

The Neural Representation of Voluntary Task-Set Selection in Dynamic Environments

David Wisniewski^{1,2,5,6,*}, Carlo Reverberi^{1,8,9,*}, Anita Tusche^{1,10,11} and John-Dylan Haynes^{1,2,3,4,5,6,7}

¹Bernstein Center for Computational Neuroscience Berlin, ²Berlin Center for Advanced Neuroimaging, ³Cluster of Excellence NeuroCure, ⁴Department of Neurology, Charité – Universitätsmedizin Berlin, Berlin, Germany, ⁵SFB 940 Volition and Cognitive Control, Technische Universität Dresden, Dresden, Germany, ⁶Berlin School of Mind and Brain, ⁷Department of Psychology, Humboldt-Universität zu Berlin, Berlin, Germany, ⁸Psychology Department, Università Milano-Bicocca, Milan, Italy, ⁹Milan Center for Neuroscience, Milan, Italy, ¹⁰Department of Social Neuroscience, Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany and ¹¹Division of the Humanities and Social Sciences, California Institute of Technology, CA, USA

*David Wisniewski and Carlo Reverberi contributed equally to this work.

Address correspondence to David Wisniewski, Charité – Universitätsmedizin Berlin, Bernstein Center for Computational Neuroscience, Haus 6, Philippstrasse 13, 10115 Berlin, Germany. Email: david.wisniewski@bccn-berlin.de

Downloaded from https://academic.oup.com/cercor/article/25/12/4715/3302214 by guest on 03 November 2018

When choosing actions, humans have to balance carefully between different task demands. On the one hand, they should perform tasks repeatedly to avoid frequent and effortful switching between different tasks. On the other hand, subjects have to retain their flexibility to adapt to changes in external task demands such as switching away from an increasingly difficult task. Here, we developed a difficulty-based choice task to investigate how subjects voluntarily select task-sets in predictably changing environments. Subjects were free to choose 1 of the 3 task-sets on a trial-by-trial basis, while the task difficulty changed dynamically over time. Subjects self-sequenced their behavior in this environment while we measured brain responses with functional magnetic resonance imaging (fMRI). Using multivariate decoding, we found that task choices were encoded in the medial prefrontal cortex (dorsomedial prefrontal cortex, dmPFC, and dorsal anterior cingulate cortex, dACC). The same regions were found to encode task difficulty, a major factor influencing choices. Importantly, the present paradigm allowed us to disentangle the neural code for task choices and task difficulty, ensuring that activation patterns in dmPFC/dACC independently encode these 2 factors. This finding provides new evidence for the importance of the dmPFC/dACC for task-selection and motivational functions in highly dynamic environments.

Keywords: fMRI, multivariate decoding, prefrontal cortex, task difficulty, task-set

Introduction

Despite the staggering complexity of the environments that humans are faced with every day, we are able to reach desired goals efficiently and with the limited resources available (Shallice and Burgess 1991; Duncan 2001; Miller and Cohen 2001). This ability requires highly flexible task selection processes that balance different competing task demands.

On the one hand, some factors favor repeated performance of the same task. For instance, if a task has a low difficulty or if it is highly rewarded, it will be chosen again. Numerous studies on reward-based decision-making (e.g., Kennerley et al. 2006; Boorman et al. 2013) associate reward-based task choices with areas such as the dorsal anterior cingulate cortex (dACC, Hampton and O'Doherty 2007; Hayden, Heilbronner, et al. 2011). Furthermore, switching between tasks or task-sets is associated with costs due to task-reconfiguration processes (Monsell, 2003). Subjects that are asked to choose between

tasks voluntarily tend to avoid frequent switching and stay in the same task (Arrington and Logan 2004), likely due to the associated switch costs.

On the other hand, some factors favor flexibly switching between tasks. The conditions under which tasks are performed might change quickly in dynamic environments, creating the necessity for adapting to these changing circumstances. For example, the difficulty of a task might increase or the reward outcome might decrease. This will lead to switching away from the currently performed task in favor of finding an alternative, potentially easier one (Daw et al. 2006; Cohen et al. 2007), a process that has also been associated with neuronal responses in the dACC (Hayden, Pearson, et al. 2011).

Task selection processes are therefore influenced both by factors that favor repeated performance of the same task and by factors that favor switching flexibly between alternative tasks (Müller et al. 2007). Here, we investigated how subjects select tasks when these opposing demands need to be balanced in a dynamically changing environment. We developed a difficulty-based choice task where subjects chose 1 of the 3 task-sets on a trial-by-trial basis, while their brain responses were measured using functional magnetic resonance imaging (fMRI). The choices were influenced by the task difficulty that changed dynamically over time. Importantly, in contrast to previous studies (Hampton and O'Doherty 2007), changes in task difficulty were also influenced interactively by the choices subjects made—a key feature of dynamic naturalistic environments. We used multivariate pattern analysis (MVPA) methods in order to investigate which brain areas contained information about choices and the main variable influencing these choices, task difficulty. Notably, our paradigm allowed us to independently assess and compare the brain areas encoding task choices and task difficulty. In line with previous fMRI results (Forstmann et al. 2006; Hampton and O'Doherty 2007; Haynes et al. 2007) and recent theories (Pessoa 2009; Holroyd and Yeung 2012), we found that medial prefrontal brain regions, including the dACC, play a central role in voluntary task selection in such dynamic environments. A cluster in the dorsomedial prefrontal cortex (dmPFC), encompassing parts of the dACC, was the only region that encoded both task choices and the main variable influencing these choices, task difficulty. This finding emphasizes the role of the dmPFC/dACC in decision-making in dynamic environments and its role in linking motivational functions and goal-directed behavior (e.g., Holroyd and Yeung 2012; Shenhav et al. 2013).

Materials and Methods

Participants

Nineteen subjects took part in the experiment (10 females, mean age: 26 years). All subjects volunteered to participate and had normal or corrected-to-normal vision. Subjects gave written informed consent and were paid 35€ for participation. The experiment was approved by the local ethics committee. One subject was excluded from the sample due to a structural brain anomaly. All other subjects had no history of neurological disorders or structural brain anomalies. Data of 2 further subjects had to be discarded, due to excessive head movement during scanning (>15 mm, all other subjects <3 mm) and exceedingly high error rates (34.5%, >8 SD from mean error rate across all subjects, which was 5.48%), respectively.

Overview

Subjects chose freely between 3 different task-sets on a trial-by-trial basis while the task difficulty changed dynamically over time. Difficulty increased when the same task-set was chosen repeatedly, and generally decreased when a task-set was not chosen. Difficulty changes were therefore directly influenced by the subjects' choices. Within this dynamic environment, subjects needed to balance competing task demands: repeatedly performing the same task-set in order to exploit a comparatively low difficulty level and to avoid switch costs, and flexibly switching between task-sets in order to find the task-set with the lowest difficulty at any given point in time. Subjects were instructed to perform with as few errors as possible.

Experimental Paradigm

The experiment was implemented using Matlab Version 7.11.0 (The MathWorks) and the Cogent Toolbox (<http://www.vislab.ucl.ac.uk/cogent.php>). Every trial consisted of 2 parts: task execution and the selection of the next task. Task execution started with a fixation cross presented for 500 ms, followed by a target stimulus displayed for 3500 ms (Fig. 1A). The target was a picture of an object, drawn from 1 of the 3 categories: musical instruments, means of transportation, or furniture. Each category contained 3 objects, and the presentation order of objects was pseudorandomized. Objects were presented in 9 different difficulty levels (Fig. 1B). Difficulty was varied by adding different amounts of independent and identically distributed Gaussian noise to the pictures. Independent behavioral piloting data ensured that each object was clearly visible at the lowest difficulty level and was invisible at the highest difficulty level. Piloting data also ensured that the difficulty manipulation was comparable over all object categories. Thus, no category was easier or harder than the others, controlling for low-level perceptual confounds of different visibility in the data. Below the object, 4 colored squares (red, green, blue, and gray) were presented at 4 fixed positions (Fig. 1C). Each position was assigned to 1 of the 4 buttons, which were operated with the index and middle fingers of both hands. Subjects were instructed to perform 1 of the 3 task-sets. Each task-set consisted of rules linking the object categories with colors (Fig. 1C): for example, task-set 1 was "If you see means of transportation, press the button at the red position. If you see a piece of furniture, press the button at the green position. If you see a musical instrument, press the button at the blue position." The second and third rule-sets were permutations of this color-category mapping. Note that the gray button was never assigned to a category throughout the experiment. It was merely included to balance left-handed and right-handed button presses. The assignment of the 4 colors to the 4 positions was pseudorandomized in each trial, to avoid motor preparation of responses. Participants were instructed to react as quickly and accurately as possible. The target stayed on screen for 3500 ms. Then, a blank screen was presented for 4000 ms.

