

Human dorsal anterior cingulate cortex neurons mediate ongoing behavioural adaptation

Sameer A. Sheth^{1*}, Matthew K. Mian^{1*}, Shaun R. Patel^{1,2}, Wael F. Asaad³, Ziv M. Williams¹, Darin D. Dougherty⁴, George Bush⁴ & Emad N. Eskandar¹

The ability to optimize behavioural performance when confronted with continuously evolving environmental demands is a key element of human cognition. The dorsal anterior cingulate cortex (dACC), which lies on the medial surface of the frontal lobes, is important in regulating cognitive control. Hypotheses about its function include guiding reward-based decision making¹, monitoring for conflict between competing responses² and predicting task difficulty³. Precise mechanisms of dACC function remain unknown, however, because of the limited number of human neurophysiological studies. Here we use functional imaging and human single-neuron recordings to show that the firing of individual dACC neurons encodes current and recent cognitive load. We demonstrate that the modulation of current dACC activity by previous activity produces a behavioural adaptation that accelerates reactions to cues of similar difficulty to previous ones, and retards reactions to cues of different difficulty. Furthermore, this conflict adaptation, or Gratton effect^{2,4}, is abolished after surgically targeted ablation of the dACC. Our results demonstrate that the dACC provides a continuously updated prediction of expected cognitive demand to optimize future behavioural responses. In situations with stable cognitive demands, this signal promotes efficiency by hastening responses, but in situations with changing demands it engenders accuracy by delaying responses.

Human cognition is characterized by the ability to parse and evaluate a stream of constantly changing environmental stimuli so as to choose the most appropriate response in evolving conditions. The dACC is thought to be important in regulating cognitive control over goal-directed behaviour. Various theories postulate its involvement in linking reward-related information to action^{1,5,6}, monitoring for conflict between competing responses^{2,7,8} or detecting the likelihood of error commission^{3,9,10}. Despite substantial information from studies using lesions, functional magnetic resonance imaging (fMRI) and event-related potentials, the neurophysiological basis of its regulatory role remains the subject of considerable debate.

We studied dACC function with a combination of fMRI, single-neuronal recordings and observations of behaviour before and after lesion in human subjects undergoing surgical cingulotomy, a procedure in which a precise, stereotactically targeted lesion is created in the dACC. Microelectrode recordings, which are routinely performed during the procedure^{11,12}, allowed us to record from individual dACC neurons. Six subjects participated, and in four of these we also obtained a preoperative fMRI with the same task. In four we recorded behavioural responses using the same task immediately after cingulotomy.

Subjects performed the multi-source interference task¹³, a Stroop-like task in which they viewed a cue consisting of three numbers and had to indicate, by pressing a button, the unique number ('target') that differed from the other two numbers ('distracters') (Fig. 1a). By varying the position of the target and the identity of the distracters, the task

established four distinct trial types (Fig. 1b), which were presented to the subject randomly. These trial types contained three levels (type 0, 1 and 2 trials; Fig. 1b) of cognitive interference, operationally defined here as the tendency of an irrelevant stimulus feature (for example, the position of the target) to impede simultaneous processing of the relevant stimulus feature (for example, the identity of the target).

When high-interference (type 2) trials were compared with non-interference (type 0) trials, there was an increased fMRI signal within the dACC (Fig. 1d), indicating increased neuronal population activity during trials with greater cognitive interference. Other cortical regions known to be involved in this decision-making network, such as the

