials (INC), and optimal repetitions. Changes in reaction times (RT) between search and repetition revealed a behavioral shift after the first reward (CO1) (**Fig. 1C**; see **supplementary Fig. S1** for detailed analyses) (Procyk and Goldman-Rakic, 2006; Procyk et al., 2000). Interestingly the behavioral shift effect (difference between CO1 and CO2 trials) increased with the length of search periods (see figure legend). This originated from reduced RT in successive search trials for monkey P, and from increased RT in CO2 for monkey M (ANOVA, $p < 0.05$). Analyses of break fixations in repetition revealed that monkeys evaluated break of fixations as such and not as selection errors, and suggested adaptations specific to these execution errors (see notes on performance in experimental procedures). Overall, behavioral data show that monkeys efficiently used each type of feedback.

ACC feedback-related unit activity

Neural activity was recorded in the dorsal bank of the cingulate sulcus within or anterior to the rostral cingulate motor area (**Fig. S2**). Analyses of 546 ACC unit activities (349 and 197 in monkeys M and P) recorded for sufficient numbers of trials revealed a majority (N=280; 51%) of neuronal activity with significant activation following feedback onset. Feedbacks were preceded by different levels of activity in search and repetition. 130 neurons out of 546 showed significant differences between search (INC and CO1) versus repetition (COR) trials (**Fig. 2**). 64% of these cells also showed feedback-related activity. Individual examples are shown in supplementary data. 58% of cells had higher activity in search trials. Pre-feedback changes occurred at different times before or after the touch on targets (**Fig. S3**). Although some individual cells had increased event-related activity in repetition, on average the cell population showing higher activity in repetition than

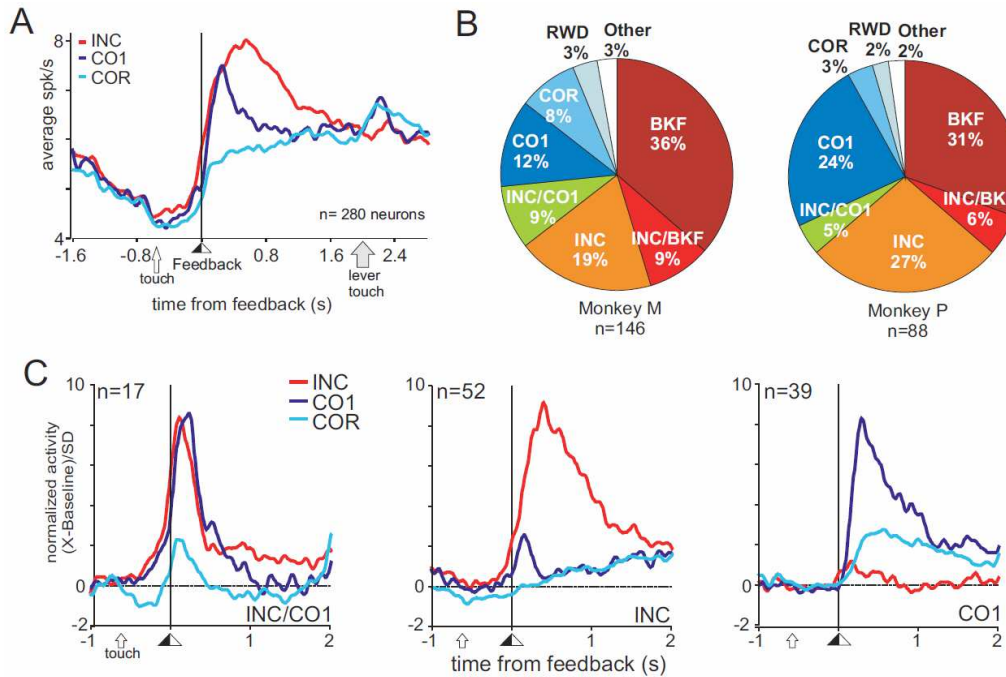


Figure 3. Populations of feedback-related activity. **A.** The grand average unit activity for 280 feedback-related units showing activation for CO1 and errors (right). Time of target touch and time of lever touch are indicated on the x-axes. **B.** Main types of feedback-related activity for each monkey. Note that most activities are related to uncertain feedbacks (BKF -break of fixation-, INC, and CO1). Very few were specific to reward per se (RWD) or reward in repetition (COR). 'Other' includes responses to CO1 and BKF, COR and BKF... The repartitions of types are not different in the two monkeys (X-squared = 7.6164, df = 7, p-value = 0.3676). **C.** Population histograms for three main categories aligned on INC, CO1, repetition (COR), and break of fixation (BKF) -, by evaluating feedback-preferences at the time of feedback, and by grouping cells in different populations.

search revealed a global reduction of neuronal activity during search in anticipation of feedbacks (Fig. 2, center).

Computing the grand average over the 280 feedback-related activities clearly showed that both negative and positive feedbacks in search, but not positive feedback in repetition, induced increased ACC activity (Fig. 3A). Note that the average activity is sustained after errors, but not after the first reward (CO1). We analyzed individual feedback-related data by separating four types of trials - incorrect (INC), first reward (CO1), reward in

The grouping procedure was performed on 234 cells (The analysis had two steps, one to detect significant activation for each feedback and one to group cells by their preference for feedbacks; this procedure was also compared to automatic hierarchical clustering Figure S6 and notes. Both methods gave similar results; the following proportions are issued from the first method - see Experimental Procedures, figures S4 and S5 and supplementary notes for details and discussions). The different populations were recorded in overlapping regions. The assessment of feedback preferences revealed a majority of ACC activities related to INC, CO1 and/or BKF feedbacks, indicating a bias toward processing feedbacks of search periods (exploration) or BKF feedbacks (Fig. 3B). Major groups of ACC neurons were: (i) those that discriminated all feedbacks in search from positive feedback in repetition regardless of whether outcomes in search were positive or negative.

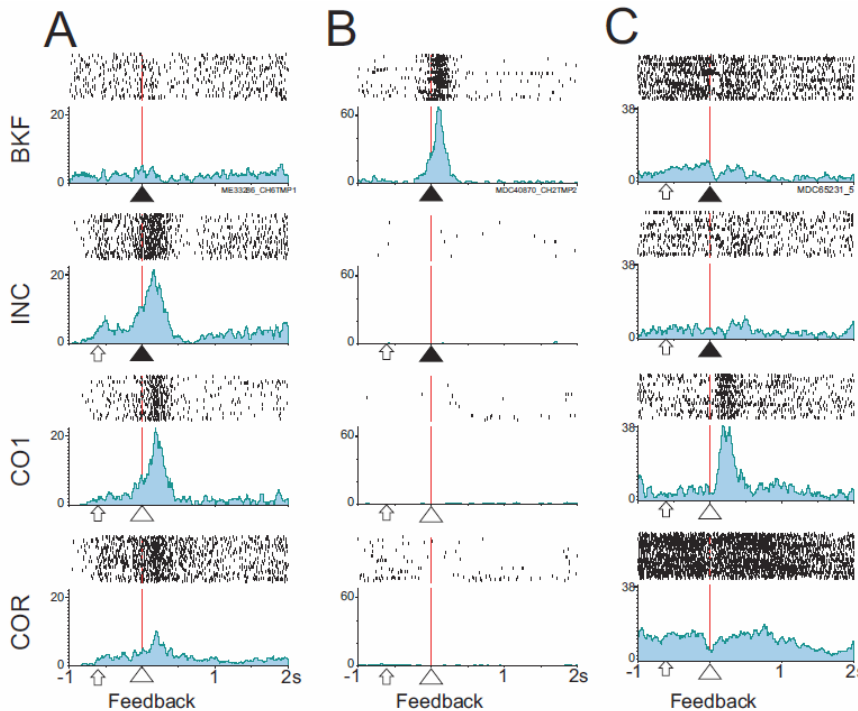


Figure 4. Examples of feedback-related single units. Three examples coding for INC/CO1 (A), BKF (B), and CO1 (C). Histograms are aligned on feedback with rasters on top. Activity is shown for four different feedbacks indicated on the left of the figure. Arrow heads indicate negative (black) or positive (reward; white) feedbacks.