A choice screen was presented next for selection of the next task-set for 2000 ms. Here, subjects could freely choose which task-set to perform subsequently. Each task-set was assigned a number that subjects learned prior to scanning, and these numbers were presented at fixed positions on screen, which were assigned to a button using a pseudorandomized mapping in each trial. Finally, a blank screen was presented for either 2000 or 4000 ms before the next task execution screen was presented.

Overall, subjects performed 306 trials during scanning, divided into 6 runs of 51 trials each. The whole experimental session lasted 85 min on average. Subjects also performed a training session outside the scanner 2–3 days prior to scanning. They were given 2 h to memorize the rules, get acquainted with the whole range of difficulty levels, and develop a choice strategy. This ensured that subjects were familiar with the dynamic environment, understood how their choices influenced that environment and that learning effects were minimized during scanning.

Task Difficulty

The manipulation of task difficulty was central to this paradigm. Each run started with a task execution screen in task-set 1 in the lowest difficulty level. The range of difficulty levels was from 1 (easiest) to 9 (hardest). Given that subjects did not choose the task-set in the first trial, that trial was excluded from all further analyses. Task choices influenced how difficulty changed using an algorithm that worked as follows: (1) The difficulty of the chosen task-set always increased by discouraging subjects from staying in a task-set across many trials. (2) The difficulty change of both non-chosen task-sets depended on the current difficulty level. The higher the current task difficulty, the stronger the decrease of difficulty for the non-chosen task-sets, with a maximum decrease of 3 levels. The only exception to this rule was if the current difficulty was at the lowest level. Here, in order to discourage subjects from switching constantly between task-sets, the difficulty of the non-chosen task-sets was set to increase faster than for the chosen task-set. Random noise was added to the difficulty changes for the non-chosen tasks, in order to prevent subjects from choosing task sets in a fixed sequence (e.g., task-set 1, task-set 2, task-set 3, task-set 1, task-set 2 etc.). Changes in the difficulty level thus depended on both the current choice and the current state of the environment. For more detailed information on the difficulty change algorithm, see Supplementary Figure 1. To keep the overall difficulty low, subjects needed to track and use several pieces of information. The difficulty level of the current task-set would help them to decide whether to switch or not. The history of the 2 alternative task-sets (i.e., when were they last performed and how difficult were they then?) would help them to decide which task-set to choose in switch trials, since the longer a task-set has not been performed the more likely it was at a low difficulty level. Combined with an understanding of the dynamics of the environment, subjects could then balance the benefits of staying in a task-set (exploit a low difficulty, low switch costs) and the benefits of switching flexibly between task-sets (potentially lower difficulty of alternative task-sets, but switch costs). Note that subjects were only given coarse instructions on how difficulty changed across trials. They were told that difficulty will increase if a task is repeated, and that non-chosen tasks will generally, but not always, decrease in difficulty. It was up to the subjects to develop a more precise model of the task during the 2 h training session.

We used 3 different task-sets in order to investigate task selection independent of task switching. Switching away from a task-set did not determine what would be the next task-set given that there were always 2 options after a switch. And given the similarity of the task-sets, we avoided the emergence of stable preferences, which would have confounded task choice analyses.

Image Acquisition

Functional imaging was conducted on a 3-T Siemens Trio (Erlangen, Germany) scanner with a 12-channel head coil. For each run, 304 T_2^* -weighted echo-planar images (EPIs) were acquired (time repetition = 2000 ms, time echo = 30 ms, flip angle 90°). Each volume consisted of 32 slices, separated by a gap of 0.75 mm. Matrix size was 64 × 64, and field of view was 192 mm, which resulted in an in-plane resolution of 3 mm². Voxel size was therefore 3 × 3 × 3.75 mm. The first 3 images of each run were discarded to account for magnetic susceptibility artifacts.

Data Analysis

Behavior

Behavioral data were analyzed using Matlab (Version 7.11.0) and SPSS (Release 19.0.0). For each subject, basic task performance was assessed by calculating mean reaction times (RTs) and mean error rates across

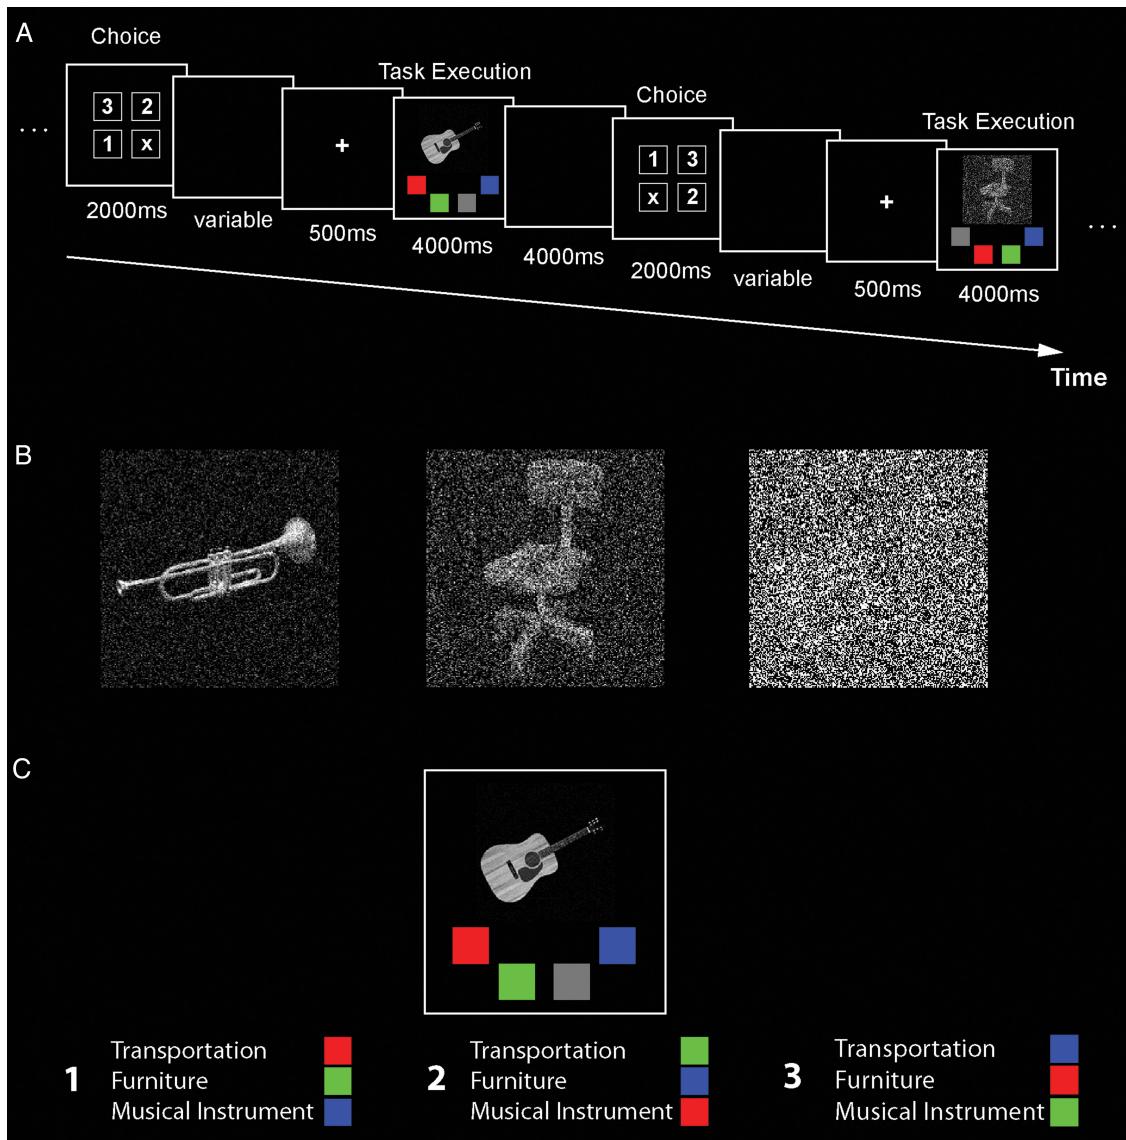


Figure 1. Experimental paradigm. (A) Trial structure. Subjects were presented with randomly chosen visual objects from 1 of the 3 categories. Their task was to apply 1 of the 3 different rule-sets that mapped object categories to colors and to press the corresponding button. Subjects freely chose which rule-set to apply in every trial. (B) Sample stimuli. Objects were presented in 9 different difficulty levels. Difficulty was varied by adding independent and identically distributed Gaussian noise to the pictures, which resulted in a stepwise reduction in visibility. (C) Task execution screen. An object was presented centrally with 4 colored squares below it. Each square position corresponded to one button, with the color-button mapping changing randomly between trials. Only red, green, and blue squares were task relevant, the gray square was included to balance left and right button presses and was never associated with a category.

all runs. Moreover, we analyzed subjects' choices and the factors influencing them. Task choices could be conceptualized as consisting of 2 parts: deciding to switch (to a different task-set) or stay (in the same task-set), and in case of a switch choice also choosing 1 of the 2 currently non-chosen task-sets. The current task difficulty was assumed to be one of the crucial factors influencing the decision to switch or stay. Subjects should have been more likely to switch in difficult compared with easy trials. To test this hypothesis, we ran a linear regression analysis using difficulty as the independent and the probability to switch to another task-set as the dependent variable.