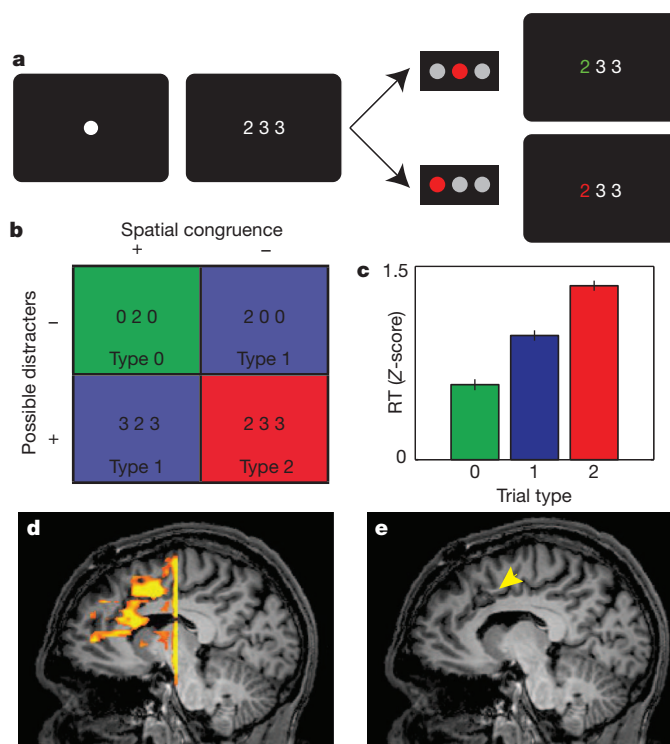


Figure 1 | Behavioural task, fMRI and subject performance. **a**, The multi-source interference task. **b**, The four trial types, based on the presence or absence of spatial congruence between the position of the target and correct button response and on whether the distracters are possible (1, 2, 3) or impossible (0) button choices. **c**, Increase in RTs with increasing cognitive interference in the cue ($P < 10^{-20}$, ANOVA; error bars indicate s.e.m., $n = 1,545$). **d**, Representative example showing greater fMRI activation in the dACC during high-interference trials than in non-interference trials. **e**, A parasagittal slice depicting the lesion (arrowhead), which was also the site of microelectrode recordings.

¹Nayef Al-Rodhan Laboratories, Department of Neurosurgery, Massachusetts General Hospital, Boston, Massachusetts 02114, USA. ²Department of Anatomy and Neurobiology, Boston University School of Medicine, Boston, Massachusetts 02118, USA. ³Department of Neurosurgery, Alpert Medical School, Brown University and Rhode Island Hospital, Providence, Rhode Island 02912, USA. ⁴Department of Psychiatry, Massachusetts General Hospital, Boston, Massachusetts 02114, USA.

*These authors contributed equally to this work.

dorsolateral pre-frontal cortex (DLPFC), were similarly activated to a greater degree in the high-interference condition (Supplementary Fig. 1). The spatial distribution and magnitude of these changes were similar to those previously observed in healthy volunteers^{14,15}, suggesting that this function is spared in the dACC in our subject population and is comparable to that in normal subjects. We registered the postoperative MRI (Fig. 1e) with the preoperative fMRI, confirming that the recording and lesion site localized together with the region of fMRI activation.

During intraoperative microelectrode recordings, subjects performed the task accurately, with an error rate of 1.4%. Reaction times (RTs) were modulated by degree of interference in a dose-dependent fashion (Fig. 1c and Supplementary Fig. 2; $P < 10^{-20}$, analysis of variance (ANOVA)). The trial-type-dependent RTs and low error rates were consistent with the tendency to sacrifice speed for accuracy that is often observed in Stroop-like tasks^{8,16}.

We recorded 59 well-isolated single dACC neurons, with an average baseline firing rate of 5.7 ± 0.7 (mean \pm s.e.m.) spikes per second. We identified three distinct subpopulations of neurons on the basis of their maximal task-responsiveness: first, those firing preferentially before the cue ($n = 12$; 20%); second, those firing preferentially after presentation of the cue ($n = 24$; 41%); and third, those firing preferentially after the behavioural choice ($n = 23$; 39%). The largest group, or cue-responsive neurons, showed distinct modulation of firing on the basis of the degree of interference present in the cue ($P = 0.02$, ANOVA; Supplementary Table 1). Paralleling the pattern for RTs, firing rates for type 2 trials were higher than those for type 1 trials, which were higher than those for type 0 trials. This effect during the cue epoch was observable at the level of individual neurons (Fig. 2a, b), as well as at the cue neuron population level (Fig. 2c, d). Inclusion of the entire recorded neuronal population produced similar effects (Supplementary Fig. 3), and using raw rather than normalized rates did not change this result. Neuronal activity within the dACC was thus correlated with the degree of cognitive interference present in the cue.

The trial-type-dependent modulation in firing rate could be a consequence of dACC neuron sensitivity to either the amount of conflict engendered by the cue^{2,17} or the number of potential responses activated by the cue¹⁸ (Supplementary Note 1). To distinguish between these possibilities, we identified trials in which the number of potential

responses remained constant (two), but the amount of conflict varied (one or two types of conflict). Firing rates were significantly higher ($P = 6.4 \times 10^{-3}$, Mann–Whitney test) in higher-conflict trials, indicating that dACC neurons were encoding conflict itself rather than the potential number of responses (Supplementary Fig. 4). RTs for the higher-conflict trials were also significantly higher ($P = 1.5 \times 10^{-4}$, *t*-test), providing behavioural evidence for the increase in perceived conflict. In a two-way ANOVA (with degree of conflict as one variable and the number of possible responses as the other) including all trials, the degree of conflict was a significant independent predictor of firing rate ($P = 5.7 \times 10^{-3}$), whereas the number of possible responses was not ($P = 0.11$).