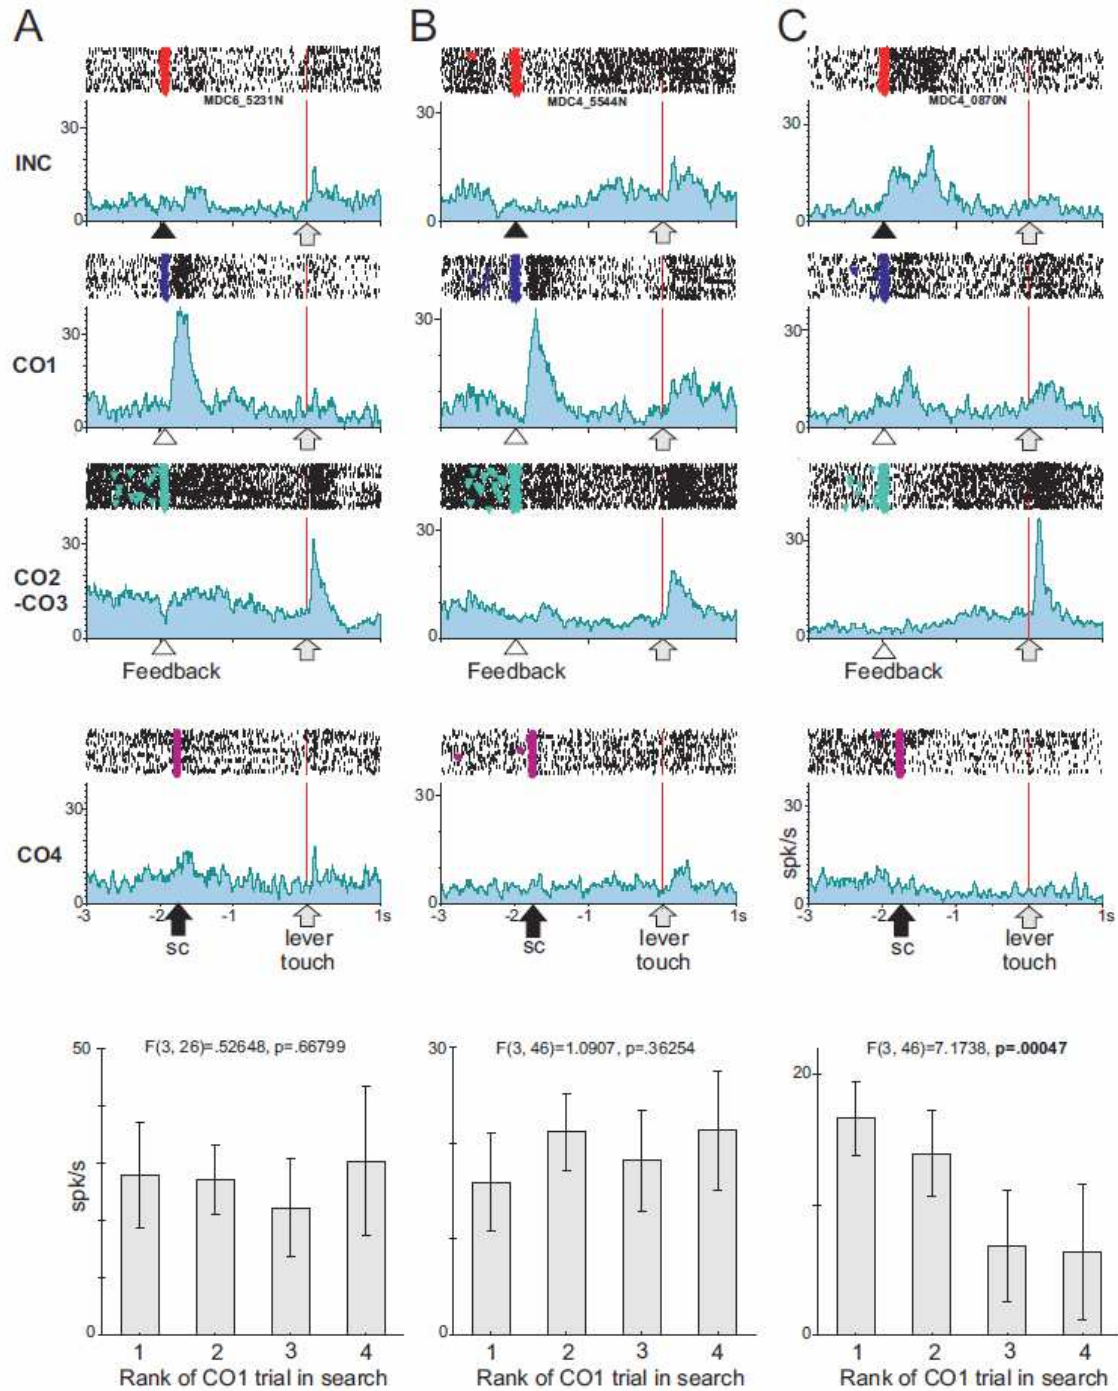


Figure 5. Three examples of reinforcement-related signals transfer in single units. The examples illustrate the transfer of feedback-related activity to lever-onset from the end of search (CO1) to repetition (CO2-CO3). Each column represents data for one single unit. Activity is aligned on lever touch. The three first rows present peristimulus histograms for INC, CO1, 2nd and 3rd correct in repetition (CO2, CO3). In the fourth row the activity is aligned on the lever that follows the signal to change response (SC; upward black arrow) i.e. the first lever touch of a problem. Note that feedback related activity disappears in repetition, and that lever-related activity increases. However, the lever-related activity disappears as soon as the monkey enters a new search following the SC. In the last row, histogram of average activity measured for the four ranks of CO1 trials in search periods. Rank 1 corresponds to discovery in the first trial in search. Ranks 2, 3 and 4 correspond to discoveries after 1, 2 and 3 incorrect trials respectively. Statistical values are the results for one-way ANOVAs.

INC/CO1 neurons (7% in the two monkeys taken together) had increased activity for all INC feedbacks and the first reward (feedbacks in search) (Fig. 3C left; Fig.

4A). Note that on average these activities started before feedback onset which might relate to the anticipatory components of monitoring. (ii) Conversely,




Figure 6. Signal transfer in population activity. **A.** Population activity for the different types of trials in an average problem for neurons showing significant activity at the lever touch (color indicates mean firing rate). Trial types are presented chronologically from top to bottom. The feedback-related response observed in search (search= grey area behind trial type names) is then transferred towards the lever touch after the first reward, after which no more feedback-related response is detectable. The signal to change in CO4 is indicated: black arrow, sc; lever touch in CO4 appears after 1.4s and is not visible on the figure. **B.** Averaged standardized unit activity measured at the feedback and lever touch epochs (grey boxes on the abscissa in A). Note the shift after the first reward, CO1.

COR neurons (6%) reacted only to rewards in the repetition period. These two groups probably concern a subpopulation of the activities reported to vary between search and repetition (Procyk et al., 2000). (iii) Classically described error-related activity following incorrect choices (INC) represented 22% of the feedback-related population (Fig. 3C middle), i.e. 9.5% of total recorded neurons as previously reported (Amiez et al., 2005; Ito et al., 2003). Of all feedback-related activities 8% increased after BKF and INC and were classified as INC/BKF (Fig. 3B). (iv) 34% of feedback-related neurons were said BKF-related (Fig. 3B and 4B). (v) Most remarkably, a critical population reacted significantly only to the first reward (CO1) (17%). The same reward given in repetition (COR) elicited no change in activity for these neurons (Fig. 3C right, 4C, 7). Thus this activity appears with the discovery of the correct target, and the shift between search and repetition. The activity related to CO1 could be a reflection of positive prediction errors as describe for mesencephalic neurons.

In such case CO1 responses should vary according to the probability to be rewarded, i.e. to the rank of CO1 trials in search periods. In other words, reward discovery in the first trial of search periods and reward discovery at the end of search periods (after targets have been eliminated by trial and error) should elicit different neural responses. A large majority (30 out of 31 neurons tested; 97%) of CO1 cells showed no effect of rank of the CO1 trial in search periods and thus did not vary according to prediction errors (ANOVA and Kruskal-Wallis, $p < 0.05$) (Fig. 5A, B). However, 4 out of 17 INC/CO1 neurons (23%) showed an effect. 3 of these 4 neurons showed a progressive decreased of activity with the rank of CO1 trials (see an example in Fig. 5C).

Transfer of feedback-related activity to trial initiation

We further observed that 13.5% of feedback-related neurons ($n=38/280$) significantly increased their activity after the lever touch following a rewarded choice ($>3SD$ from baseline before lever touch – see methods). The effect was also visible on grand average histograms (Fig. 3A). We named this activity lever-related but note that the lever touch is immediately followed by fixation point onset that indicates trial initiation. Figure 5 shows three example units. Two of these units showed increased activity after the first correct feedback (panels A and B). This activity disappeared for COR trials while activity at the initiation of the following trial increased. Interestingly, when the animals were informed of the termination of the repetition period (signal to change - SC) and thus that a new correct target was to be found, the activity after the lever touch disappeared. These neurons ($n=38$) had higher lever-related responses after CO1 or COR rewards than after errors. They were selected without initial consideration for feedback-related activity. Remarkably, the average activity of these neurons revealed minor increases for incorrect feedbacks, a maximum peak for CO1 feedbacks, and no feedback-related activity during repetition (Fig. 6A-B). 24/38 activities (18 in monkey M; 6 in monkey P) were significantly increased after INC and/or CO1 feedbacks. In summary, the population was reactive to the first reward, after which it reacted to the initiation of the next surely rewarded trial and not to the reward itself. Activity at the lever disappeared again when the change in reward contingencies were signaled.

High gamma oscillations related to feedbacks

Feedback-related ACC activity was also found in local field potentials. For 50% of recordings sites ($N=55/110$; monkey M: 34; monkey P: 21), the LFPs revealed significant increased power in high gamma bands after feedbacks (see example of time-frequency diagrams in figure 7). These oscillations from 60 to 120Hz, and often higher, are very similar to those described using




Figure 7. Simultaneous feedback-related LFPs and single unit. Average firing rate of a single neuron (left) and time-frequency composition of LFPs (right) recorded simultaneously by the same electrode. Time-frequency diagrams show the normalized spectral content of LFP from -1s to +1s around feedback onset, expressed in number of sigma of the baseline (-300 to -100ms). The unit is activated for the first reward only whereas the LFPs show an increased power in the high gamma frequency band (60-120Hz) for both incorrect and first correct trials.

intracranial or epidural electrodes in humans (Edwards et al., 2005; Mainy et al., 2006). We focused on variations related to feedback onset for INC, CO1, and COR, although changes were also found after BKF. The increased gamma was significantly higher during search - INC and/or CO1 feedbacks - than repetition (COR) in about 73% of sites (N=40/55; monkey M: 26; monkey P: 14). This includes 20% (N=11; monkey M: 5; monkey P: 6) of sites for which gamma increase was present only for CO1 feedbacks, and 36% (N=20; monkey M: 12; monkey P: 8) for INC. An increase in gamma bands was observed during the repetition period or during both search and repetition periods in 27% of sites (N=15; monkey M: 8; monkey P: 7). Regarding recording locations, no clustering was observed among the different feedback-related activity. Interestingly, feedback-related high gamma oscillations and simultaneously recorded unit activity were not always sharing the exact same functional properties (Fig. 7). Yet, when feedback-related unit activity was recorded, feedback-related gamma was simultaneously observed in 58% of cases. When non feedback-related activity was recorded, feedback-related gamma was found in only

32% of cases (One-sided proportions comparison test: $p < 0.007$). In summary, high gamma oscillations revealed a high incidence of activity related to the processing of both negative and positive feedbacks, when those feedbacks are relevant for behavioral adaptation. Remarkably, the transfer of neural responses from feedbacks to lever touch also appeared at the level of gamma oscillations (Fig. 8A-B). Whereas we observed feedback-related gamma increases during search periods (INC and CO1), the gamma power increased and peaked in relation to lever touch after the first reward and during repetition. As for unit activity, ACC gamma activity related to outcome processing transferred to the initiation of trials once the positive value of a choice was established.

DISCUSSION

Recordings of unit activities and LFPs lead to four main findings on the role of ACC in behavioral adaptation.

First, the same ACC region processes and discriminates different negative and positive events. Remarkably, and in addition to previous reports of ACC responses to task-related rewards or free rewards (Amiez et al., 2005; Ito et al., 2003; Niki and Watanabe, 1979), we show a high incidence of specific responses to the first reward in a problem. Thus, ACC responses to positive feedback depend on its context which, in the PS task, varies between uncertain (CO1) and highly expected (COR). It is possible that the bias of the literature toward a role for ACC in negative feedback processing derives in part from experimental biases, i.e. non equivalence of the relevance of positive and negative outcomes. It is remarkable that most CO1 neuronal activities did not vary with prediction errors, although variations according to prediction errors are found for other types of activity. The CO1 signal thus appears as categorical information and not as a scalar measure of discrepancy between outcome and expectation as observed with other protocols (Matsumoto et al., 2007). One possibility then is that CO1 activity signals the end of the exploratory period and the shift toward a repetition mode of behavior.




Figure 8. Average LFPs and signal transfer. **A.** Time-frequency diagrams showing the averaged signals from 8 sites with INC/CO1 gamma activity. Frequencies are shown from 20 to 150Hz for clarity. Note the disappearance of feedback-related activity for COR and lever-related gamma activity appearing after CO1 and maximum after COR. **B.** Averaged gamma signal between 60-120Hz for the same 8 sites. Note the change in feedback-related activity and the appearance of lever-related activity from search to repetition trials. Profiles are expressed in standard deviations compared to respective baselines (grey boxed on abscissa). Profiles are produced separately for feedback and lever (see experimental procedures).

related mechanisms that are also observable in LFP gamma oscillations. High gamma oscillations are a good marker for studying structure-function specificities in humans, and seem to correlate strongly with functional magnetic resonance imaging bold signal (Edwards et al., 2005; Mainy et al., 2006; Niessing et al., 2005). LFP data can thus explain recent reports of increased bold signal in human ACC for both correct and incorrect performance in a learning paradigm (Walton et al., 2004), and can also clarify the observed absence of

feedback-related ACC activation when measures for negative and positive feedbacks are directly contrasted (van Veen et al., 2004). We predict, from LFP recordings, that using the PS task in humans will reveal feedback-related increased bold signals in the same ACC area for both incorrect and first correct trials. However, in light of the present data we conclude that ACC encodes and discriminates both positive and negative outcomes.