As mentioned above, subjects also had to consider their choice history during task-set selection. This was due to the fact that the longer a task-set was not chosen and performed, the more likely it was at a low difficulty level. This information could have been used to improve performance by choosing the task-set that has not been performed for a longer time in switch trials. To test this assumption, we compared the mean number of trials that passed since the chosen

task-set and the non-chosen task-set have been performed last. If subjects indeed used their choice history to make decisions, the number of trials for the chosen task-set should be larger than for the non-chosen task-set. Note that choice history could also have been used in previous studies to maximize outcomes (e.g., [Hampton and O'Doherty 2007](#); [Boorman et al. 2009](#)). Contrary to the present task, the environment in these studies changed unpredictably and was not as directly affected by subjects' choices.

Neuroimaging

Functional data analysis was performed using SPM8 (<http://www.fil.ion.ucl.ac.uk/spm>), unless stated otherwise. We first unwarped, realigned, and slice time corrected all volumes. Preprocessed data were then entered into a general linear model (GLM; [Friston et al. 1994](#)). Using MVPA methods ([Haynes and Rees 2006](#); [Kriegeskorte et al. 2006](#)), we performed 2 independent analyses to investigate the neural encoding of task choices and the encoding of task difficulty, respectively.

Neuroimaging Analysis of Choices

Quality Checks

In a first step, we analyzed the neural encoding of task choices. As subjects chose freely which task-set to perform, we only had limited control over which task-sets were chosen throughout a run. It was therefore required to run quality checks in order to ensure that we could run a choice analysis in each subject. For instance, although very unlikely, subjects could have chosen to perform only task-set 1 throughout the whole run. We thus first checked whether there were a sufficient number of trials for each task choice in each run, and excluded a run from the analysis if any regressor could only be estimated from fewer than 6 trials. Moreover, we excluded subjects in which fewer than 5 runs remained after that minimal-trial-number check, to ensure reliable run-wise cross-validation for support vector classification (see below). In fact, all runs from all subjects passed this basic quality check.

We then tested whether task choices were influenced by the current task difficulty. Due to the fact that we used 3 instead of 2 task-sets, the choice to switch or stay was not equivalent to task choices, *and we only expected difficulty to influence the switch/stay aspect of choices*. Yet, to rule out the possibility of residual correlations between these 2 factors, we explicitly tested whether we could predict task choices from the current task difficulty. We split the behavioral dataset into a training dataset (5 runs) and a test dataset (remaining run) for each subject. We then fitted a multinomial logistic regression function to the average difficulty level of trials in which the subject chose task-set 1, task-set 2, and task-set 3. We computed the same values for the remaining independent test run and used the multinomial logistic regression function to predict the data from that test run. This procedure was repeated 6 times, always using a different run as the test dataset. The average accuracy across predictions in each cross-validation step was computed and tested against chance level (33%) using a one sample *t*-test. This procedure was specifically designed to mimic the run-wise cross-validation approach used in the searchlight decoding analysis described below. We reasoned that if there were potential confounds in the searchlight decoding, they should be clearly visible in this behavioral analysis given that the analysis approach was highly similar and behavioral data are usually less noisy than fMRI data. Two subjects showed significant predictions ($P < 0.05$) and were excluded from further analyses to ensure that the neural encoding of task choices was independent from the current difficulty level. An analysis of variance (ANOVA) confirmed the bias in 1 of the 2 subjects, $F_{2,5} = 4.32$, $P < 0.05$, and was therefore slightly less sensitive than our proposed approach.

GLM and SVC

To examine the neural encoding of choices, for each subject, a GLM was used to estimate 3 regressors of interest in each run: choices of task-set 1, task-set 2, and task-set 3 during the task selection phase. The regressors were convolved with the canonical hemodynamic response function (HRF). Task selection was modeled as an event at the time point when subjects had all the necessary information to make a choice, at the onset of the task execution screen just prior to the choice screen. We chose this time point instead of the choice screen because at the choice screen onset the decision was likely already made. Subjects might have made a choice at any point in time between the task screen onset and the button press indicating their choice (RT_{Choice}). To ensure that we captured all relevant cognitive processes, we locked the onset of the HRF function to the earliest possible point in time where choice processes could have started. Given the long training and subjects' experience with the paradigm, it seems very likely that they indeed made their choices at the earliest possible point in time. Moreover, in a previous study investigating choice representations using MVPA (Hampton and O'Doherty 2007), the authors also chose a similar approach. They decoded choices from brain activity close to the earliest point in time at which a choice could have been made in order to ensure relevant decision-related activity was captured. Note that our analysis was independent of motor preparation processes, as response mappings on the choice screen were pseudorandomized across trials.

For each subject, we then applied multivariate pattern classification using a support vector classifier (SVC) with a linear kernel and a fixed

regularization parameter ($C = 1$) on the parameter estimates of the GLM (Cox and Savoy 2003; Mitchell et al. 2004; Kamitani and Tong 2005; Haynes and Rees 2006), as implemented in LIBSVM (<http://www.csie.ntu.edu.tw/~cjlin/libsvm>). More precisely, we applied a searchlight decoding approach (Kriegeskorte et al. 2006; Haynes et al. 2007), which makes no a priori assumptions about informative regions. We first defined a sphere with a radius of 4 voxels around each measured voxel in the acquired volumes. For each condition (choices of task-set 1, task-set 2, or task-set 3), we extracted parameter estimates for each of the N voxels in the given sphere, thus yielding an N -dimensional pattern vector. This was done for each run independently. Pattern vectors from 5 of the 6 runs (training dataset) were then used to train the SVC to discriminate activation patterns of the 3 conditions. The classification performance was then tested using the remaining independent run (test dataset). We repeated this procedure 6 times with every run being the test dataset once and therefore achieved a 6-fold cross-validation. Splitting the dataset into training and test datasets and run-wise cross-validation was necessary to control for potential problems of overfitting. We then calculated the mean prediction accuracy across the cross-validation steps and assigned this value to the central voxel of the sphere. The classification was repeated for every sphere in the measured brain volume, resulting in a three-dimensional accuracy map for each subject. The resulting accuracy maps were then normalized to a standard brain (Montreal Neurological Institute [MNI] EPI template as implemented in SPM8) and resampled to an isotropic resolution of $3 \times 3 \times 3$ mm. Normalized images were smoothed with a Gaussian kernel with 6 mm full-width at half-maximum, to account for differences in localization across subjects. The accuracy maps of each subject were then entered into a random-effects group analysis and statistically tested using voxel-wise one sample *t*-tests against chance level. As the SVC was performed on 3 task choice conditions, chance level was 33%. We applied a statistical threshold of $P < 0.001$, corrected for multiple comparisons at the cluster level (FWE, $P < 0.05$).

Neuroimaging Analysis of Difficulty

Quality Checks

Next, we analyzed the neural encoding of a major factor influencing choices, the current task difficulty. We expected task difficulty to be strongly correlated with the decision to switch or stay. To ensure that difficulty decoding results did not merely reflect processes related to task switching, we restricted this analysis to stay trials. Given a high probability of staying in trials with a low difficulty, this analysis was performed on the easiest 2 difficulty levels only. In all higher difficulty levels, we had too few (< 6 per run) stay trials to include them into the analysis, as parameter estimates would be too unreliable with so few trials. We then ran the same quality checks and applied the same minimal-trial-number criterion as in the choice decoding analysis on the easiest 2 difficulty levels. Two subjects had to be excluded from the analyses due to an insufficient number of trials (< 6 per run).

GLM and SVC

To investigate the neural encoding of task difficulty, functional data of each subject were analyzed using a GLM that modeled 2 conditions of interest: trials with the difficulty level 1, and trials with the difficulty level 2. Parameter estimates were convolved with the canonical HRF. Task difficulty was modeled as an event at the onset of the task execution screen, similar to the choice decoding analysis. Note that this analysis was independent of motor preparation processes as the response mapping on the task execution screen was pseudorandomized across trials. We then applied a searchlight decoding approach similar to the choice decoding described above, to identify areas that encode the current difficulty level. The classification of difficulty was based on 2 alternatives, resulting in a chance level of 50%. Again, a statistical threshold of $P < 0.001$, corrected for multiple comparisons at the cluster level (FWE, $P < 0.05$), was applied at the group level to identify areas that encoded task difficulty across subjects.