Current models of dACC function, whether predicated on conflict monitoring^{2,8,17}, reinforcement learning^{3,19} or reward-based decision making^{1,20,21}, require that future dACC activity reflect past experience, but modulation of dACC firing on the basis of recent history has not been demonstrated at the single-neuronal level. To determine whether dACC neuronal firing rates are influenced by previous activity, we separated type 0 and type 2 trials depending on whether they were immediately preceded by a trial containing interference (type 1 or 2) or not (type 0). In both cases, dACC neuronal activity increased more rapidly after the cue when the preceding trial contained interference (Fig. 3a, b). The average magnitude of the cue neuron signal was greater in trials preceded by interference. This finding also held for the entire neuronal population (Supplementary Fig. 5) and was not altered by using raw rather than normalized rates.

The association between previous and current trial activity was maintained across all successive trial pairs. On a trial-by-trial basis including all trial types, previous trial activity during the cue period was significantly correlated with current trial activity ($r = 0.15$, $P = 2.0 \times 10^{-11}$ for cue neurons; $r = 0.12$, $P = 3.3 \times 10^{-16}$ for all neurons; Supplementary Note 2), demonstrating that dACC neurons encode information about both the current task context and the recent past. This neural–neural correlation was not simply an effect of drift in the recordings (Supplementary Note 3).

The behavioural correlate of this neuronal pattern of activity depended on the identity of the current trial. RTs for type 0 trials were positively correlated with previous trial activity ($r = 0.13$,

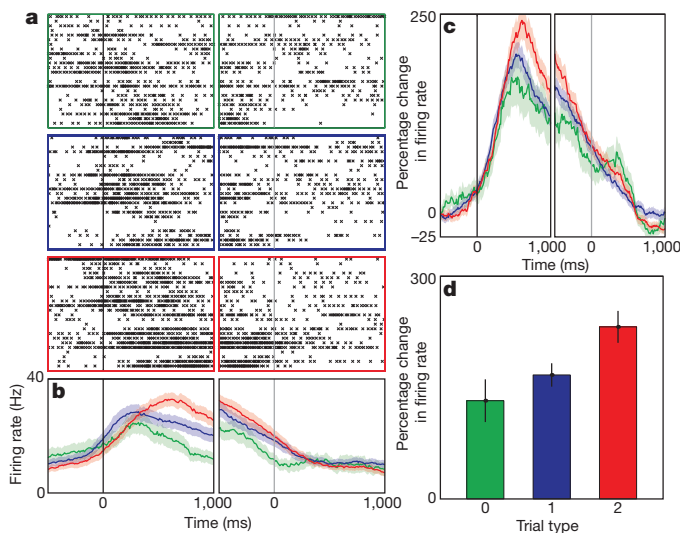


Figure 2 | Individual and population neuronal responses. **a**, Example neuron showing modulation of firing based on cue-related interference. Rasters for type 0 (green), 1 (blue) and 2 (red) trials are shown aligned to the cue (black line) and choice (grey line). **b**, Average firing rates of the same neuron. Error bars (s.e.m., $n = 72$) are depicted with shading. **c**, **d**, Firing rates (**c**) and averaged activity within a window 200 ms wide centred 500 ms after the cue (**d**) for all cue-related neurons. Neuronal firing increased with cognitive interference ($P = 0.02$, ANOVA), correlating with RT.