Studying the relationship between unit activity and LFP is crucial. For instance, one key aspect in the debates over the role of ACC has been the discrepancy between human and monkey experiments, the later giving little - if any - support to the conflict monitoring hypothesis (Botvinick et al., 2004; Ito et al., 2003; Nakamura et al., 2005; Rushworth et al., 2004). One can argue that discrepancies come from the use of different recording techniques that give access to distinct aspects of ACC, one related to its output (spike), the other to its input (Bold). LFPs reflect mainly synaptic potentials (unless local multiple spikes are highly correlated) – that is the input of a given cortical area as well as its local intracortical processing - and seem to be a better predictor of Bold increases than single spike firing (Buzsaki, 2006; Logothetis, 2003; Logothetis and Wandell, 2004). Indeed, it has been proposed that although in many cases spikes, LFPs, and Bold can correlate with each other, in some possible instances Bold/LFPs and spike firing can vary independently (in case of modulation, habituation...) (Logothetis, 2003). Therefore, and although clear logical counter-arguments have already been developed (Nakamura et al., 2005), the LFPs/unit comparison needs to be discussed in the

Second, ACC unit activity reflects important feedback-

context of debates over ACC function and for comparisons between human and non-human experiments.

According to the conflict monitoring theory, conflict would arise when an unexpected event conflicts with an expected one, and hereby would be detected at the level of performance feedback (van Veen et al., 2004). Our data show that overall, at unit and LFPs levels, the ACC is responding for positive and negative feedbacks when those are relevant for adaptation. Overall the majority of recordings show signals (spike or LFPs) that discriminate between positive and negative feedbacks which suggest a role for ACC in processing the valence of relevant feedbacks. It is hard to fit conflict detection and specific feedback-related signals. Moreover, conflict detection theory hardly explains the shifts in activity from feedback to trial initiation. We thus propose that explaining ACC feedback-related activity in terms of reinforcement learning mechanisms is more parsimonious.

Third, the transfer of reinforcement-related information from feedback to trial initiation reveals that ACC participates in fast learning mechanisms. It adjoins previous report of ACC activity modulated by reward prediction (Amiez et al., 2005; Matsumoto et al., 2007). The transfer is comparable to the reinforcement learning effect observed for dopaminergic neurons whose activity related to unexpected rewards transfers to conditioned stimuli predicting future rewards (Schultz, 2000). The major differences in our protocol are that ACC activity transfer relates to trial initiation, and that it occurs within a few hundred milliseconds. Indeed, fast valuation of action is an important requirement for organisms to correctly exploit resource discovery, or learn from dangerous or painful incidents. Satoh and colleagues found that mesencephalic activity can relate to trial initiation and be modulated across trial and error similarly to our lever-related activity (Satoh et al., 2003). This supports close relationships between mesencephalic and ACC neurons. However shifts in neural activity toward conditioned stimuli have also been observed in the locus coeruleus that might have important roles in fast adaptation (Bouret and Sara, 2005). In any case we show here the first evidence of a link between ACC feedback-related neural activity and subsequent neurophysiological change in behavioral valuation. It suggests that in our protocol the neural response to the uncertain reward concerns its behavioral relevance.

ACC lesions impair the integration of reward history, and ACC unit activity reflects reward-action associations, average expected values, and negative reward prediction errors (Amiez et al., 2005; Amiez et al., 2006; Kennerley et al., 2006; Matsumoto et al., 2003). In addition, activity of ACC neurons reflects reduction in rewards and action shifting (Shima and Tanji, 1998). In this context, the

present data clearly supports a role for ACC feedback-related activity in updating action values (Kennerley et al., 2006; Rushworth et al., 2004). Recordings of mesencephalic dopaminergic neurons during trial-and-error task revealed reward prediction error-related signals that could influence ACC feedback-related activity (Matsumoto et al., 2007; Satoh et al., 2003). The present findings suggest however that ACC computes various types of reward-related signals that might be more categorical than prediction errors. Models of the relationship between ACC and phasic dopaminergic signals (Holroyd and Coles, 2002) must thus take into account much more than detection of behavioral errors. Our finding on the varieties of feedbacks encoded in ACC, about learning effects, and other reports of ubiquitous modulations related to reward prediction errors, suggest that beyond direct dopamine-ACC relations, feedback detection and evaluation emerge from interactions within larger networks (Haruno and Kawato, 2006; Ullsperger and von Cramon, 2006). In particular, studying the top-down influences of ACC and the differential impacts of dopamine on different cortical targets (pyramidal and inhibitory neurons) and on ventral striatum will give essential information (Seamans and Yang, 2004; Tanaka, 2006). Such research might explain the relationships between the bi-directional coding in dopaminergic neurons (activation and inhibition for positive and negative errors respectively) and the segregated, feedback specific, uni-directional coding in ACC. Overall, the present data suggest important interactions but rule out the idea of ACC feedback-related signals as being simple reflectors of dopaminergic inputs. Simultaneous cortical and mesencephalic recordings will be needed to further study the interactions.

Finally, adaptation of values goes with adaptation of behavior. Our data show that ACC discriminates between different types of feedback allowing appropriate behavioral adaptations (e.g. stay after CO1, shift in response after INC, increase control on fixation after a BKF (See notes about unspecific activity for CO1 and INC)). Activity shows different dynamics after feedback depending on the feedback valence. ACC also codes for visual signals indicating a need to engage explorations (Amiez et al., 2005). Thus the output of ACC is unlikely to be devoted to one specific type of adaptation (for instance shifting response, or pure valuation in terms of reward). Recent data show that another area of the medial frontal cortex, the pre-SMA, has an important role in correctly shifting from automatic to controlled behaviors (Isoda and Hikosaka, 2007). The respective role and collaboration between the two areas are a major issue (Akkal et al., 2002; Rushworth et al., 2004). It is likely that the specificity of ACC relates in part to its position within the reward system and to the use of outcome information for action value adjustments and

behavioral regulation or global changes in goal-directed policy.

Computational modeling suggests that ACC participates in the regulation of cognitive control notably exercised by lateral prefrontal cortex, and that it receives and/or controls modulatory signals that mark transitions between routine and non-routine behaviors (Aston-Jones and Cohen, 2005; Cohen et al., 2004; Dehaene et al., 1998). As such it is a key component of a so-called multiple demand network or global workspace (Dehaene et al., 1998; Duncan, 2006). Recent works in humans suggest that the medial frontal cortex, including part of ACC, with the anterior insula/frontal operculum participate in maintenance of task sets or task control (Braver et al., 2003; Dosenbach et al., 2006). We previously showed that ACC expresses different processing states between search (exploration) and repetition (exploitation) (Procyk et al., 2000), and the present data reveal that these differences are covering different events of trials (delay, target onset, movement, and feedback expectation). ACC then produces specific signals in anticipation or at the onset of events especially when task control is high and when actions are not yet valued, that is during exploration. In similar conditions lateral prefrontal cortex activity is modulated with analogous dynamics between search and repetition (Procyk and Goldman-Rakic, 2006). Our data are compatible with a role for ACC within a regulatory loop involving both noradrenergic and dopamine-based reinforcement learning mechanisms. ACC reward-discovery signals might serve as a trigger to shift from an exploratory to an exploitative mode, and conversely for error-related signals (Aston-Jones and Cohen, 2005; Dehaene et al., 1998). This interpretation can also be formulated in light of the recent demonstration that ACC feedback-related activity correlates with the volatility or uncertainty of the reward environment and with the subjects learning rates (Behrens et al., 2007). Specific ACC states and dedicated processing of outcomes in uncertain environments would convey the tuning of meta-parameters such as learning/adaptation rate (Behrens et al., 2007; Doya, 2002). Although we do not have direct evidence of the consequences of ACC feedback-related activity on lateral prefrontal cortex or on other structures including neuromodulatory systems, ACC itself is likely to be involved in anticipating and characterizing multiple relevant events to trigger appropriate reactions that *in fine* aim a behavioral valuation and possibly global behavioral regulation. Thus the function we attribute to ACC activations is clearly not only to evaluate feedbacks but is also to participate in monitoring the different steps of the task at hand to optimize action adaptation and valuation. A dysfunction of these mechanisms represents the core feature of cognitive alterations observed in addiction and mental illness.

CONCLUSION.

The ACC produces signals that discriminate between various behaviorally relevant positive and negative feedbacks suggesting a role in triggering appropriate adaptations. Specific signals might be at the origin of transitions between different behavioral policies (exploration and exploitations). Transfer of reinforcement signals from time of reinforcement to trial initiation reveals mechanisms that pertained to reinforcement learning and possibly reflects direct interactions between ACC information processing and monoaminergic functions.

Our data reinforce the proposal that ACC is important for establishing action valuations. But they also emphasize a combined role in monitoring events/actions for behavioral regulation when task control is high, underlining the intimate link between control and action valuation.

EXPERIMENTAL PROCEDURES

Housing, surgical, electrophysiological and histological procedures were carried out according to the European Community Council Directive (1986) (Ministère de l'Agriculture et de la Forêt, Commission nationale de l'expérimentation animale) and Direction Départementale des Services Vétérinaires (Lyon, France). Each animal was seated in a primate chair within arm's reach of a tangent touch-screen (Microtouch System) coupled to a TV monitor. In the front panel of the chair, an arm-projection window was opened, allowing the monkey to touch the screen with one hand. A computer recorded the position and accuracy of each touch. It also controlled the presentation via the monitor of visual stimuli (color shapes), which served as light-targets (CORTEX software, NIMH Laboratory of Neuropsychology, Bethesda, Maryland). Eye movements were monitored using an Iscan infrared system (Iscan Inc. USA). Four target items (disks of 5mm in diameter) were used: upper left (UL), upper right (UR), lower right (LR), lower left (LL) (**Fig. 1A**). A central white disk served as fixation point (FP). The lever was disposed just below the FP.

Task. Two male rhesus monkeys were trained in the Problem Solving Task (PS) (**Fig. 1A-B**). Monkeys had to find by trial and error which target, presented in a set of four, was rewarded. Each trial started by the onset of a starting target named 'lever'. The animal had to start a trial by touching the lever and holding his touch. The FP appeared and the animal had to fixate it with his gaze. A delay period (2s) followed, and ended by the simultaneous onset of the four targets. At the FP offset the animal made a saccade toward a target, fixated it (0.5s), and then touched it following the GO signal. All targets switched off at the touch, and a 0.6s delay followed before the feedback was given. A reward (fruit juice) was delivered for choosing the correct target (positive feedback; white arrow head). If a choice in one trial was incorrect (no reward, negative feedback; black arrow head), the monkey could select another target in the following trial and so on until the solution was discovered (search period). Each touch was followed by an interval of 2s.