Neuroimaging Analysis: Convergence of Information

Due to the fact that we were able to independently assess task choice and difficulty representations in the brain, we could also address the issue of whether choice and difficulty information overlapped in some brain regions. Whenever we make decisions that are influenced by an environmental factor (e.g., rewards or difficulty), our brain should represent that factor as well as the ensuing decision itself. One might speculate that there might be a brain area that has access to both pieces of information and might even be able to integrate them to guide behavior. We can partly assess this notion in our experiment. One necessary condition that such a brain area needs to meet is that it has access to information about the environmental factor and the choice, a hypothesis that we can test using decoding methods. In our paradigm, such a brain area should encode both task choices and task difficulty. We therefore tested for a convergence of information using small volume correction (Worsley et al. 1996). Firstly, we extracted the areas encoding choices from the group level choice decoding analysis and then assessed whether those areas also contained information about difficulty, as assessed in an independent decoding analysis. Secondly, we extracted the areas encoding difficulty from the group level difficulty decoding analysis and assessed whether the same areas also contained information about choices, again assessed in an independent decoding analysis.

Results

Behavioral Results

Task Performance

Subjects performed a difficulty-based choice task in which they chose freely between 3 task-sets on a trial-by-trial basis in a dynamic environment. Subjects needed to balance different competing task demands in order to keep difficulty at a low level.

Each trial consisted of 2 phases: (1) task execution, in which subjects saw an object and applied a rule-set, and (2) task selection, in which subjects indicated which rule-set they wanted to apply in the next trial. Thus, 2 RTs were measured in each trial: The time to apply a rule-set to the stimulus on the task execution screen (RT_{Exec}) and the time to select a task-set at the choice screen (RT_{Choice}). We averaged the respective RTs over all trials with a difficulty from 1 to 5. All higher difficulty levels were rarely reached (<1 trial per run on average) and were therefore excluded. Mean RT_{Exec} was 1537 ms ($SD = 211$ ms). Mean RT_{Choice} was 964 ms ($SD = 168$ ms). We tested whether RTs rose with increasing difficulty, using 2 independent linear regression analyses on RT_{Exec} and RT_{Choice} . Results showed a significant effect of difficulty on RT_{Exec} ($b = 0.69$, $t_{(78)} = 8.47$, $P < 0.001$), but no effect on RT_{Choice} ($b = 0.01$, $t_{(78)} = 0.11$, $P > 0.05$; Fig. 2A).

The overall error rate during the scanning session was 5.48% ($SD = 3.57\%$), suggesting that subjects successfully remembered and applied the rule-sets. As expected, results of a linear regression indicated that error rates increased with difficulty ($b = 0.66$, $t_{(78)} = 7.78$, $P < 0.001$; Fig. 2B). Error trials were excluded from the following analyses.

RT Switch Costs

We also tested for switch costs, that is, an increased RT_{Exec} subsequent to a task switch (Monsell 2003). Note that we could not test for task switch costs in the narrow sense of the word. We did not use different tasks, but rather different stimulus-response mappings within the same task, which might have influenced observable switch costs. We did not observe switch costs on RT_{Exec} ($P > 0.05$). There are 2 potential causes for

this result. First, switch costs are reduced if subjects can select the tasks themselves, compared with cued task switching (Arrington and Logan 2004, 2005). Secondly, long preparation periods greatly reduce switch costs in cued (Monsell 2003) and voluntary task switching (Arrington and Logan 2004, 2005). Given that both factors were present in our paradigm, it was unlikely to find switch costs in RTs. Note that even despite this fact, some evidence in favor of switch costs is reported in the task selection analysis (see below).

Task Selection

We assumed that task choices were based on at least 2 key pieces of information. The current task difficulty should influence the choice to switch away or stay in the current task-set. The “choice history” should influence to which of the 2 possible task-sets to switch to. Combined with the acquired knowledge of the difficulty dynamics during training, this would allow subjects to balance competing task demands: repeatedly performing the same task-set in order to exhibit a low difficulty level and to avoid switch costs, and flexibly switching between task-sets in order to find the easiest task-set at a given point in time.

First, we checked whether the decision to switch or stay was random or whether it was influenced by the current task difficulty. If subjects chose randomly between staying in the same task-set and switching to a different task-set, then the probability to stay should equal the probability to switch, which has been reported previously in studies on free choices (Arrington and Logan 2004; Forstmann et al. 2005, 2006; Haynes et al. 2007; Soon et al. 2008). Based on this notion, we calculated an expected probability distribution of run lengths, that is, the number of subsequent trials in the same task-set given that choices were made randomly (Arrington and Logan 2004). This distribution was then compared with the actual run length distribution using the Kolmogorov-Smirnov test to assess the randomness of choices (Fig. 2C). Results indicate that these distributions differ significantly ($P < 0.001$), confirming that subjects did not choose to switch or stay randomly in the present paradigm. Next, we directly tested whether switch/stay choices were influenced by task difficulty. A linear regression of the switch rate showed a significant positive relationship of switch/stay choices with the current difficulty ($b = 0.39$, $t_{(46)} = 2.80$, $P = 0.006$; Fig. 2D). These results strongly suggest that the decision to switch or stay was guided by the current task difficulty.

We then investigated whether the choice history influenced decision-making in this task. Owing to the dynamics of the environment, the longer a task-set has not been performed, the more likely it was at a low difficulty level. This information could have been used by the subjects to select 1 of the 2 available task-sets in switch trials. Thus, the number of trials that have passed since the chosen task-set has been last performed should have been larger than that of trials that have passed since the alternative non-chosen task-set has been last performed. The mean number of trials that have passed since the chosen task-set was last performed was 4.37 ($SD = 0.89$; Fig. 2E), which was significantly larger than the expected value if choices were random (2 trials, $t_{(15)} = 10.73$, $P < 0.001$), and also larger than the same value for the non-chosen task-set (2.99 trials, $SD = 0.66$, $t_{(15)} = 13.41$, $P < 0.001$). Additionally, we ran a linear regression of task choices in switch trials to look for a relationship with the number of trials that have passed