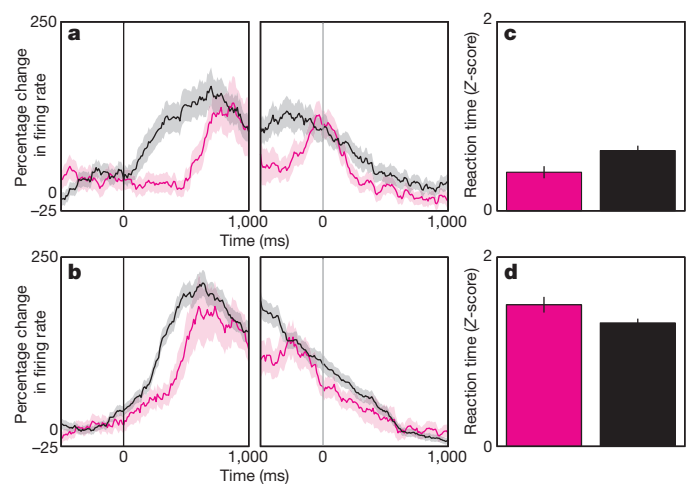


Figure 3 | Effect of previous trial on dACC firing and RT. **a**, Activity was greater for current non-interference trials immediately preceded by interference trials (1,2→0, black) than non-interference trials (0→0, purple). **b**, Similarly, activity was greater for high-interference trials preceded by interference trials (1,2→2, black) than by non-interference trials (0→2, purple). **c**, RTs for non-interference trials were shorter ($P = 0.008$) when preceded by non-interference trials (0→0, purple) than by interference trials (1,2→0, black). **d**, RTs for high-interference trials were shorter ($P = 0.04$) when preceded by interference trials (1,2→2, black) than by non-interference trials (0→2, purple). Shading and error bars represent s.e.m. ($n = 816$).

$P = 9.2 \times 10^{-3}$), meaning that previous elevated activity (consistent with a previous difficult trial) predicted a longer RT on the current non-interference trial. RTs during type 2 trials, however, were negatively correlated with previous trial activity ($r = -0.09$, $P = 8.4 \times 10^{-3}$), meaning that previous elevated activity predicted a shorter RT on the current interference trial. Taken together, these findings predict that RTs for a particular trial type will be shorter when the preceding trial is of the same type, and longer when of a different type (Supplementary Note 4).

The behavioural responses bore out these predictions. RTs during non-interference trials were shorter when preceded by another non-interference (0→0) trial than by an interference (1,2→0) trial (Fig. 3c and Supplementary Fig. 6a). Conversely, RTs during high-interference trials were shorter when preceded by an interference (1,2→2) trial than by a non-interference (0→2) trial (Fig. 3d and Supplementary Fig. 6b). This dependence of dACC neuronal activity on history provides a neurophysiological basis for the current data and for previous observations of behavioural adaptations, known as micro-adjustments, conflict adaptation or the Gratton effect^{2,4,8,17,22,23}.

We analysed post-cingulotomy task performance and thereby captured the acute behavioural manifestations of a precise, stereotyped, reproducible lesion to the previously recorded area (Supplementary Fig. 7). The error rate after cingulotomy was 1.3%, indicating that subjects had not changed in their ability to perform the task. The post-cingulotomy RT distribution was similar to that observed before the lesion (Fig. 4a and Supplementary Fig. 8a): longer RTs were associated with increasing interference ($P < 10^{-12}$, ANOVA). Furthermore, there was no difference in mean RT before and after cingulotomy ($P = 0.76$, *t*-test). Thus, the dACC lesion did not significantly disrupt the subjects' ability to perform the task, nor did it affect the dependence of RT on the cognitive load presented by the current stimulus.

Cingulotomy, however, caused abolition of the history-dependent modulation of RTs. The pre-lesion trial-to-trial adaptations in RT were significantly reduced after cingulotomy for both non-interference trials ($P = 2.2 \times 10^{-8}$, bootstrap test) and high-interference trials ($P = 7.1 \times 10^{-3}$). Consistently, the difference in RT attributable to the previous trial that existed before the lesion (Fig. 3c, d and Supplementary Fig. 6a, b) disappeared after the lesion (Fig. 4b, c and Supplementary Fig. 8b, c); that is, RTs did not depend on the preceding trial type. This effect was observable at the population level and at the level of individual subjects (Supplementary Fig. 9). Thus, although dACC lesions did not globally degrade subject performance, they eliminated the dependence of behavioural responses on recent experience.

These trial-to-trial behavioural adjustments (faster RT when the preceding trial type was of the same type as the current one) are concordant with those reported by others^{2,4,8,17,22,23}. During high-interference trials




Figure 4 | Abolition of behavioural adaptation after a targeted dACC lesion. RTs were recorded after cingulotomy, in which a stereotactic lesion was created precisely in the region of the dACC from which fMRI signals and microelectrode recordings were obtained. **a**, RTs followed a dose-response pattern by trial type ($P < 10^{-12}$, ANOVA) similar to that obtained before the lesion (Fig. 1c). Error bars represent s.e.m. ($n = 572$). **b**, **c**, However, behavioural adaptations (the influences of previous trial identity on current trial RTs) were abolished for both non-interference trials ($P = 0.54$) (**b**) and high-interference trials ($P = 0.53$) (**c**).

preceded by a high-interference trial, however, we observed increased single-neuron activity in the dACC, whereas others have reported a decreased blood-oxygen-level-dependent fMRI signal^{12,17}. This apparent discrepancy may be explained by the fact that the peak fMRI signal (which occurs 5–7 s after the cue) reflects input synaptic activity evoked by both the appearance of the cue and evaluation of the response (which all occur within the first 1 s after the cue). In contrast, we recorded output spiking activity occurring within 500–750 ms after the cue. These complementary measures may therefore reflect the spatiotemporal dynamics of conflict processing in the dACC.