The animal had to search for the correct target by trial and error. After discovery the animal was allowed to repeat the response. In 90% of cases, after the third repetition, a red flashing signal (the four targets in red) indicated the start of a new problem, i.e. a search for a new

correct target. In 10% of cases the repetition lasted for 7 or 11 trials. A problem was composed of two periods: a “search” period that included all incorrect trials up to the first correct touch, and a “repetition” period wherein the animal was required to repeat the correct touch. The different types of trials are indicated (-: no reward; +: reward; INC: incorrect; CO: correct).

Notes on performance. The two monkeys worked on the PS task with optimal performances: average number of trials in search: 2.4 ± 0.15 trials in monkey M, 2.65 ± 0.23 in monkey P (optimal 2 or 2.5; see (Procyk and Goldman-Rakic, 2006), in repetition: 3.14 ± 0.7 trials in monkey M, 3.4 ± 0.55 in monkey P. A break of fixation (BKF) in repetition (24 and 20% of repetition trials for monkeys M and P) did not induce particular change in response selection in the following trial, preserving good overall performance in repetition. This suggests that monkeys detected and evaluated BKF as such and not as a potential error in selecting a target. This is a behavioral parallel of the discrimination seen at the neural level. Based on a previous publication (Amiez et al., 2005; Procyk and Goldman-Rakic, 2006) we hypothesized that reactions to BKF would vary depending on when the BKF occurred within trials (early when monkeys did not yet overtly selected a target or later when monkeys were selecting the target by gaze). Effectively, analyses showed that, in repetition, monkeys were less likely to repeat a BKF or to make an INC trial after a BKF done at target selection time than after a BKF early in a trial (Chi-Square on numbers of INC, BKF, COR, and Other after early -before 900ms after trial start - and late - between 1800 and 2500ms - BKF: Chi= 79.42, df= 4, $p= 2.30911e-16$, and Chi= 84.18, df= 4, $p= 2.265524e-17$ for monkeys M and P respectively).

RT were significantly different for CO1 and CO2 (i.e. at the shift between search and repetition) for the two monkeys (t-test, df=27, monkey M: $t=-5.03$, $p<0.0001$; P: $t=-4.8$, $p<0.0001$; calculated for 28 days).

Recordings. Monkeys were implanted with a head-restraining device, and an atlas-guided craniotomy was done to expose an aperture over the prefrontal cortex. A recording chamber was implanted with its centre placed at stereotaxic anterior level +31. Neuronal activity was recorded using epoxy-coated tungsten electrodes (1–4 M Ω at 1 kHz; FHC Inc, USA). One to four microelectrodes were placed in stainless steel guide tubes and independently advanced into the cortex through a set of micromotors (Alpha-Omega Engineering, Israel). Neuronal activity was sampled at 13 kHz resolution and LFP at 900 Hz. Recordings were referenced on the guide tubes in contact with the dura and containing the microelectrodes. Recordings sites (see **fig. S2**) covered an area extending over about 6mm (anterior to posterior), in the dorsal bank of the anterior cingulate sulcus, at stereotaxic antero-posterior levels superior to A+30, and at depths superior to 4.5mm from cortical surface. Locations were confirmed by anatomical MRI and histology. This corresponds to a region recorded in previous reports and in which error-related activity has been observed (Amiez et al., 2006; Ito et al., 2003; Procyk et al., 2000). This part of the anterior cingulate cortex lies at the same anterior level than the SEF, and includes part and goes anterior to the rostral cingulate motor area (CMar) as evaluated from previous publications (Ito et al., 2003; Shima et al., 1991).

Unit activity. Single activity was identified using online spike sorting (MSD, AlphaOmega). The activity of single neurons was compared with respect to different events and outcomes resulting from different conditions by using averaged PSTH and trial by trial spike counts (NeuroExplorer, Nex Technology, USA and MatLab – The MathWorks Inc. - home made scripts). PSTHs had a binning of 0.01s and were Boxcar averaged. Neural activity was considered to be significantly different between conditions if it exceeded 5 standard deviations of the mean difference between trial types taken during the window -600 / -200 ms preceding event alignment time, and remained above this threshold for more than six 0.01s bins. MatLab, R v2.5.0 (R Foundation for Statistical Computing), and Statistica (StatSoft Inc.) were used for analyses and graphics.

Neuronal activity was studied at the level of target touch and tested for having significant difference between trials in search periods (INC and CO1 combined) and trials in repetition periods (COR). Activity in each trial was measured in a post-touch epoch (+0ms to +400ms from touch time; the window was limited to +400 to avoid changes in activity related to feedbacks). A 1-way ANOVA, Search vs Repetition, at $p<0.01$ was used to select activity that discriminated between the two periods.

Neuronal unit activity was evaluated for being feedback-related or lever-related. For feedback-related measurements, the average activity aligned on the different feedback types (INC, CO1, COR, and BKF) were standardized to the mean and standard deviation of the baseline activity taken from -600 to -200ms before feedback onset. Resulting data were expressed in number of sigma (SD) of the baseline. Data was analyzed with a first test (TEST1): an activity was said feedback-related when, during the window +60 +800 it passed 5SD (of the baseline) for more than 6 bins (6 x 0.01ms). For Lever-related measurements, the average activity aligned on the first lever touch present after the different feedback types (INC, CO1, and COR) were standardized to the mean and standard deviation of the baseline activity taken from -300 to -20ms before lever touch. Resulting data were expressed in number of sigma (SD) of this baseline. An activity was said lever-related when the average activity measured in the window +50 +350ms after lever touch passed 3SD of the baseline.

Grouping, Clustering: feedback-related activities were grouped in different types with the following strategy: - if one activity showed significant at the 5SD test (TEST1) for only one feedback, then it was defined as being specific for this feedback – if one activity was significant at TEST1 for more than one feedback we ran a post-hoc analysis on differences between feedbacks (TEST2): the standardized average epoch activity was calculated for all feedbacks in the time window showing significant bins. If the standardized epoch averages passed 5SD for several feedbacks and the difference between standardized averages was inferior to 3SD then the activity was defined as being related to the different feedbacks: for instance, a standardized epoch activity measured at 12.667, -0.6868, -3.5318, and 56.362 for INC, CO1, COR, and BKF respectively would be BKF, whereas a standardized epoch activity measured at 6.6393, 0.13089, -2.0113, and 9.5247 would be defined as being INC/BKF. Note that in the first case the activity is also significant for INC feedback; yet it is 5 times higher in standardized activity for BKF. To account for these differences and also for equivalence we defined TEST 2. The choice of a threshold at 3SD is somewhat arbitrary although we visually verified the test with PSTH. The raw statistical data for classification of CO1, INC, and INC/CO1 activity illustrates the implementation of the classification for neurons classified in each category after TEST2 (**Fig. S4, S5**). This classification incorporates specificity of activity and tendencies for activity to be more related to one particular feedback. Although it is not ideal it gives clear ideas about tendencies within populations. Data obtained directly from TEST1 are presented in figure S5.

We evaluated the grouping by feedback preference procedure by using an alternative method, hierarchical clustering. The details are given in supplementary notes and figure S6. In short both methods led to similar grouping of activity and feedback preferences.

Local Field Potentials (LFPs). We analyzed the local field potentials (LFPs) at sites where unit activities have been recorded. Recordings at single recordings sites were analyzed independently, and selected for having correct recording quality (regarding electrical noise in particular). Although variations at different frequency band were observed in time frequency diagrams, we focused this study on high gamma band that showed unique relations to feedback onset.

LFP signals were evaluated with the software package for electrophysiological analysis (ELAN-Pack) developed at the laboratory INSERM U821 laboratory (previously U280; Lyon, France; <http://www.lyon.inserm.fr/821>). LFPs were visualized and analyzed in the form of time frequency graphics. Each set of data represents the activity recorded at one site (one depth for one track) in the ACC. The

time frequency graphics contained the averaged power of each frequency ranging from 2Hz to 150Hz in 2Hz steps across all the trial for the entire data set. Data were analyzed in the time-frequency domain by convolution with complex Gaussian Morlet's wavelets with a ratio f/δ_f of 7. We applied a Blackman window (=100) to eliminate border effects. Trials were aligned on feedback onset (reward or no reward). Data were baseline subtracted. The average signal at all frequencies taken in the last 500ms of trials was subtracted (common baseline) to data because it seemed to be the more neutral period in terms of power variations for the different frequency bands.

The statistical evaluations of post-feedback increases in high gamma oscillations were performed as following. Analyses were made on time-frequency data computed for three types of feedbacks: INC, CO1, and COR. The acceptability threshold was defined at $p < 0,001$ for Kruskal and Wilcoxon tests. The analyzes followed 3 systematic steps: 1) We applied for each type of feedback a Wilcoxon test in order to isolate significant changes in the high gamma frequency band (60 to 120Hz). The level of gamma frequency power at the time and just after the feedback (from -100 to 400ms) was compared to the averaged signal included between -300 and -100ms before the feedback. A profile was generated to determine if these changes corresponded to increases. 2) We then applied a Kruskal-Wallis test on the files where a statistical increase in gamma band had been found to determine the statistical differences in gamma power between the three events (INC, CO1, COR). To rank the signals related to each type of feedback from the higher to the lower (in order to categorize "gamma signals") we compared the averaged power measured from -100 to 400ms and from 60 to 120Hz. 3) Classification: the files where a significant increase in the gamma band was observed were classified in different categories according to the type of feedback generating the signal and according to the rank of each signal in term of power. We obtained 5 main categories of gamma related to:

1. incorrect trials : 36.4% of the files with a gamma signal (INC, that means that the increase of the power in the gamma band occurs only after the incorrect trials or that the level of the power of the gamma signal observed is statistically higher for this condition than for the other generating such a signal),
2. first correct (CO1) : 20%,
3. correct of the repetition period (COR) : 21.8%,
4. the trials of the search period (INC/CO1 where the powers of the gamma signals following these two feedbacks are equal, if there is an increase of the power of the gamma band after a correct trial belonging to the repetition period, it is statistically less powerful) : 16.4%,
5. To these 3 types of feedbacks (ALL, in this case, the power of the signals are equal in the three conditions): 5.4%.

LFP time-frequency diagrams and profiles. The diagrams and profiles presented in figure 8 were made independently for feedback onset and lever touch. The averaged raw power signal for the frequency band 60-120Hz are expressed in standard deviations compared to a baseline taken from the averaged signal (from -500 to -200ms before feedback and from -900 to -750ms before the lever touch for feedback and lever-related signals respectively). Because of occasional jitters in the time delay between feedback onset and lever touch, and because of occasional breaks of fixation just after lever touch, we aligned separately on feedback onset and lever touch and removed for lever-related activity trials for which a break of fixation occurred just after lever touch.

Grand average unit activity (**Fig. 2, 3A**) and average gamma profiles (**Fig. 7-8**) were smoothed with a loess fitting (Statistica soft.).

Aknowledgement.