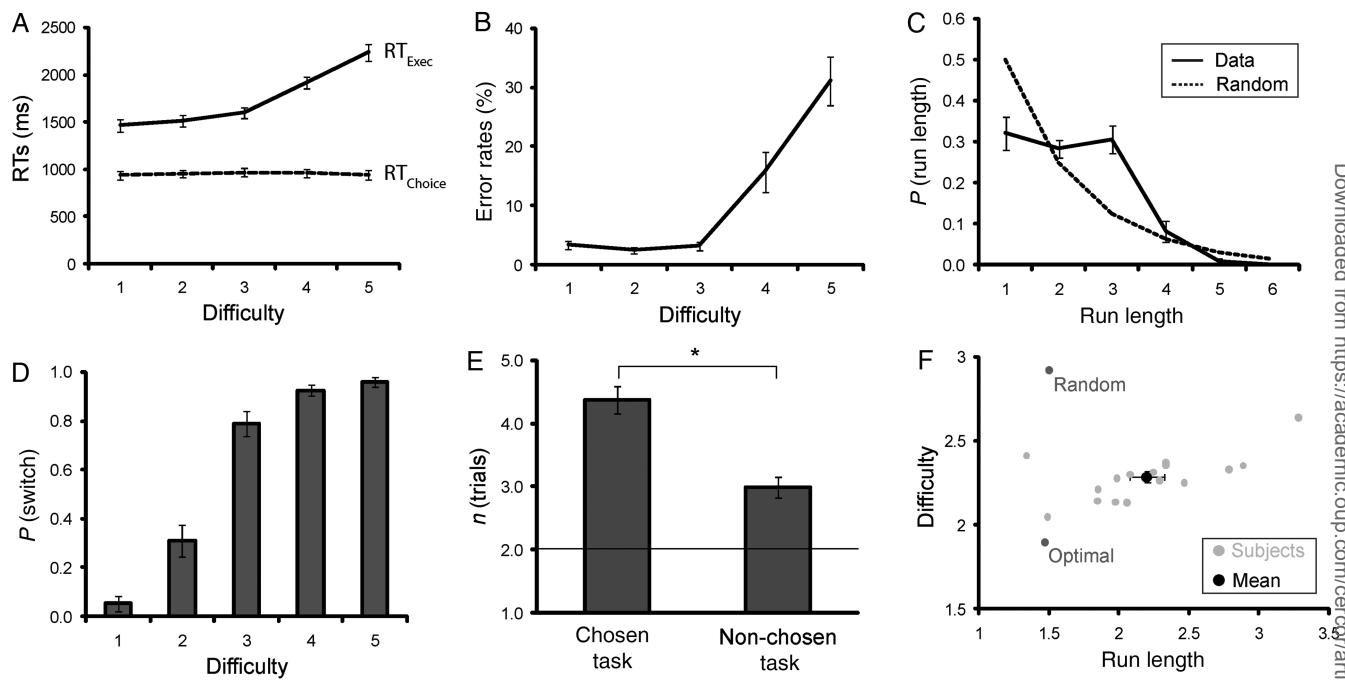


Figure 2. Behavioral results. (A) Reaction times (RTs). We measured 2 RTs per trial: The time to execute the task (RT_{Exec}) and the time to choose a task-set from the choice screen (RT_{Choice}). RT_{Exec} shows a strong increase for trials with a task difficulty above 3, while there was no effect of difficulty on RT_{Choice}. (B) Error rates. The graph depicts error rates as a function of difficulty. A strong increase in error rates is shown for trials with a task difficulty above 3. (C) Run length distribution. A run is a series of trials in the same task-set. Depicted is the distribution of run lengths from 1 to 6. The dashed line corresponds to the expected distribution if choices were made randomly. The solid line depicts the actual run length distribution, which clearly deviates from the random distribution. (D) Probability to switch away from a task-set, $P(\text{switch})$, as a function of difficulty. We observed a steep, monotonic increase in $P(\text{switch})$ with increasing difficulty. This indicates that switch/stay choices were influenced by the current task difficulty. (E) Choice history effect. Due to the dynamics of task difficulty in this paradigm, the longer a task-set has not been performed the more likely this task-set was at a low difficulty level. This information could have been used in switch trials, to decide to which of the 2 alternative task-sets to switch. The graph depicts the average number of trials that have passed since the chosen and non-chosen task-sets have been performed last. The asterisk indicates a significant difference, $P < 0.001$. The expected value for random choices is shown as a line for comparison. As can be seen, subjects indeed used their choice history to decide to which task-set to switch. (F) Simulations. We assessed subjects' performance by calculating the average difficulty across trials and the average run length for each subject. This approach took into account that it was effortful to perform difficult trials as well as switching regularly between task-sets. A successful subject should reach a low average difficulty while switching rarely between task-sets. Light gray dots represent individual subjects' data, and the black marker represents the mean values across all subjects. The "Random" marker shows the performance of a random task choice simulation for comparison. All subjects clearly performed better than the random choice simulation. The "Optimal" marker shows the performance of an optimal task choice simulation that picked the easiest task-set every trial. This simulation outperforms all subjects by showing the lowest average difficulty level, but it also switches between task-sets more often than most subjects. This possibly suggests that subjects chose to stay longer in one task-set, accepting more difficulty trials, in exchange for fewer task switches. All error bars represent standard errors.

since the chosen and non-chosen tasks were last performed. This analysis yielded results that were not significant ($b = 0.16$, $t_{(30)} = 1.97$, $P = 0.057$). Still, results seem to suggest that subjects used their choice history to determine which task-set to choose in switch trials and waited about 4 trials before choosing a task-set again. Taken together, this supports our initial notion that subjects tracked the current task difficulty and their choice history in order to select task-sets.

Subjects were free to decide how to keep difficulty at an overall low level. In order to assess their success in doing so, we calculated the average difficulty across all trials and compared this value to the average run length (Arrington and Logan 2004), which is a global measure for the number of task switches performed. This approach took into account that it was disadvantageous to perform difficult trials as well as to switch often between task-sets. Subjects were free to weight these 2 factors, for example, switching more often between task-sets in order to perform easier trials. Ranging from 1 to 9, the average difficulty was 2.29 (SD = 0.14). The average run length was 2.20 (SD = 0.50). There was considerable variability across subjects in the observed behavior (Fig. 2F). Performance was then compared with a "random choice simulation," in which a random task-set was chosen in each trial. The

resulting mean difficulty was 2.92 and the mean run length was 1.50. As expected, subjects' mean difficulty was significantly lower than what was realized for random choices ($t_{(15)} = -17.37$, $P < 0.001$), indicating that subjects did not choose task-sets randomly. Performance was also compared with an "optimal choice simulation," in which the easiest of the 3 task-sets was chosen in each trial. The resulting mean difficulty was 1.9 and the mean run length was 1.47. Subjects had a significantly higher mean difficulty ($t_{(15)} = 11.07$, $P < 0.001$), as well as a significantly higher run length ($t_{(15)} = 5.86$, $P < 0.001$). Note that the optimal choice simulation always chose the easiest task-set, regardless of switch costs. This suggests that subjects did not behave as expected if there were no costs associated with switching. Instead, subjects were willing to perform more difficult trials in order to avoid frequent switches between task-sets, in line with previous experiments reporting similar findings in voluntary task switching experiments (Yeung 2010). This suggests that subjects might have experienced switch costs. Alternatively, this behavior may be attributed to uncertainty about the status of the non-chosen task-sets: Subjects remained longer in the current task-set, to be certain that the alternative tasks-sets reached a low difficulty level before switching to them.

Multivariate Decoding of Task Choices

Distributed activation patterns in the right dACC (Brodmann area 32), extending to the dmPFC (Brodmann area 9), were found to predict the task choices (Fig. 3A, mean decoding accuracy 9.35% above chance, $t_{(13)} = 5.78$, $P < 0.001$). This accuracy level is similar to previous work decoding task-sets from local spatial activation patterns in the prefrontal cortex (Bode and Haynes 2009; Reverberi et al. 2012a). A complete list of decoding results is given in Table 1.

To provide further evidence for the validity of the analysis, we repeated the searchlight decoding, now randomly assigning the labels in the test dataset (Tusche et al. 2010). We repeated this analysis 1000 times and extracted a null distribution. We then extracted the decoding accuracy from the cluster in the dmPFC/dACC and compared this accuracy value against the null distribution of the same ROI. We found the accuracy to be significantly above chance ($P < 0.05$).

We also performed a conventional univariate analysis to test whether there was any difference in voxel activations between the 3 task choices, using an one-way ANOVA. No significant

results were found (voxel threshold $P < 0.001$ uncorrected, minimal cluster size: 20 voxels).

As an additional test of the independence of choice and difficulty signals, we repeated the searchlight decoding. We now additionally regressed out the effect of task difficulty, to account for possible effects of the task difficulty on the choice decoding in each individual subject (as suggested by Todd et al. 2013). More specifically, we regressed out the effects of the current task difficulty experienced while subjects made their task choice for the next trial. Given that we controlled for the influence of task difficulty directly in the neural data, no subjects were excluded based on their behavioral data in this analysis and the sample size was therefore larger compared with the original decoding analysis (2 subjects more). As can be seen in Figure 3B, regressing out difficulty from the signal does not strongly impact decoding results. In fact, results from both analyses are largely overlapping, lending further support for the independence of choice and difficulty signals found in our analysis. Regressing out the difficulty even seems to improve results slightly, which might be due to the difference in sample sizes. In order to explore the differences between these 2 analyses more formally, we performed an additional analysis. We calculated the difference map between our original analysis and the control analysis for each subject that was included in the original analysis. These difference maps were then assessed statistically at the group level using a simple t -test. In this analysis, the sample sizes were equal as the comparison was within-subjects. We found no voxel in which the 2 maps differed significantly (at the same threshold as the 2 analyses). This shows that there are no strong differences between results provided by our original analysis and the control analysis. Taken together, both analyses therefore provide converging evidence for the independence of task choice and difficulty signals in the dACC/dmPFC.

The task choice classification was above chance for a model based on the onset of the task screen preceding the choice. This could reflect an early choice for the next task at the onset of the task screen or immediately following completion of the trial. Note that fMRI cannot resolve such short timing differences accurately. We decided to perform post hoc tests and directly compare our model with 2 alternative models, in which brain responses were time-locked to (1) the task execution reaction time (RT_{exec} , on average 1537 ms after the task screen onset in case subjects made their choice after they finished performing the task) and (2) the choice screen onset (in case

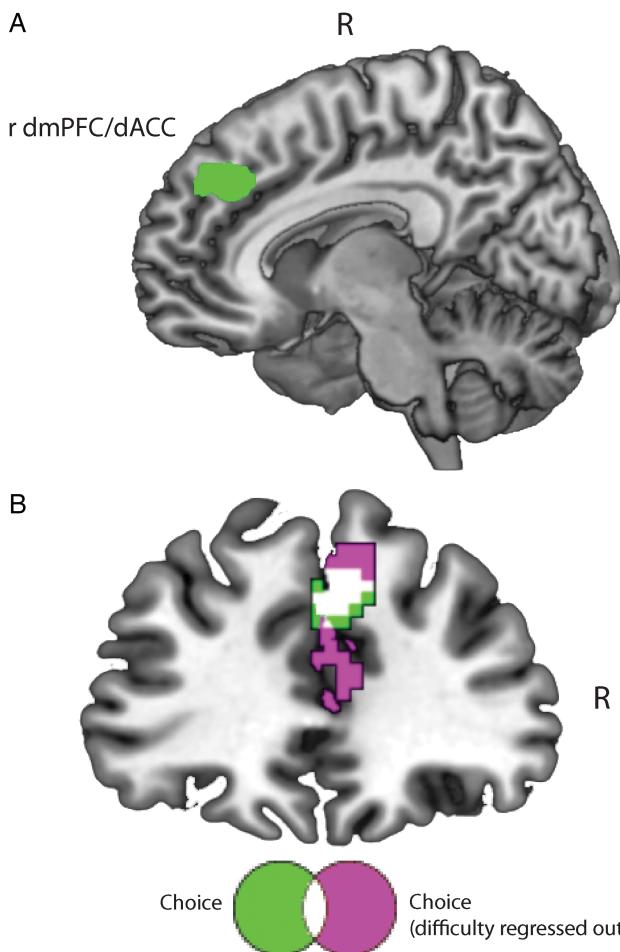


Figure 3. Results of the task choice decoding analysis. (A) Spatial activation patterns in the right dACC (Brodmann area 32) and dmPFC (Brodmann area 9) predicted task choices ($P < 0.001$, corrected at the cluster level, FWE $P < 0.05$). (B) Comparison and overlap of choice decoding (green) and an additional choice decoding control analysis in which task difficulty was regressed out of the neural data for each subject before running the decoding (violet). The control procedure was proposed by Todd et al. (2013) and leads to largely overlapping results.