In this task, increasing interference within the cue could variably be interpreted as representing increasing conflict between competing potential responses⁸, increasing likelihood of error commission³ or increasing energetic cost of decision making²⁰. Consistent with these theories, we find that dACC activity is correlated with these manifestations of cognitive demand. However, although the dACC is modulated by the cognitive load within the context of the current task, its function is not essential for generating the load-dependent behavioural response^{16,24}, because interference-dependent behaviour was not altered after dACC ablation. In contrast, an intact dACC is required for trial-to-trial behavioural adjustments.

Previous studies have proposed that the dACC monitors for conflict between competing responses^{7,8} and drives behaviour towards efficient strategies²⁰. Whereas previous human studies using fMRI^{12,17}, event-related potentials²⁵ and lesions²⁶ have implicated the dACC in this function, single-unit data supporting this theory have been lacking. Moreover, non-human primate single-unit recordings^{18,21} and lesion²⁴ studies have arrived at opposite conclusions and cast doubt on the conflict-monitoring theory²⁷. We reconcile these issues by demonstrating fMRI and single-neuronal conflict signals in the human dACC, and also behavioural adjustments that disappear after a precisely targeted dACC lesion. Our results support the view that the dACC is specifically responsible for providing a continuously updated account of predicted demand on cognitive resources. This account is particularly sensitive to relative shifts in situational complexity from instance to instance, weighted by the recent past. The salient influence of current dACC activity on future neuronal activity and behaviour permits the implementation of behavioural adjustments that optimize performance. In situations in which cognitive demands remain constant, this signal facilitates efficiency by accelerating responses. In situations involving rapidly changing demands, it promotes accuracy by retarding responses.

METHODS SUMMARY

Six human subjects (four males, ages 37.5 ± 5 years) undergoing stereotactic cingulotomy for treatment-refractory obsessive-compulsive disorder provided informed consent and enrolled. The study was approved by our Institutional Review Board. The surgical procedure produces a stereotypical lesion with an average volume of 3.58 ± 1.24 cm³, centred 9 mm lateral to the midline, 18 mm anterior to the anterior commissure, and 30 mm superior to the anterior-commissure-posterior-commissure plane²⁸.

During surgery, subjects performed the multi-source interference task^{13,14}, which was presented on a computer monitor with a customized software package in MATLAB (MathWorks)²⁹. Each trial contained a stimulus consisting of three integers from 0 to 3. One number (the unique 'target') differed from a pair of 'distracter' numbers. Subjects were asked to report, by pressing a button, the identity (rather than the position) of the target (left button for 1, middle for 2, right for 3).

Functional MRI data were analysed with Brain Voyager software (Brain Innovation). Anatomical and functional data were registered and transformed into common Talairach space. A general linear model was constructed by using predictors modelled by convolution with a standard haemodynamic response function. Single-subject repeated-measures ANOVAs were performed on a voxel-wise basis. Multiple comparisons were accounted for by using a cluster constraint with regional false-positive probability $P < 10^{-4}$, requiring a cluster of at least seven contiguous voxels with $P < 0.05$.

For microelectrode recordings, an array of three tungsten microelectrodes (500–1,500 k Ω ; FHC) was attached to a motorized microdrive (Alpha Omega

Engineering). On reaching the cingulate cortex, microelectrodes were held in place and monitored for about 5 min to assess signal stability. Putative neurons were not screened for task responsiveness. Analogue data were amplified, bandpass filtered between 300 Hz and 6 kHz, sampled at 20 kHz and spike-sorted (Offline Sorter).

Full Methods and any associated references are available in the online version of the paper.

Received 9 October 2011; accepted 17 May 2012.

Published online 24 June 2012.