We thank JP Joseph, H Kennedy, O Bertrand, J Sallet, J Vezoli and P-E Aguera for help on data analyses, reading and discussion of the manuscript, and K. Knoblauch for important helps with statistics. EP is

funded by NRJ foundation, Région Rhône-Alpes, Agence National de la Recherche, Centre National de la Recherche Scientifique, and Institut National de la Santé et de la Recherche Médicale ; RQ is funded by Facultad de Medicina Universidad de Valparaiso, proyecto MECESUP UVA-106, and Fondation pour la Recherche Médicale; MR is founded by Ministère de l'Enseignement Supérieur et de la Recherche.

REFERENCES

Akkal, D., Bioulac, B., Audin, J., and Burbaud, P. (2002). Comparison of neuronal activity in the rostral supplementary and cingulate motor areas during a task with cognitive and motor demands. *Eur J Neurosci* 15, 887-904.

Amiez, C., Joseph, J. P., and Procyk, E. (2005). Anterior cingulate error-related activity is modulated by predicted reward. *Eur J Neurosci* 21, 3447-3452.

Amiez, C., Joseph, J. P., and Procyk, E. (2006). Reward encoding in the monkey anterior cingulate cortex. *Cereb Cortex* 16, 1040-1055.

Aston-Jones, G., and Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu Rev Neurosci* 28, 403-450.

Behrens, T. E., Woolrich, M. W., Walton, M. E., and Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nat Neurosci* 10, 1214-1221.

Botvinick, M. M., Cohen, J. D., and Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: an update. *Trends Cogn Sci* 8, 539-546.

Bouret, S., and Sara, S. J. (2005). Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends Neurosci* 28, 574-582.

Braver, T. S., Reynolds, J. R., and Donaldson, D. I. (2003). Neural mechanisms of transient and sustained cognitive control during task switching. *Neuron* 39, 713-726.

Buzsaki, G. (2006). *Rhythms of the brain* (Oxford, Oxford University Press).

Cohen, J. D., Aston-Jones, G., and Gilzenrat, M. S. (2004). A systems-level perspective on attention and cognitive control. In *Cognitive Neuroscience of attention*, M. I. Posner, ed. (Guilford), pp. 71-90.

Daw, N. D., Niv, Y., and Dayan, P. (2005a). Actions, Policies, Values, and the basal ganglia. In *Recent Breakthroughs in basal ganglia research*, E. Bezard, ed. (Nova Science Publishers).

Daw, N. D., Niv, Y., and Dayan, P. (2005b). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat Neurosci* 8, 1704-1711.

Dehaene, S., Kerszberg, M., and Changeux, J. P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proc Natl Acad Sci U S A* 95, 14529-14534.

Dosenbach, N. U., Visscher, K. M., Palmer, E. D., Miezin, F. M., Wenger, K. K., Kang, H. C., Burgund, E. D., Grimes, A. L., Schlaggar, B. L., and Petersen, S. E. (2006). A core system for the implementation of task sets. *Neuron* 50, 799-812.

Doya, K. (2002). Metalearning and neuromodulation. *Neural Netw* 15, 495-506.

Duncan, J. (2006). EPS Mid-Career Award 2004: brain mechanisms of attention. *Q J Exp Psychol (Colchester)* 59, 2-27.

Edwards, E., Soltani, M., Deouell, L. Y., Berger, M. S., and Knight, R. T. (2005). High gamma activity in response to deviant auditory stimuli recorded directly from human cortex. *J Neurophysiol* 94, 4269-4280.

Gehring, W. J., and Willoughby, A. R. (2002). The medial frontal cortex and the rapid processing of monetary gains and losses. *Science* 295, 2279-2282.

Haruno, M., and Kawato, M. (2006). Heterarchical reinforcement-learning model for integration of multiple cortico-striatal loops: fMRI examination in stimulus-action-reward association learning. *Neural Netw* 19, 1242-1254.

- Holroyd, C. B., and Coles, M. G. (2002). The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol Rev* 109, 679-709.
- Isoda, M., and Hikosaka, O. (2007). Switching from automatic to controlled action by monkey medial frontal cortex. *Nat Neurosci* 10, 240-248.
- 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.
- Johnston, K., Levin, H. M., Koval, M. J., and Everling, S. (2007). Top-down control-signal dynamics in anterior cingulate and prefrontal cortex neurons following task switching. *Neuron* 53, 453-462.
- Kawato, M., and Samejima, K. (2007). Efficient reinforcement learning: computational theories, neuroscience and robotics. *Curr Opin Neurobiol*.
- Kennerley, S. W., Walton, M. E., Behrens, T. E., Buckley, M. J., and Rushworth, M. F. (2006). Optimal decision making and the anterior cingulate cortex. *Nat Neurosci* 9, 940-947.
- Kondo, H., Osaka, N., and Osaka, M. (2004). Cooperation of the anterior cingulate cortex and dorsolateral prefrontal cortex for attention shifting. *Neuroimage* 23, 670-679.
- Logothetis, N. K. (2003). The underpinnings of the BOLD functional magnetic resonance imaging signal. *J Neurosci* 23, 3963-3971.
- Logothetis, N. K., and Wandell, B. A. (2004). Interpreting the BOLD signal. *Annu Rev Physiol* 66, 735-769.
- Mainy, N., Kahane, P., Minotti, L., Hoffmann, D., Bertrand, O., and Lachaux, J. P. (2006). Neural correlates of consolidation in working memory. *Hum Brain Mapp*.
- Matsumoto, K., Suzuki, W., and Tanaka, K. (2003). Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science* 301, 229-232.
- Matsumoto, M., Matsumoto, K., Abe, H., and Tanaka, K. (2007). Medial prefrontal cell activity signaling prediction errors of action values. *Nat Neurosci* 10, 647-656.
- Nakamura, K., Roesch, M. R., and Olson, C. R. (2005). Neuronal activity in macaque SEF and ACC during performance of tasks involving conflict. *J Neurophysiol* 93, 884-908.
- Niessing, J., Ebisch, B., Schmidt, K. E., Niessing, M., Singer, W., and Galuske, R. A. (2005). Hemodynamic signals correlate tightly with synchronized gamma oscillations. *Science* 309, 948-951.
- Niki, H., and Watanabe, M. (1979). Prefrontal and cingulate unit activity during timing behavior in the monkey. *Brain Res* 171, 213-224.
- Procyk, E., and Goldman-Rakic, P. S. (2006). Modulation of dorsolateral prefrontal delay activity during self-organized behavior. *J Neurosci* 26, 11313-11323.
- 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.
- Roelofs, A., van Turenout, M., and Coles, M. G. (2006). Anterior cingulate cortex activity can be independent of response conflict in Stroop-like tasks. *Proc Natl Acad Sci U S A* 103, 13884-13889.
- Rushworth, M. F., Behrens, T. E., Rudebeck, P. H., and Walton, M. E. (2007). Contrasting roles for cingulate and orbitofrontal cortex in decisions and social behaviour. *Trends Cogn Sci*.
- Rushworth, M. F., Walton, M. E., Kennerley, S. W., and Bannerman, D. M. (2004). Action sets and decisions in the medial frontal cortex. *Trends Cogn Sci* 8, 410-417.
- Satoh, T., Nakai, S., Sato, T., and Kimura, M. (2003). Correlated coding of motivation and outcome of decision by dopamine neurons. *J Neurosci* 23, 9913-9923.
- Schultz, W. (2000). Multiple reward signals in the brain. *Nat Rev Neurosci* 1, 199-207.
- Schultz, W. (2006). Behavioral theories and the neurophysiology of reward. *Annu Rev Psychol* 57, 87-115.
- Seamans, J. K., and Yang, C. R. (2004). The principal features and mechanisms of dopamine modulation in the prefrontal cortex. *Prog Neurobiol* 74, 1-58.
- Shima, K., Aya, K., Mushiaki, H., Inase, M., Aizawa, H., and Tanji, J. (1991). Two movement-related foci in the primate cingulate cortex observed in signal-triggered and self-paced forelimb movements. *J Neurophysiol* 65, 188-202.
- Shima, K., and Tanji, J. (1998). Role for cingulate motor area cells in voluntary movement selection based on reward. *Science* 282, 1335-1338.
- Sutton, R. S., and Barto, A. G. (1998). Reinforcement learning: an introduction (Cambridge, MA London, England, MIT Press).
- Tanaka, S. (2006). Dopaminergic control of working memory and its relevance to schizophrenia: a circuit dynamics perspective. *Neuroscience* 139, 153-171.
- Ullsperger, M., and von Cramon, D. Y. (2006). The role of intact frontostriatal circuits in error processing. *J Cogn Neurosci* 18, 651-664.
- van Veen, V., Holroyd, C. B., Cohen, J. D., Stenger, V. A., and Carter, C. S. (2004). Errors without conflict: implications for performance monitoring theories of anterior cingulate cortex. *Brain Cogn* 56, 267-276.
- Vidal, F., Burle, B., Bonnet, M., Grapperon, J., and Hasbroucq, T. (2003). Error negativity on correct trials: a reexamination of available data. *Biol Psychol* 64, 265-282.
- Walton, M. E., Devlin, J. T., and Rushworth, M. F. (2004). Interactions between decision making and performance monitoring within prefrontal cortex. *Nat Neurosci* 7, 1259-1265.

Correspondence. Correspondence and requests for materials should be addressed to E.P. (procyk@lyon.inserm.fr)

Neuron, Volume 57

Supplemental Data

Behavioral Shifts and Action Valuation in the Anterior Cingulate Cortex

René Quilodran, Marie Rothé, and Emmanuel Procyk

CONTENT:

- **Supplementary notes.** Notes on RT. Note on Automatic Hierarchical clustering. Cell grouping, classification and preference. Alternative interpretations of feedback-related activity.
- **Figure S1.** Reaction time data
- **Figure S2.** Locations and proportions of feedback-related activity
- **Figure S3.** Single unit examples showing different activity in Search and Repetition periods
- **Figure S4.** Examples and procedures for Test1 and Test2
- **Figure S5.** Raw data and classification from Test1 and Test2
- **Figure S6.** Comparative analyses – hierarchical clustering

SUPPLEMENTARY NOTES

Notes on RT. Previous experiments using the Problem-solving task showed major differences in Reaction Times (RT) between search and repetition trials. Our previous work (Procyk and Goldman-Rakic, 2006; Procyk and Joseph, 1996; Procyk et al., 2000) shows that those effects are always present although somewhat variable depending on individuals. To show this variability and the average effects, we included in figure S1-A data pooled from 7 monkeys trained in the task detailed in the manuscript. The figure shows the variability between individuals with the overall change in RT variance and values between search and repetition. The two animals used in the present study (monkeys M and P) are also presented. For those two monkeys the tendency was a decrease in reaction times during search and overall longer RT in repetition. The details for each monkey are shown in B together with statistics. Overall we observed that some monkeys speed up their reaction times in approach of the solution (i.e. with increasing number of errors) or slow down. Both reflect anticipatory phenomena; although the first might reveal increasing urge to get the reward, and the second time to recall previous failures. We included behavioral performances obtained in humans showing, on average, an effect similar to the one observed in monkeys (Fig S1-C).

Note on Automatic Hierarchical clustering.