Table 1

Brain regions encoding choices and difficulty

Brain region	Side	Cluster size	Accuracy above chance			MNI coordinates		
			M	SE	t-value	x	y	z
Choice decoding								
Dorso-medial PFC/ dorsal ACC	R	209	9.35	1.62	5.78	6	32	34
Difficulty decoding								
Ventro-lateral PFC	R	596	11.84	1.79	6.60	60	32	10
Medial lateral PFC	L	1615	13.4	1.64	8.15	-12	52	52
Occipital cortex	Bilateral	1606	12.23	1.5	8.16	-6	-93	10

Note: Results are shown for a statistical threshold of $P < 0.001$, FWE corrected at the cluster level ($P < 0.05$).

The mean accuracies and SEs with the corresponding t-value are displayed for each cluster. The coordinates of the peak voxel in each cluster are also shown.

subjects made their choice when the task choice options were presented). We extracted the ROI coordinates from the dmPFC/dACC cluster identified in the original analysis and tested whether accuracies were significantly above chance in the same ROI in the model locked to the task execution reaction time. Decoding accuracies were significantly above chance (mean decoding accuracy 6.03% above chance, $t_{(13)} = 2.52$, $P < 0.05$). However, decoding accuracies were still significantly below the results from our original analysis (mean decoding accuracy 9.35% above chance, $t_{(13)} = 1.85$, $P < 0.05$). For the second model time-locked to the choice screen onset, the ROI analysis revealed no significant above-chance accuracies in the dmPFC/dACC. Thus, possibly, the format of choice-selective signals is either not accessible to our decoding analysis or is represented elsewhere in the brain at this stage. Nevertheless, these additional checks suggest that our original model provides the best explanation of the data, lending further support to our choice of analysis parameters.

As shown in Figure 2F, there was considerable variance between subjects' behavior. We reasoned that subjects might differ in the strategies they used in our task. Given the long training, some subjects might have developed and followed a clear strategy, leading to little variability in their choice behavior. Others might have a less clear strategy, leading to more variability in behavior. For that reason, we assessed subjects' between run variance in mean difficulty by calculating the standard deviation of run-by-run mean difficulties for each subject. We found the variance in performance to be correlated with subjects' decoding accuracies in the right dmPFC/dACC ($r_{(12)} = -0.55$, $P = 0.04$). Interestingly, there was no significant correlation of decoding accuracies with the actual mean difficulty ($r_{(12)} = -0.27$, $P > 0.05$). This might suggest that subjects following a clear strategy, even if it was not highly successful, have a higher choice decoding accuracy in the dmPFC/dACC.

These findings point towards the involvement of the dmPFC/dACC in representing the choice of the future task-set. Alternatively, this finding is also compatible with the possibility that dmPFC/dACC is only representing all variables affecting choices, that is, the estimated difficulty level of all 3 task-sets and the choice history. However, given that we used a searchlight approach, all this information would have to be available in each and every searchlight to produce successful predictions. In either case, the finding shows a key role of the dmPFC/dACC in task choice.

Multivariate Decoding of Task Difficulty

In a next step, we investigated whether activation patterns also contained information about one major determinant of these task choices, that is, the current level of difficulty. We found information about the current difficulty level in the right ventro-lateral PFC (mean decoding accuracy 11.84% above chance, $t_{(13)} = 6.60$, $P < 0.001$, Fig. 4), a large cluster in the left PFC also encompassing the left anterior insula and spanning into the right dmPFC (mean decoding accuracy 13.40% above chance, $t_{(13)} = 8.15$, $P < 0.001$), and the bilateral occipital cortex (mean decoding accuracy 12.23% above chance, $t_{(13)} = 8.16$, $P < 0.001$).

As a further control analysis, we repeated the multivariate decoding of difficulty, randomly assigning the test labels using the same procedure as in the choice decoding. We found all clusters identified in the difficulty decoding

analysis to be significantly above chance ($P < 0.05$). We also repeated the difficulty decoding, only regressing out choice signals beforehand, similar to the choice decoding analysis. The additional regressor does not seem to strongly impact results, as there is considerable overlap with the original difficulty decoding analysis (see Supplementary Fig. 2 for more information).

Contrary to the choice analysis, an additional univariate analysis of task difficulty (contrast: difficulty level 2 > difficulty level 1) revealed increased activation for higher levels of difficulty in the posterior dmPFC, right precentral gyrus, posterior cingulate cortex, right vIPFC, and bilateral occipital cortex (voxel threshold $P < 0.001$, corrected at a cluster level, FWE $P < 0.05$). We also found a decrease of activation with increasing difficulty in the anterior and ventro-lateral PFC, right dorsal lateral PFC, right angular gyrus, bilateral posterior cingulate gyrus, and bilateral mid-temporal gyrus (voxel threshold $P < 0.001$, corrected at a cluster level, FWE $P < 0.05$). Therefore, results obtained in the decoding analysis can partly be explained with the expected univariate signal increase or decrease in the identified areas.

Convergence of Information

We were able to investigate the neural representations of task choices and difficulty independently, which allowed us to address the issue of whether choice and difficulty information converges in some brain areas. If a brain area integrates information about the environmental factors influencing choices with the choices themselves, such convergence is one necessary condition that brain area needs to meet. So far, we have shown that the right dmPFC/dACC encoded choices. We assessed whether this region also contained information about task difficulty. We repeated the task difficulty analysis, this time applying small volume correction so that only the right dmPFC/dACC was considered. Results showed significant encoding of difficulty within this cluster ($P < 0.005$ uncorrected, corrected for multiple comparisons at the cluster level, FWE $P < 0.05$, Fig. 5). In addition, we assessed whether the clusters identified in the difficulty decoding analysis (left PFC, right vIPFC, and occipital cortex) contained information about task choices, using the same approach. Here, we did not find any significant results ($P < 0.005$ uncorrected, corrected for multiple comparisons at the cluster level, FWE $P < 0.05$). Therefore, only the right dmPFC/dACC encoded both task choices and the environmental factor influencing these choices, that is, task difficulty.

To further validate these results, we created a brain mask only including voxels that showed an overlap of choice and difficulty decoding (Fig. 5). We then tested whether decoding accuracies within these voxels were significantly above chance. This procedure was proposed by Etzel et al. (2013), as one cannot infer from searchlight decoding results that a larger cluster identified in such an analysis is also encoding information. Results showed that voxels in which choice and difficulty decoding results overlapped showed accuracies that were significantly above chance for choice ($t_{(13)} = 5.18$, $P < 0.001$) and difficulty ($t_{(13)} = 4.10$, $P = 0.001$).

We also directly tested whether the overlap of information was limited to the voxels where an overlap was observed or whether the whole right dmPFC/dACC exhibited this pattern of results. If the convergence of information was limited to the

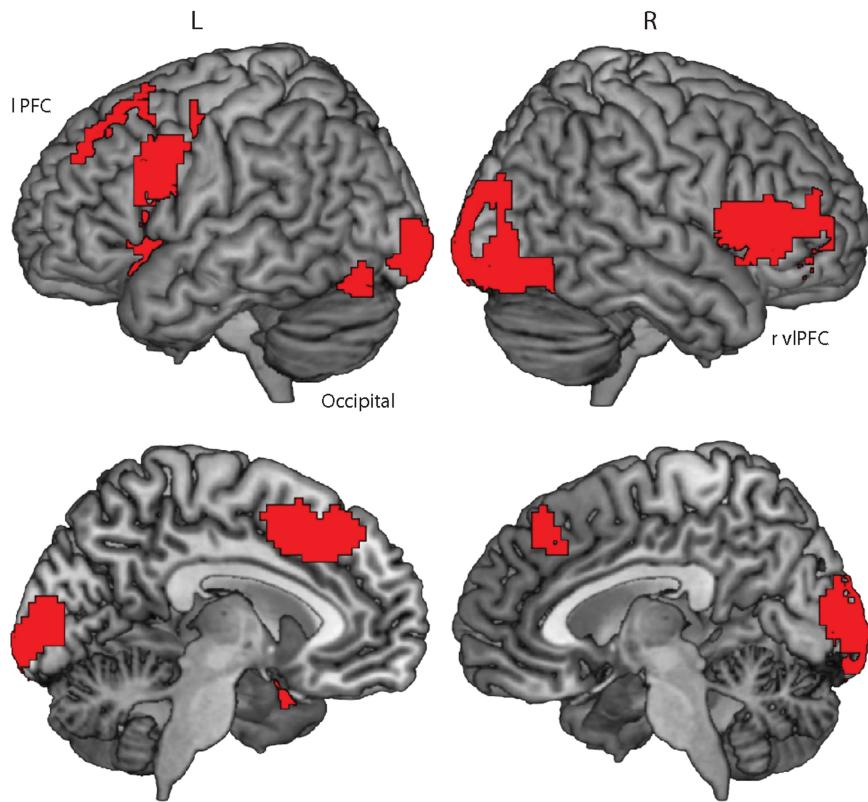


Figure 4. Results of the difficulty decoding analysis. Spatial activation patterns in a large cluster in the left prefrontal cortex (lPFC) encompassing lateral and medial areas, lateral occipital cortex, and right ventro-lateral prefrontal cortex (r vIPFC) encoded the current difficulty level of the trial ($P < 0.001$, corrected at the cluster level, FWE $P < 0.05$).