- Williams, Z. M., Bush, G., Rauch, S. L., Cosgrove, G. R. & Eskandar, E. N. Human anterior cingulate neurons and the integration of monetary reward with motor responses. *Nature Neurosci.* **7**, 1370–1375 (2004).
- Botvinick, M., Nystrom, L. E., Fissell, K., Carter, C. S. & Cohen, J. D. Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature* **402**, 179–181 (1999).
- Brown, J. W. & Braver, T. S. Learned predictions of error likelihood in the anterior cingulate cortex. *Science* **307**, 1118–1121 (2005).
- Gratton, G., Coles, M. G. & Donchin, E. Optimizing the use of information: strategic control of activation of responses. *J. Exp. Psychol. Gen.* **121**, 480–506 (1992).
- Hayden, B. Y. & Platt, M. L. Neurons in anterior cingulate cortex multiplex information about reward and action. *J. Neurosci.* **30**, 3339–3346 (2010).
- Narayanan, N. S. & Laubach, M. Neuronal correlates of post-error slowing in the rat dorsomedial prefrontal cortex. *J. Neurophysiol.* **100**, 520–525 (2008).
- Botvinick, M. M., Cohen, J. D. & Carter, C. S. Conflict monitoring and anterior cingulate cortex: an update. *Trends Cogn. Sci.* **8**, 539–546 (2004).
- Carter, C. S. & van Veen, V. Anterior cingulate cortex and conflict detection: an update of theory and data. *Cogn. Affect. Behav. Neurosci.* **7**, 367–379 (2007).
- Carter, C. S. *et al.* Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science* **280**, 747–749 (1998).
- Nieuwenhuis, S., Schweizer, T. S., Mars, R. B., Botvinick, M. M. & Hajcak, G. Error-likelihood prediction in the medial frontal cortex: a critical evaluation. *Cereb. Cortex* **17**, 1570–1581 (2007).
- Davis, K. D., Hutchison, W. D., Lozano, A. M., Tasker, R. R. & Dostrovsky, J. O. Human anterior cingulate cortex neurons modulated by attention-demanding tasks. *J. Neurophysiol.* **83**, 3575–3577 (2000).
- Davis, K. D. *et al.* Human anterior cingulate cortex neurons encode cognitive and emotional demands. *J. Neurosci.* **25**, 8402–8406 (2005).
- Bush, G. & Shin, L. M. The Multi-Source Interference Task: an fMRI task that reliably activates the cingulo-frontal-parietal cognitive/attention network. *Nature Protocols* **1**, 308–313 (2006).
- Bush, G., Shin, L. M., Holmes, J., Rosen, B. R. & Vogt, B. A. The Multi-Source Interference Task: validation study with fMRI in individual subjects. *Mol. Psychiatry* **8**, 60–70 (2003).
- Bush, G. *et al.* Functional magnetic resonance imaging of methylphenidate and placebo in attention-deficit/hyperactivity disorder during the multi-source interference task. *Arch. Gen. Psychiatry* **65**, 102–114 (2008).
- Fellows, L. K. & Farah, M. J. Is anterior cingulate cortex necessary for cognitive control? *Brain* **128**, 788–796 (2005).
- Kerns, J. G. *et al.* Anterior cingulate conflict monitoring and adjustments in control. *Science* **303**, 1023–1026 (2004).
- Nakamura, K., Roesch, M. R. & Olson, C. R. Neuronal activity in macaque SEF and ACC during performance of tasks involving conflict. *J. Neurophysiol.* **93**, 884–908 (2005).
- Holroyd, C. B. & Coles, M. G. The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol. Rev.* **109**, 679–709 (2002).
- Botvinick, M. M. Conflict monitoring and decision making: reconciling two perspectives on anterior cingulate function. *Cogn. Affect. Behav. Neurosci.* **7**, 356–366 (2007).
- Ito, S., Stuphorn, V., Brown, J. W. & Schall, J. D. Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science* **302**, 120–122 (2003).
- Mayr, U., Awh, E. & Laurey, P. Conflict adaptation effects in the absence of executive control. *Nature Neurosci.* **6**, 450–452 (2003).
- Ridderinkhof, K. R. Micro- and macro-adjustments of task set: activation and suppression in conflict tasks. *Psychol. Res.* **66**, 312–323 (2002).
- Mansouri, F. A., Buckley, M. J. & Tanaka, K. Mnemonic function of the dorsolateral prefrontal cortex in conflict-induced behavioral adjustment. *Science* **318**, 987–990 (2007).
- Gehring, W. J. & Fencsik, D. E. Functions of the medial frontal cortex in the processing of conflict and errors. *J. Neurosci.* **21**, 9430–9437 (2001).
- di Pellegrino, G., Ciaramelli, E. & Ladavas, E. The regulation of cognitive control following rostral anterior cingulate cortex lesion in humans. *J. Cogn. Neurosci.* **19**, 275–286 (2007).
- Cole, M. W., Yeung, N., Freiwald, W. A. & Botvinick, M. Cingulate cortex: diverging data from humans and monkeys. *Trends Neurosci.* **32**, 566–574 (2009).
- Rauch, S. L. *et al.* Volume reduction in the caudate nucleus following stereotactic placement of lesions in the anterior cingulate cortex in humans: a morphometric magnetic resonance imaging study. *J. Neurosurg.* **93**, 1019–1025 (2000).
- Asaad, W. F. & Eskandar, E. N. A flexible software tool for temporally-precise behavioral control in Matlab. *J. Neurosci. Methods* **174**, 245–258 (2008).