In addition to the main analysis used to characterized activity, we evaluated the grouping of cells by using an automatic hierarchical clustering algorithm of the sort used in molecular biology for gene classification. Note that the hierarchical clustering does not aim at an absolute partitioning of cases but at defining a hierarchy of classes of cases. *Hierarchical clustering* clusters cases based on distance measures and amalgamation or linkage rules.

Cells were classified using the standardized measure of activity obtained for the four feedbacks (INC, CO1, COR, BKF; see methods). The distance chosen was the $1-r$ of Pearson (with r the coefficient of correlation; varies between 0 and 2). The Ward's method was used for the aggregation measurements (This method uses an analysis of variance approach to evaluate the distances between clusters. In short, this method attempts to minimize the Sum of Squares (SS) of any two (hypothetical) clusters that can be formed at each step. [Refer to Ward (1963) J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58, 236.]. From Statistica, StatSoft library).

The results of the classification are presented in figure S6. The upper diagram in A is the vertical dendrogram of the clustering, with the distance of aggregation on the y axis. Cases (cells) are on the abscissa. The lower diagram is a 3D representation of the matrix of distances measured by the automatic method and ordered according to the dendrogram. Values in blue represent short distances. The diagonal is the cell-to-cell correspondence (distance = 0). The blue zones (squares) that appear on the diagonal delimitate clusters for which delimitations can be seen on the dendrogram. Those particular clusters have been chosen from the distance matrix and from a threshold of distance taken on the dendrogram (red dashed line). The choice of a threshold is usually taken above small distances and where large jumps of distances can be observed.

6 clusters were thereby defined and an average population peristimulus histogram was computed for each (bottom). Histograms are shown for each cluster and each feedback (color line). We basically found the main groups defined with the original method: BKF, INC/BKF, INC, INC/CO1, CO1, COR/RWD. The clustering as used here does not discriminate clearly COR from RWD type and the type OTHER, at least at the chosen threshold.

Figure S6B shows the proportions of cells in each cluster found with the two methods (original and hierarchical clustering), combining COR and RWD and removing OTHER. Proportions are identical (Pearson's Chi-squared test: Chi-squared test: $\text{ChiSq} = 11.23664$; $\text{df} = 6$; $p = 0.081$).

In summary the two methods lead to similar identification of main group properties. However, with both methods one should keep in mind that boundaries of classes are hardly absolute.

Cell grouping, Classification and Preference.

Classification is usually understood as a mean to identify separate and clearly delineated categories with distinct properties. Although we do emphasize the existence of separate encoding for specific feedbacks we do not want to say that cells are all clearly functionally segregated. We rather support the idea that ACC unit activities have preferences for particular feedbacks. In functional terms it means that overall ACC has a function in detecting, evaluating and discriminating between different feedbacks. At the unit level, activities are tuned for particular feedbacks (this could be compared to encoding of information in the motor cortex (Georgopoulos, 1995)). The concept 'preference' would thus be preferable to 'class' or 'category' to identify neuronal activities. This, we believe, better reflects the neurophysiological mechanisms. Each neuron contributes at different level to one or more neural populations depending on the feedback. The fact that neurons can be highly selective for feedback reveals that those feedbacks are well differentiated by the structure. This interpretation is also compatible with the idea that different ACC outputs have different consequences. However, the diversity of functional properties of feedback-related activity might also be found in the dynamics of firing, phasic or tonic activations reflecting different functional roles (see below).

Alternative interpretations of feedback-related activity.

Unspecific activity in response to negative and positive feedbacks has been interpreted in terms of requirement or need to direct attention towards feedbacks (Matsumoto et al 2007). In sight of the ensemble of data presented here and because INC/CO1 neurons are active earlier we would favor a refinement of this interpretation and propose that unspecific feedback-related activity (for INC and CO1, i.e. during search) reflects a local setting of the neural state that originates in the high search-related and pre-feedback activation and that optimizes the reception, processing, and eventually transmission of feedback-related information. This mechanism would relate to so-called monitoring. This is supported by lower averaged latencies of the non specific responses (see Figure 3C).

However one should look at the complete data, and in particular at the diversity of feedback preferences. At the present time we cannot decipher whether INC/CO1 ACC activity reflects a specific process or whether it reflects a sub-sample of neurons equally active for negative and positive feedbacks. The two interpretations might not be completely exclusive since it is conceivable that (as proposed in the Cognitive Control Loop theory) detections and evaluation performed by the ACC would trigger attentional and adaptive mechanisms.

One perspective would be to look at multiple dimensions of unit activity. It appears that three properties could be integrated to characterize and interpret feedback-related activity: feedback preference, latency, and duration. It is possible that the processes supported by unit activity depend on what they code for, when, and for how long, reflecting the local cortical dynamic and processing. Future research will be needed to understand the specificity and meaning of such dynamics.

References.

Georgopoulos, A. P. (1995). Current issues in directional motor control. *Trends Neurosci* 18, 506-510.

SUPPLEMENTARY FIGURES

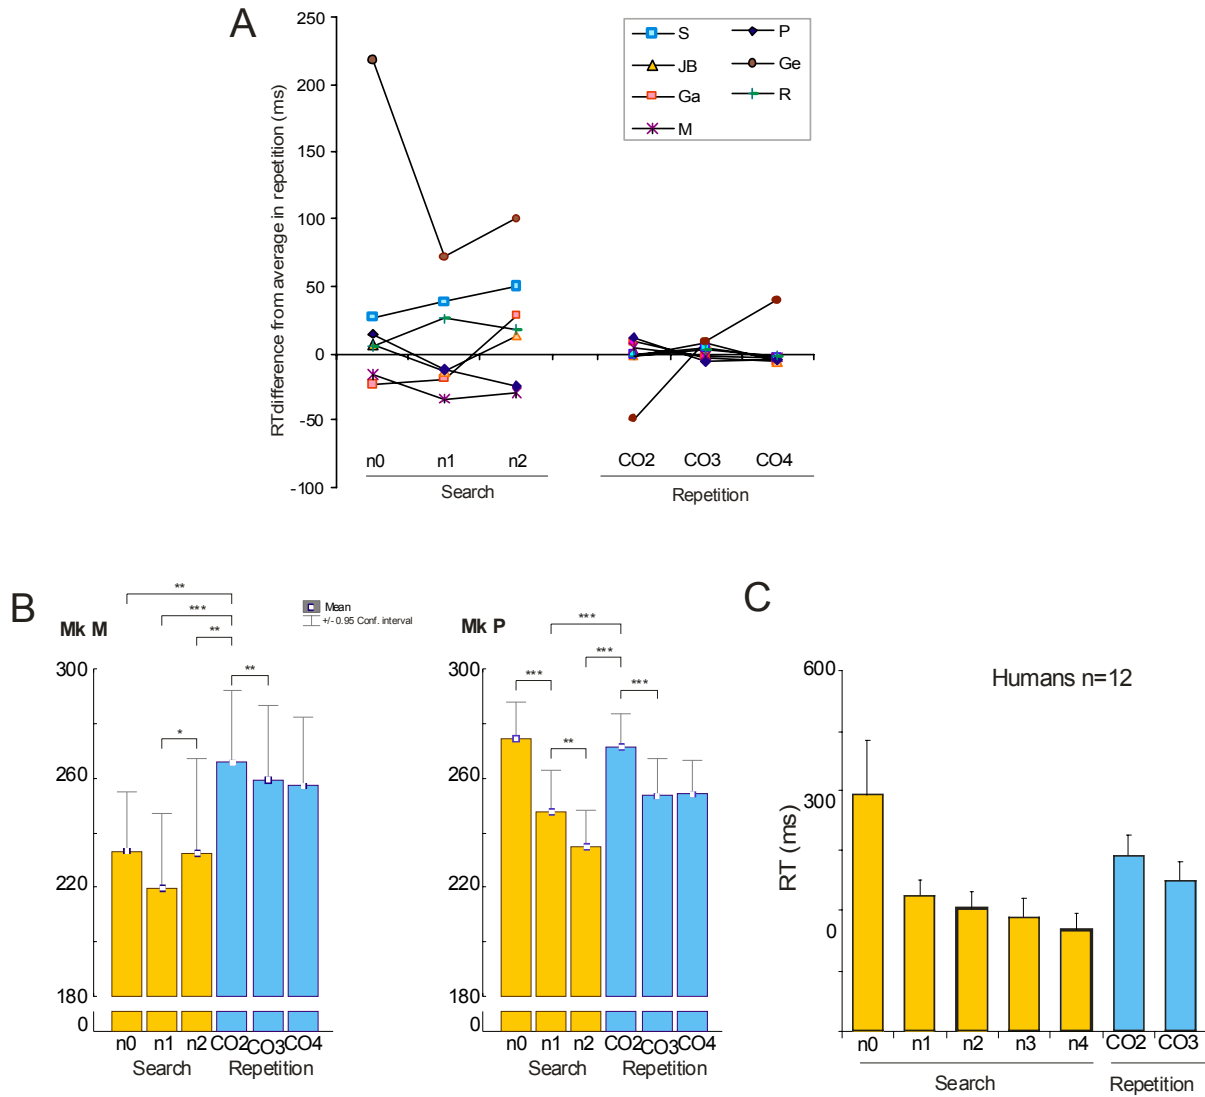


Figure S1. A. Reaction time data from 7 animals trained to the Problem Solving task. The average reaction times (RT) are normalized to the average reaction time measured in Repetition. Trials in Search and Repetition are shown for the different ranks of trials. 'n0' is the first trial in a problem. 'n1' is the second trial pertaining to the Search. CO2 is the first repetition trial. **B.** Behavioral performance of the two animals used in the study. The average reaction times (RT) for target touches during the search and repetition period for each monkey, calculated from 28 sessions of recordings. Paired t-test, * = $p < 0.05$; ** = $p < 0.01$; *** = $p < 0.001$; $df = 27$. **C.** Behavioral data obtained from 12 human subjects performing a Problem Solving task with 5 targets. Movements were displacement of a computer mouse used to move a cursor toward the chosen target. These data are presented for the purpose of comparison with monkey data. Differences between search and repetition were mainly due to RT at n0 (Data obtained by J. Sallet). See also supplementary notes on RT.