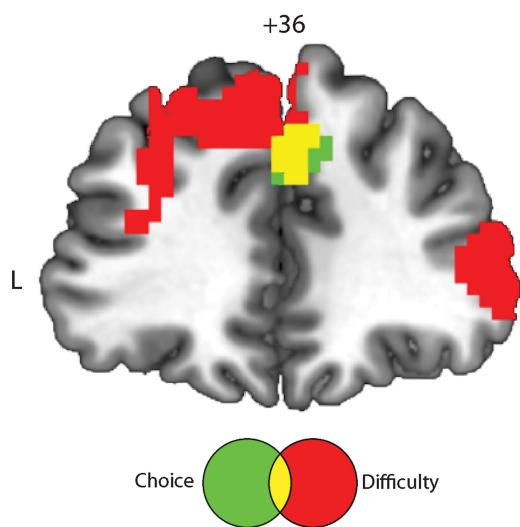


Figure 5. Convergence of information. Depicted are brain regions that were informative about task choices (green) and the current level of difficulty (red). Brain regions that encoded both variables were identified using small volume correction (yellow).

overlap voxels, we should not find both choice and difficulty information in a dmPFC/dACC ROI after removing the overlap voxels from it (Etzel et al. 2013). However, even after removal of the overlap voxels, we still found choice information ($t_{(13)} = 5.66$, $P < 0.001$) and difficulty information ($t_{(13)} = 3.25$, $P = 0.006$) in the right dmPFC/dACC. We can infer that the

convergence of choice and difficulty information can be seen in the whole right dmPFC/dACC cluster.

Discussion

The present study investigated the neural correlates of task selection in dynamically changing environments, using a novel difficulty-based choice task. Subjects were required to balance the opposing demands of staying in the same task-set versus switching flexibly between task-sets. On the one hand, repeated task-set performance allowed subjects to avoid switch costs. On the other hand, task difficulty changed dynamically over time: in order to find the easiest of the 3 available task-sets, subjects had to switch among them. We assessed separately which brain areas encoded task choices resulting from balancing these opposing demands and which brain areas encoded task difficulty, a major factor influencing task choices. Results showed that local spatial activation patterns in the right dmPFC/dACC independently encoded both task choices and the current level of difficulty, showing that this brain region contains information about choices ("How can I reach my goal?") and the factor influencing them ("Why do I make this particular choice?").

Decision-Making in Predictable Dynamic Environments

In most everyday situations, we are faced with environments that change dynamically over time. These changes are often predictable to some degree, partly due to the fact that our actions lead to predictable effects in these environments. Yet,

in previous work on voluntary task selection, subjects were either asked to make choices in environments that did not change at all (Haynes et al. 2007; Soon et al. 2008) or that changed unpredictably for the subjects (Hampton and O'Doherty 2007; Boorman et al. 2009). In contrast, our difficulty-based choice task was able to capture decision-making processes in predictably changing environments, as subjects were aware that a specific action will lead to a specific effect in the environment. We assessed which factors influenced such choices and behavioral data demonstrated that subjects considered both the current state of the environment (task difficulty) and their own choice history in order to select task-sets. These results are similar to previous studies on foraging (Cohen et al. 2007; Hayden, Heilbronner, et al. 2011; Kolling et al. 2012), in which subjects are also faced with a predictable dynamic environment. It has been shown that subjects adapt to trial-by-trial changes in foraging environments (Sugrue et al. 2004; Daw et al. 2006), and that they are using their choice history to do so (Kennerley et al. 2006).

Choice Encoding in the dmPFC/dACC

We demonstrated that task choices were encoded in the right dmPFC/dACC. Earlier studies demonstrated that similar areas encode unconstrained choices (Haynes et al. 2007) and reward-based choices (Hampton and O'Doherty 2007). Furthermore, it has been shown that the dmPFC compares the values of different choice options (Wunderlich et al. 2009; Hare et al. 2011). The dACC also plays an important role in foraging, as neurons in this region were found to encode foraging choices (Hayden, Heilbronner, et al. 2011) and lesions to this area lead to a drop in performance during foraging (Kennerley et al. 2006). The dmPFC and dACC are also crucial brain regions for several important subfunctions that are necessary (but not sufficient) for task selection in dynamic environments, such as action outcome prediction (Alexander and Brown 2011), the processing of uncertainty (Volz et al. 2003; Rushworth and Behrens 2008) and task difficulty (Barch et al. 1997), cost-benefit evaluation (Walton et al. 2003; Kennerley et al. 2009), conflict monitoring (Botvinick et al. 2001), and error processing (Desmet et al. 2011). We were able to show that the dmPFC/dACC independently encodes task choices and one of the major factors influencing these choices, task difficulty. It is important to note that these 2 factors are inherently correlated in most reward- and difficulty-based choice tasks. For instance, rewards and choices were confounded in a previous experiment (Hampton and O'Doherty 2007), in which reward-based choices were predicted from activity patterns in the dACC. However, it remained unclear whether choice prediction from the dACC was due to choice- or reward-related signals in that brain area. We provide evidence that this brain region, in fact, independently encodes both choices and variables affecting choices, in our case the task difficulty. Thus, one might speculate whether the dmPFC/dACC integrates the "How" and "Why" aspects during decision-making. Although we cannot test for integration directly, we can show that this region at least fulfills one necessary condition for this function: access to information about both choice and difficulty. Future experiments will need to clarify the precise computational function of this brain area during decision-making in dynamic environments. In our behavioral analysis, we also demonstrated that subjects considered the choice history to select

task-sets. The present dataset did not allow us to investigate the neural basis of this behavioral effect. Moreover, it is an intriguing question for future research whether the dmPFC/dACC also contains information about the choice history in humans, a notion supported by previous data from non-human primates (Kennerley et al. 2006). A further open question about decision-making in dynamic environments is the temporal evolution of choices in decision-making networks (e.g. Hunt et al. 2012), from value representations of choice options to the actual resulting decision. Unfortunately, fMRI lacks the temporal resolution to address this issue directly in the present study, but it might be possible to investigate temporal dynamics of decision-making in dynamic environments in future experiments using Magnetoencephalography or Electroencephalography.

Theories of dmPFC/dACC Function

Recent theories of dACC function suggest that this brain area connects motivational valence with executive control processes (Williams et al. 2004; Kennerley et al. 2006; Pessoa 2009). With its afferent connections from subcortical reward-related networks and efferent connections to prefrontal control and action selection networks, the dACC acts as a hub between reward-related/motivational networks and executive networks (Pessoa and Engelmann 2010; Rushworth et al. 2011). One prediction derived from these theories is that the dACC should have access to information about both motivational and executive variables (in our design that translates to access to both difficulty and choice information), which we confirmed using an MVPA. A more recent theory is provided by Holroyd and Yeung (2012), who claim that the ACC's main role is to maintain "options"—sequences of simple actions directed toward larger goals. These options provide the motivation for moment-to-moment actions. Again, this locates the ACC in between motivational and executive networks in the brain. In our experiment, subjects tried to correctly perform each trial, with the long-term goal of keeping the overall difficulty low so that future trials would stay easy. One might speculate that a way to reach this goal is to organize behavior into options, which would describe a sequence of task choices that keep difficulty low in the long run. Converging evidence for such an interpretation of our results comes from voluntary task switching experiments, where subjects seem to choose sequences of actions, not individual actions in each trial (Vandierendonck et al. 2012). According to Holroyd and Yeung, the specific function of the ACC then is to select which of the available choice sequences to implement. It might be that our task, which strongly emphasizes the future consequences of task choices, relies critically on the ACC precisely for this reason. In light of this theory, we might find task decoding results in the ACC, because in fact subjects organized their behavior in choice sequences. This would allow subjects to make individual task choices with relatively little effort, as they are already specified in the chosen sequence, and would also explain why we find task choice information very early in the trial while subjects were busy with the task execution. Alternatively, one might also regard task-sets, instead of choice sequences, as being the options subjects choose from. In fact, task-sets can be seen as a simpler form of behavioral options (Botvinick et al. 2009), as both abstract from simple stimulus-response mappings. This might suggest that both task-sets and choice sequences share some common neural basis. Although speculative, this

explanation of our results offers an intriguing perspective on the specific function of the ACC in dynamic environments.