Supplementary Information is linked to the online version of the paper at www.nature.com/nature.

Acknowledgements This work was supported by grants from the National Science Foundation (IOB 0645886), the National Institutes of Health (NEI 1R01EY017658-01A1, NIDA 1R01NS063249, NIMH Conte Award MH086400 and R25 NS065743), the Klingenstein Foundation, the Howard Hughes Medical Institute, the Sackler Scholar Programme in Psychobiology, the Centers for Disease Control (5R01DP000339), the Benson-Henry Institute at Massachusetts General Hospital for Mind–Body Medicine, the David Judah Fund, the McIngvale Fund, and the Center for Functional Neuroimaging Technologies (P41RR14075).

Author Contributions E.N.E., W.F.A., Z.M.W. and D.D.D. designed the study. G.B. administered and interpreted the fMRI scans. E.N.E. performed the surgical procedures, and S.A.S., M.K.M., S.R.P. and W.F.A. obtained the neuronal recordings. S.A.S. and M.K.M. analysed the data and wrote the manuscript. All authors edited the manuscript.

Author Information Reprints and permissions information is available at www.nature.com/reprints. The authors declare no competing financial interests. Readers are welcome to comment on the online version of this article at www.nature.com/nature. Correspondence and requests for materials should be addressed to E.N.E. (eeskandar@partners.org).

METHODS

Subjects. We enrolled six study subjects (four male, ages 37.5 ± 5 years (mean \pm s.e.m.)) undergoing stereotactic cingulotomy for treatment-refractory obsessive-compulsive disorder. Evaluation for surgical candidacy was conducted by a multidisciplinary team consisting of psychiatrists, neurologists and neurosurgeons. Subjects enrolled voluntarily, providing informed consent under a protocol approved by the Massachusetts General Hospital Institutional Review Board. The surgical procedure produces a stereotypical lesion with an average volume of 3.58 ± 1.24 cm³ (mean \pm s.d.), centred 9 mm lateral to the midline, 18 mm anterior to the anterior commissure, and 30 mm superior to the anterior-commissure-posterior-commissure plane²⁸. Subject participation was in no way related to clinical decision-making regarding their candidacy for surgery.

Behavioural task. During the microelectrode recording portion of surgery, subjects performed the multi-source interference task (MSIT) (Fig. 1a, b)^{13,14}. The task was presented on a computer monitor using a customized software package in MATLAB (MathWorks)²⁹. Each trial contained a stimulus consisting of three integers from 0 to 3. One number (the unique 'target') differed from a pair of 'distracter' numbers (for example 100 or 323). Subjects were asked to report, by pressing a button, the identity (rather than the position) of the target (left button for 1, middle for 2, right for 3).

Functional MRI. Functional MRI was performed before surgery by using a 3.0-T scanner (Allegra, Siemens AG) and head coil. The MSIT was presented on a screen visible by means of a tilted mirror, and controlled with MacStim 2.6 software (WhiteAnt Occasional Publishing). Scans were acquired with the following specifications: 15 coronal sections, 64×64 matrix, 3.125 mm² in-plane resolution, 5 mm thickness with 0 mm skip, 30 ms echo time, 1,500 ms repetition time, 90° flip angle, 20 cm² field of view. During fMRI, only types 0 and 2 trials (no interference and both types of interference; see Fig. 1b) were used. The task was run in a block design. Each block consisted of 24 trials of the same type. One run consisted of eight alternating blocks, with an additional five visual fixation blocks interspersed. Data were analysed with Brain Voyager software (Brain Innovation). Anatomical and functional data were registered and transformed into common Talairach space. A general linear model was constructed by using predictors modelled by convolution with a standard haemodynamic response function. Single-subject repeated-measures ANOVAs were performed on a voxel-wise basis. Multiple comparisons were accounted for by using a cluster constraint with regional false-positive probability $P < 10^{-4}$. This constraint required a cluster of at least seven contiguous voxels with $P < 0.05$.