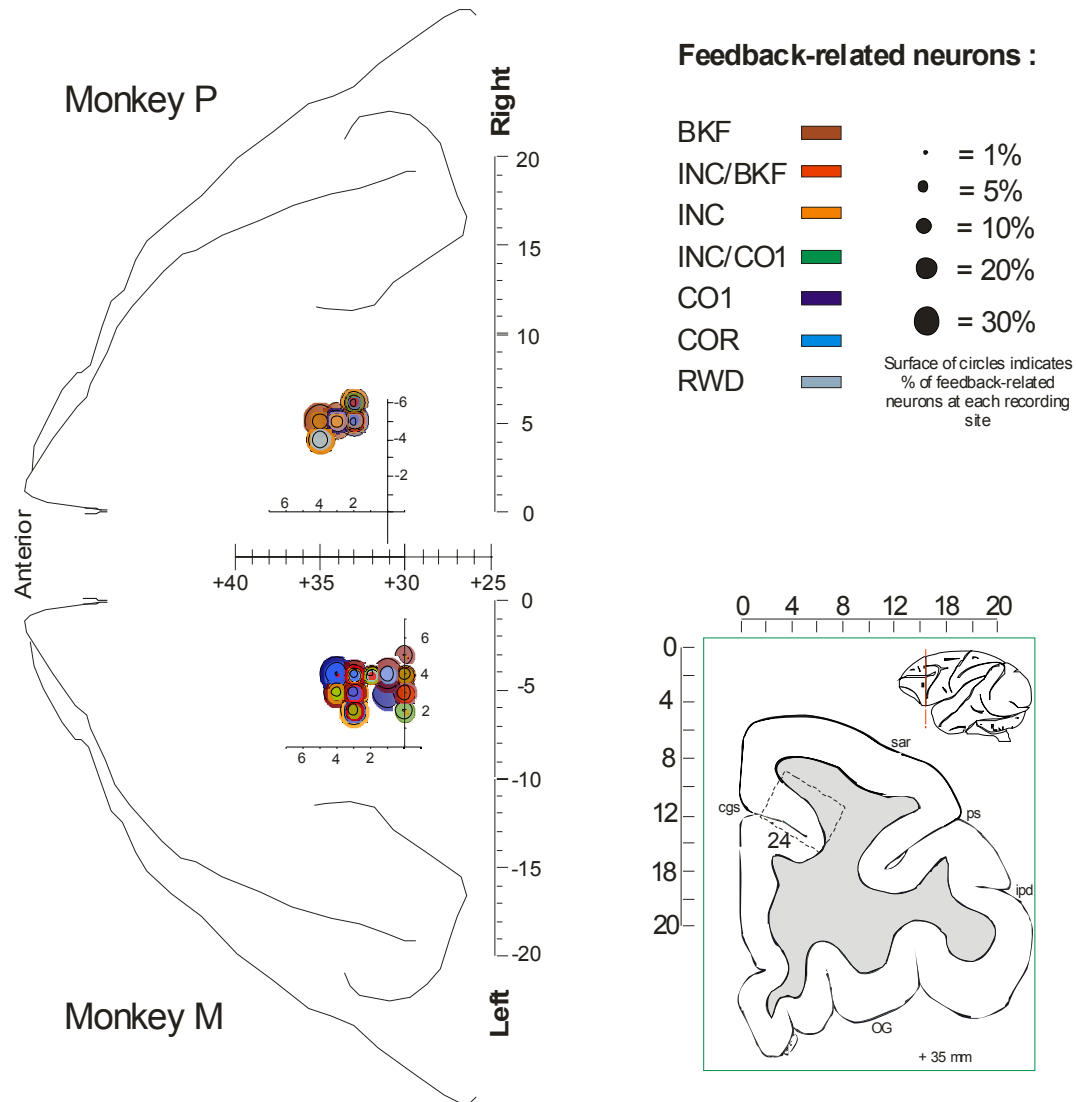


Figure S2. Locations and proportions of feedback-related activity for each recording site in the two monkeys. Coordinates in recording chambers are reported on reconstructed frontal maps as seen from the top. Stereotaxic coordinates are indicated on antero-posterior and latero-medial axes. Colors indicate categories of activity as defined in the paper. The inset at the bottom right shows a slice view adapted from the stereotaxic atlas at anterior coordinate +35 (Saleem and Logothetis 2007, A Combined MRI and Histology Atlas of the Rhesus Monkey Brain in Stereotaxic Coordinates, Academic Press) .

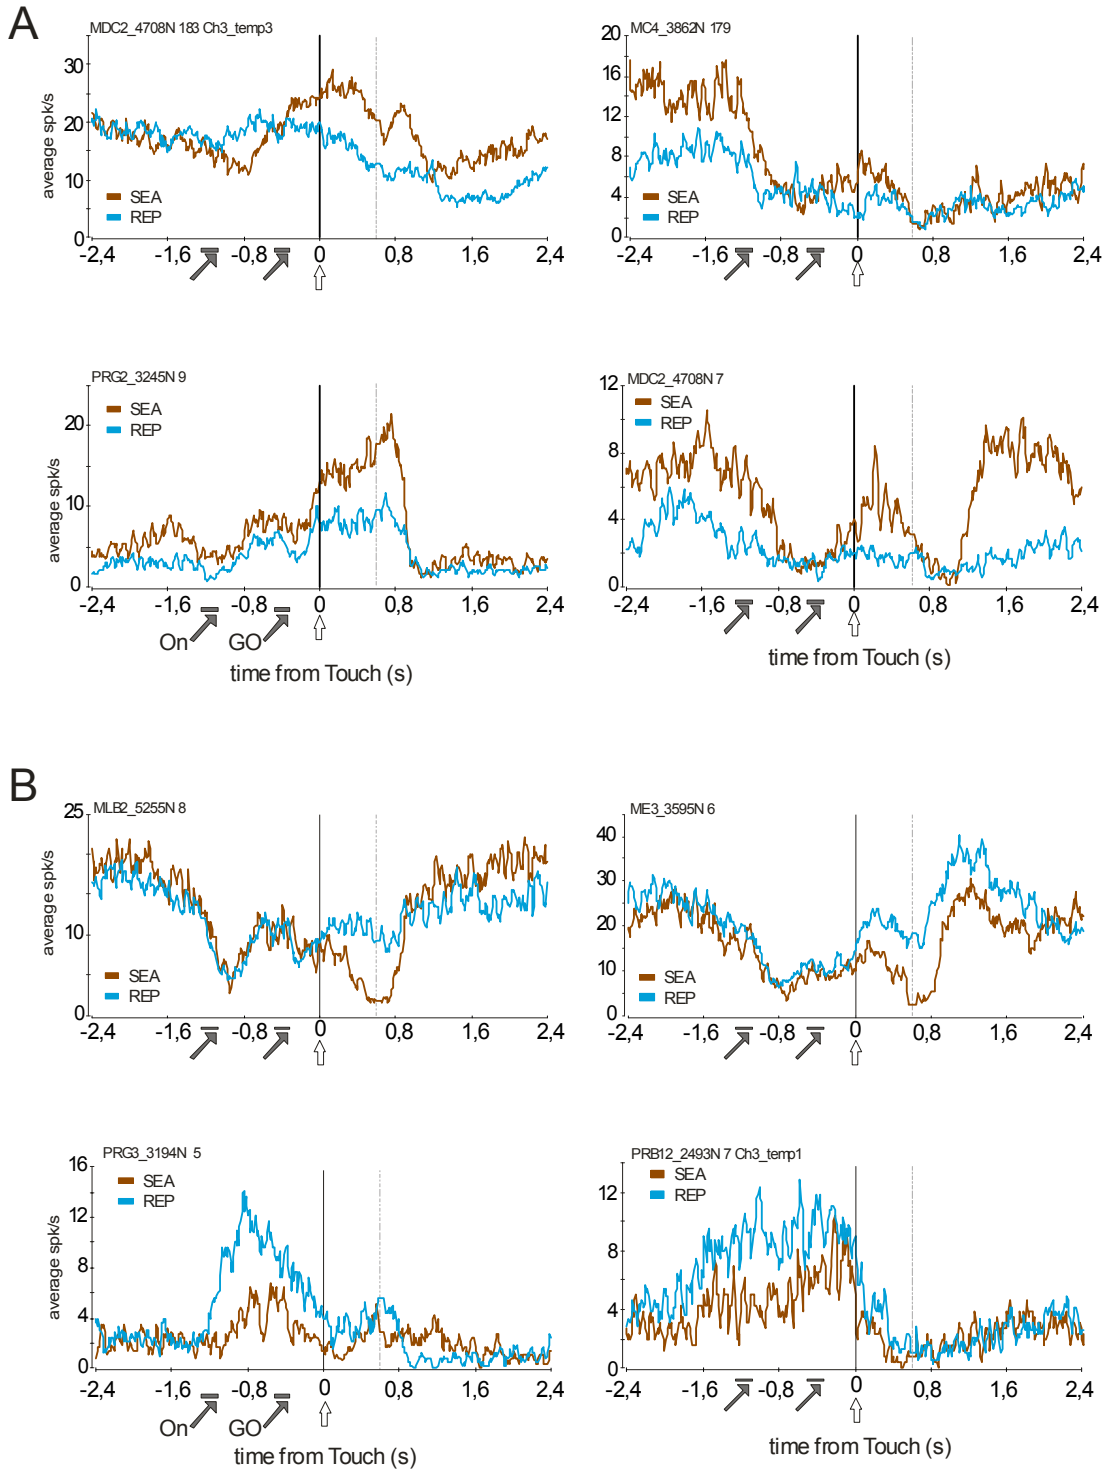


Figure S3. Single unit examples showing different activity in Search and Repetition periods before the feedbacks. **A.** Neurons with higher activity in search. **B.** Neurons with higher activity in Repetition. The average population activity for the two populations is shown in figure 2. Grey dashed line indicates time of Feedbacks.

Note for both populations the presence of differences in activity at different periods during trials (delay, target Onset (On), Touch, post-touch...). Note also that some neurons with effects Search vs Repetition before the Feedback also present feedback-related activity.

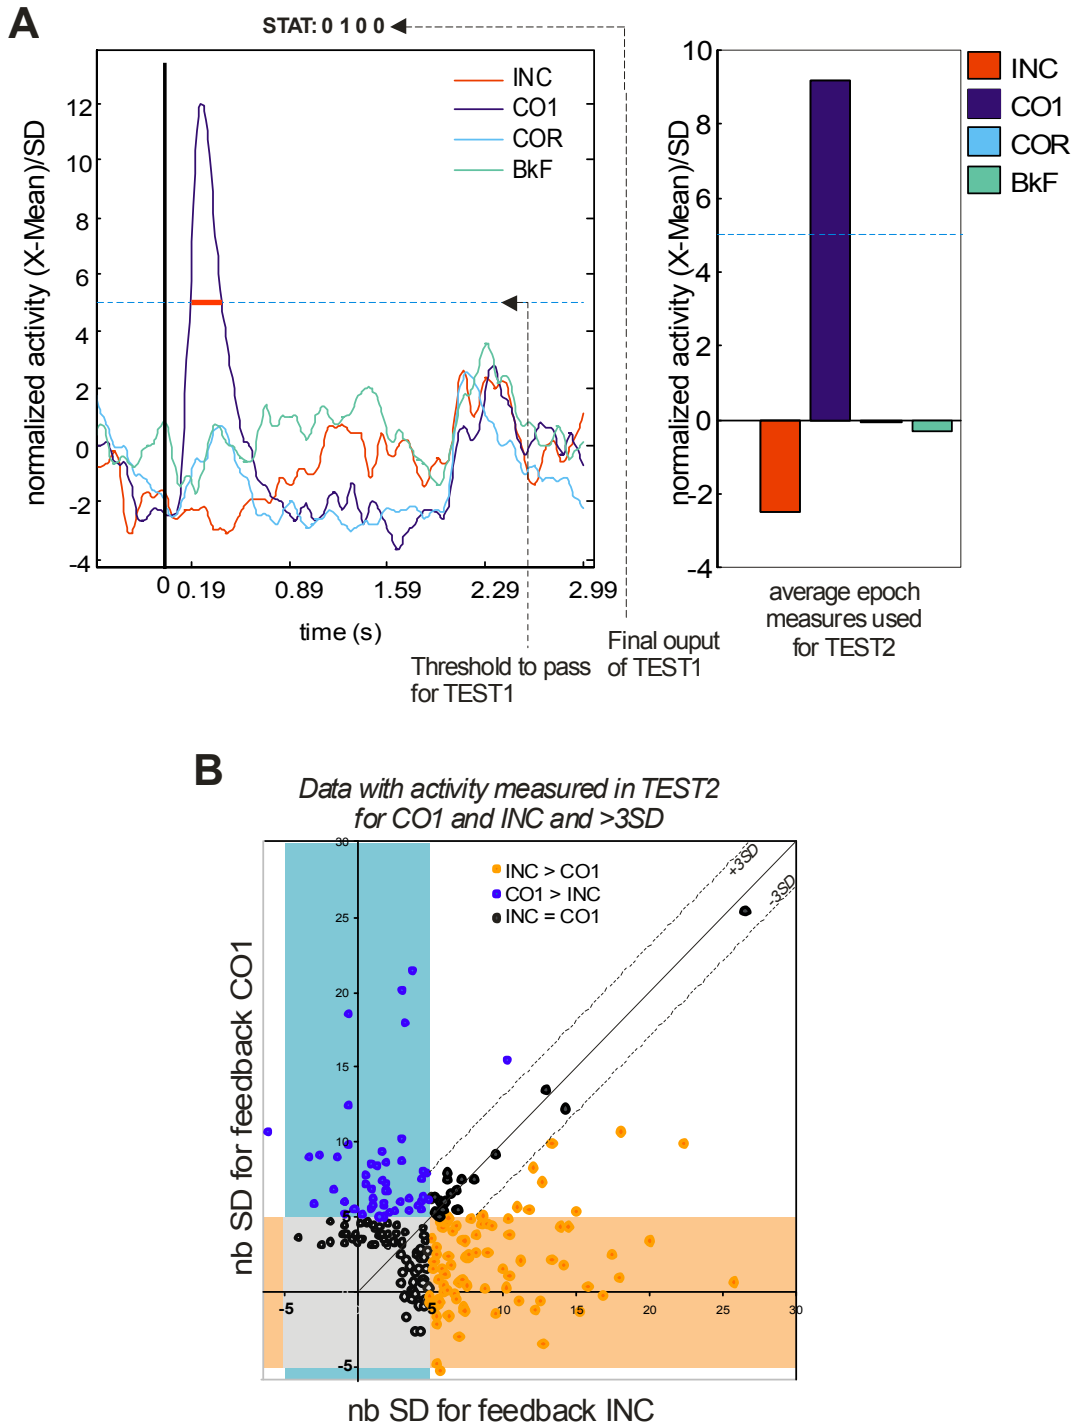


Figure S4. Illustration of the procedures for TEST1 and TEST2 **A.** Example of measures taken on one single unit activity. The activity for the 4 feedbacks are represented on the left histograms. Threshold for TEST1 is at 5SD plus a test on duration (see experimental procedures). The output of the test is indicated (STAT: 0 1 0 0) 1 and 0 denote the significance for each feedback and are used to form the table in figure S5. The histogram on the right represent the average measures used for TEST2. **B.** Plot of data obtained from 231 units having activity superior to 3SD for either INC or CO1. Activities >5SD are used for measure of feedback preferences. Those for which activity for INC and CO1 are >5SD and are not different from each other at 3SD are represented by dark plain circles. Those cells could be identified as INC/CO1 cells depending on their activity for COR and BKF. Note that this plot does not strictly reflect the final selection of INC, CO1, or INC/CO1 cells since it does not take into account data for COR and BKF conditions. See Figure S5.

A Results from TEST1 - N=280

INC	CO1	COR	BKF	NB	
1				24	INC
	1			29	CO1
		1		7	COR
			1	55	BKF
1	1			30	INC/CO1
	1	1		17	RWD
1			1	43	INC/BKF
1	1		1	24	INC/CO1/BKF
		1	1	3	OTHER
	1	1	1	4	
1		1	1	6	
1	1	1		8	
1		1		2	
		1	1	8	
1	1	1	1	20	





Figure S5. Data from TEST1 and TEST2. **A.** number of significant activity at TEST1 for at least one feedback (n=280). The table shows the numbers for each combination. 1 means significant for the corresponding feedback named on top of the table. On the right the corresponding type names are given for clarity. At the bottom the pie chart shows the proportions with output of TEST2 and Clustering (see figure S6). **B.** Raw data for each feedback CO1, COR and BKF against INC for neurons classified as INC (top), INC/CO1 (middle) and CO1 (bottom). Each unit is thus represented by 3 points.

Note that while TEST1 is used to detect significant responses, TEST2 is used to determine feedback preference. However results from TEST1 already show the presence of feedback selectivity.

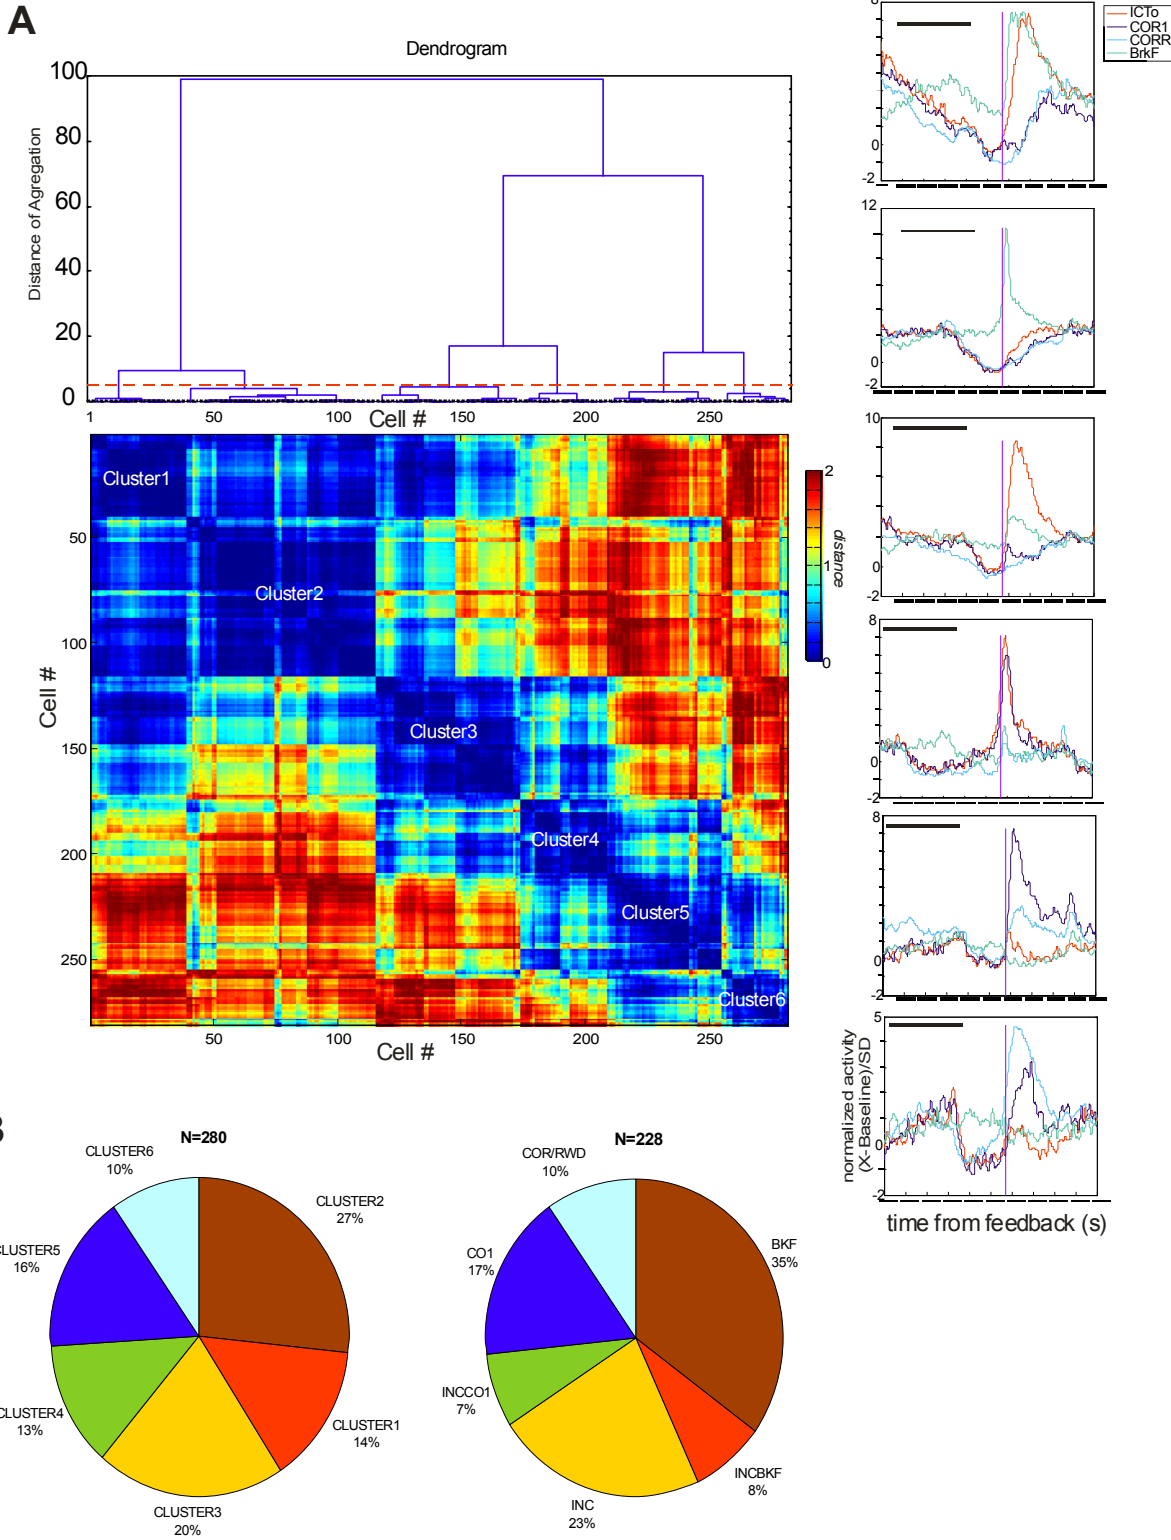


Figure S6. Supplementary analysis for neuron clustering. **A.** Results from automatic hierarchical clustering. See notes for details. On the top the dendrogram represents the tree of classification constructed by the algorithm. **B.** Pie charts show the proportions of cells in each cluster found with the two methods (original and hierarchical clustering), combining COR and RWD and removing OTHER for purpose of comparison. The two populations are not different: $Chi_{sq}=11.23664$; $df=6$ $p=0.081$.