Surgical procedure. The surgical procedure was performed with standard stereotactic techniques. A Cosman-Roberts-Wells (Integra) stereotactic frame was affixed to the patient under local anaesthesia, and a high-resolution MRI was obtained. The target for the left posterior lesion (2 cm posterior to the most anterior point of the frontal horn of the lateral ventricle, 0.7 cm lateral to midline, and 0.5 cm superior to the corpus callosum) was programmed into a neuro-navigation computer (StealthStation, Medtronic) and the stereotactic frame was then set appropriately. The patient was positioned semi-recumbent, the surgical area was prepared, and sterile drapes were applied. Local anaesthetic was infiltrated, a coronal skin incision was performed, and bilateral burr holes were drilled 1.5 cm lateral to the midline and 10.0 cm posterior to the nasion. A computerized microelectrode drive controlled by a neurophysiology system (Alpha Omega) was affixed to the frame. After dural opening, microelectrodes were lowered using the computerized drive in increments of 0.01 mm. The position of the tip of the electrodes was also monitored in real time using the stereotactic neuronavigation system. After microelectrode recordings, a thermocoagulation electrode with a 10-mm exposed tip (Cosman Medical) was lowered to the target. Lesions were performed by heating the electrode to 85 °C for 60 s. Two more pairs of lesions were then created, each 7 mm anterior and 2 mm inferior to the previous lesion.

Microelectrode recordings. For microelectrode recordings, an array of three tungsten microelectrodes (500–1,500 k Ω ; FHC) was attached to a motorized microdrive (Alpha Omega Engineering). As per routine surgical protocol, recordings were obtained from the left hemisphere. On reaching the cingulate cortex, microelectrodes were held in place and monitored for about 5 min to assess signal stability. Putative neurons were not screened for task responsiveness. Analogue data were amplified, bandpass filtered between 300 Hz and 6 kHz, sampled at 20 kHz (Alpha Omega Engineering) and spike-sorted (Offline Sorter, Plexon).

Post-lesion behaviour. Immediately after creation of the cingulate lesions, subjects again performed the MSIT. These task sessions were identical in all respects to the pre-lesion sessions except that we collected only behavioural data. Four of the six patients participated in these post-lesion sessions.

RT data analysis. RT was defined as the interval between the onset of the cue and the subject's button-press. To allow for comparisons between subjects, we normalized subject RTs relative to their individual distributions. Normalized RTs were defined as the number of standard deviations (*Z*-scores) above or below the bottom tenth centile of a subject's RT distribution. The choice of each subject's reference point for the normalization (bottom tenth centile versus median or mean) is arbitrary and does not affect the result, because it simply represents a rigid translation of the normalized values. We chose the bottom tenth centile so that values were positive, to facilitate visual comparisons.

We also computed a speed of target selection (STS) as described in ref. 24. We inverted RT to obtain $STS = 1/RT$. We then divided by the mean STS across all trials in the behavioural session. STS was thus normalized to 1, which represented an intermediate or 'average' response speed. Faster responses (for example, on low-conflict trials) are reflected as relative increases in STS (that is, $STS > 1$), whereas slower responses (for example, on high-conflict trials) are attended by decrements in the STS (that is, $STS < 1$). Use of the normalized STS, rather than raw RTs, allowed us to view data from different subjects on a common scale and to minimize between-subject differences when pooling data to calculate composite statistics.

Single-unit data analysis. Single units were isolated by first building a histogram of peak heights from the raw voltage tracings on each channel. We applied a minimum threshold of three standard deviations to exclude background noise. Action potentials were sorted by using waveform principal component analysis. Spike clusters of putative neurons were required to separate clearly from any channel noise, to demonstrate a voltage waveform consistent with that of a cortical neuron and to have at least 99% of action potentials separated by an inter-spike interval of 1 ms or more.

We recorded an average of 1.1 neurons per microelectrode. When a single channel captured more than one neuron cluster, we required a clear distinction between the two clusters to include either one as a single unit ($P < 0.01$, multivariate ANOVA across the first two principal components). In addition, we required putative dACC neurons to fire at an average rate of at least 1.0 spikes per second, to be stably active for at least 25 task trials, and to not demonstrate significant drift over the duration of the recording. We excluded single units not meeting these criteria. We did not use any multi-unit activity.

We classified neurons into three mutually exclusive categories on the basis of the timing of their peak firing rates with respect to task events. Neurons peaking in activity before stimulus presentation, during the stimulus period and after the button-push were classified as 'pre-cue', 'cue' and 'feedback' neurons, respectively.

To facilitate comparisons between neurons with different firing rates, we normalized (divided) neuronal activity by the average neuronal firing rate during a 500-ms window preceding the appearance of the fixation point. Population firing rates were computed by averaging these normalized neuronal responses with a 250-ms moving boxcar window.