It further seems to be the case that the dmPFC and dACC are recruited especially if subjects make voluntary choices (Walton et al. 2004; Forstmann et al. 2006; Haynes et al. 2007; but see Zhang et al. 2013), whereas the lateral PFC seems to be recruited more if subjects are cued on which task to perform (Bode and Haynes 2009; Reverberi et al. 2012a,b). In this context, our task is an interesting case as subjects did choose freely, while at the same time the environment often suggested which choice is best—similar to a cued task paradigm. Our choice decoding results suggest that even if the environment guides voluntary task choices, if subjects do not receive an “imperative” cue, choices will be represented in dorsal medial prefrontal areas associated with voluntary choice.

Lastly, it has been suggested that the dmPFC is organized along a posterior-to-anterior gradient (Venkatraman et al. 2009), similar to the lateral PFC (Koechlin and Summerfield 2007; Badre 2008). While posterior portions of the dmPFC were found to be activated when subjects resolved low-level response conflicts, more anterior portions were activated when subjects engaged in high-level control functions, which are required during task selection. Interestingly, the cluster identified in the choice decoding analysis is located in an anterior portion of the dmPFC, as would be predicted by the gradient hypothesis. Notably, a previous study using unconstrained voluntary choices (Haynes et al. 2007) reported results that were even more anterior, when compared with our findings. This might suggest that although both choosing tasks voluntarily with and without external guidance rely on the dmPFC, the specific portion of the dmPFC that is recruited might differ. Along the lines of Venkatraman et al. (2009), one possible interpretation is that voluntary task choices without any external guidance require even more high-level control functions than externally guided voluntary choices, which would lead to more anterior dmPFC involvement.

Conclusion

In this experiment, we investigated voluntary task selection processes, using a novel difficulty-based choice task in which subjects needed to balance different opposing task demands in a dynamically changing environment. We demonstrated that task selection was influenced both by the current state of the environment and by the subjects' choice history. Using MVPA methods, we demonstrated that the right dmPFC/dACC independently encoded task choices as well as a major factor influencing these choices, task difficulty. These findings emphasize the importance of the dmPFC/dACC for task selection and motivation, which is in line with recent theories of both dmPFC and dACC function. The results also broaden our understanding of task selection processes in naturalistic dynamic environments.

Supplementary Material

Supplementary material can be found at: <http://www.cercor.oxfordjournals.org/>.

Funding

This work was supported by the Bernstein Computational Neuroscience Program of the German Federal Ministry of

Education and Research (grant reference 01GQ1001C), the German Research Foundation (Exc 257 NeuroCure, DFG Grant KFO247, and DFG Grant SFB 940), and the Berlin School of Mind and Brain.

Notes

Conflict of Interest: None declared.

References

Alexander WH, Brown JW. 2011. Medial prefrontal cortex as an action-outcome predictor. *Nat Neurosci*. 14:1338–1344.

Arrington CM, Logan GD. 2004. The cost of a voluntary task switch. *Psychol Sci*. 15:610–615.

Arrington CM, Logan GD. 2005. Voluntary task switching: chasing the elusive homunculus. *J Exp Psychol Lear Mem Cogn*. 31:683–702.

Badre D. 2008. Cognitive control, hierarchy and the rostro-caudal organization of the frontal lobes. *Trends Cogn Sci*. 12:193–200.

Barch DM, Braver TS, Nystrom LE, Forman SD, Noll DC, Cohen JD. 1997. Dissociating working memory from task difficulty in human prefrontal cortex. *Neuropsychologia*. 35:1373–1380.

Bode S, Haynes JD. 2009. Decoding sequential stages of task preparation in the human brain. *Neuroimage*. 45:606–613.

Boorman ED, Behrens TEJ, Woolrich MW, Rushworth MFS. 2009. How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron*. 62:733–743.

Boorman ED, Rushworth MFS, Behrens TEJ. 2013. Ventromedial prefrontal and anterior cingulate cortex adopt choice and default reference frames during sequential multi-alternative choice. *J Neurosci*. 33:2242–2253.

Botvinick MM, Braver TS, Barch DM, Carter CS, Cohen JD. 2001. Conflict monitoring and cognitive control. *Psychol Rev*. 108:624–652.

Botvinick MM, Niv Y, Barto AC. 2009. Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. *Cognition*. 113:262–280.

Cohen JD, McClure SM, Yu AJ. 2007. Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philos T Roy Soc B*. 362:933–942.

Cox DD, Savoy RL. 2003. Functional magnetic resonance imaging (fMRI) “brain reading”: detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage*. 19:261–270.

Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ. 2006. Cortical substrates for exploratory decisions in humans. *Nature*. 441:876–879.

Desmet C, Fias W, Hartstra E, Brass M. 2011. Errors and conflict at the task level and the response level. *J Neurosci*. 31:1366–1374.

Duncan J. 2001. An adaptive coding model of neural function in prefrontal cortex. *Nat Rev Neurosci*. 2:820–829.

Etzel JA, Zacks JM, Braver TS. 2013. Searchlight analysis: promise, pitfalls, and potential. *Neuroimage*. 78:261–269.

Forstmann BU, Brass M, Koch I, von Cramon DY. 2005. Internally generated and directly cued task sets: an investigation with fMRI. *Neuropsychologia*. 43:943–952.

Forstmann BU, Brass M, Koch I, von Cramon DY. 2006. Voluntary selection of task sets revealed by functional magnetic resonance imaging. *J Cogn Neurosci*. 18:388–398.

Friston KJ, Holmes AP, Worsley KJ, Poline JP, Frith CD, Frackowiak RSJ. 1994. Statistical parametric maps in functional imaging: a general linear approach. *Hum Brain Mapp*. 2:189–210.

Hampton AN, O'Doherty JP. 2007. Decoding the neural substrates of reward-related decision making with functional MRI. *Proc Natl Acad Sci USA*. 104:1377–1382.

Hare TA, Schultz W, Camerer CF, O'Doherty JP, Rangel A. 2011. Transformation of stimulus value signals into motor commands during simple choice. *Proc Natl Acad Sci USA*. 108:18120–18125.

Hayden BY, Heilbronner SR, Pearson JM, Platt ML. 2011. Surprise signals in anterior cingulate cortex: neuronal encoding of unsigned

reward prediction errors driving adjustment in behavior. *J Neurosci.* 31:4178–4187.

Hayden BY, Pearson JM, Platt ML. 2011. Neuronal basis of sequential foraging decisions in a patchy environment. *Nat Neurosci.* 14:933–939.

Haynes JD, Rees G. 2006. Decoding mental states from brain activity in humans. *Nat Rev Neurosci.* 7:523–534.

Haynes JD, Sakai K, Rees G, Gilbert S, Frith C, Passingham RE. 2007. Reading hidden intentions in the human brain. *Curr Biol.* 17:323–328.

Holroyd CB, Yeung N. 2012. Motivation of extended behaviors by anterior cingulate cortex. *Trends Cogn Sci.* 16:122–128.

Hunt LT, Kolling N, Soltani A, Woolrich MW, Rushworth MFS, Behrens TEJ. 2012. Mechanisms underlying cortical activity during value-guided choice. *Nat Neurosci.* 15:470–479.

Kamitani Y, Tong F. 2005. Decoding the visual and subjective contents of the human brain. *Nat Neurosci.* 8:679–685.

Kennerley SW, Dalmubed AF, Lara AH, Wallis JD. 2009. Neurons in the frontal lobe encode the value of multiple decision variables. *J Cogn Neurosci.* 21:1162–1178.

Kennerley SW, Walton ME, Behrens TEJ, Buckley MJ, Rushworth MFS. 2006. Optimal decision making and the anterior cingulate cortex. *Nat Neurosci.* 9:940–947.

Koechlin E, Summerfield C. 2007. An information theoretic approach to prefrontal executive function. *Trends Cogn Sci.* 11:229–235.

Kolling N, Behrens TEJ, Mars RB, Rushworth MFS. 2012. Neural mechanisms of foraging. *Science.* 336:95–98.

Kriegeskorte N, Goebel R, Bandettini P. 2006. Information-based functional brain mapping. *Proc Natl Acad Sci USA.* 103:3863–3868.

Miller EK, Cohen JD. 2001. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci.* 24:167–202.

Mitchell TM, Hutchinson R, Niculescu RS, Pereira F, Wang X. 2004. Learning to decode cognitive states from brain images. *Mach Learn.* 57:145–175.

Monsell S. 2003. Task switching. *Trends Cogn Sci.* 7:134–140.

Müller J, Dreisbach G, Goschke T, Hensch T, Lesch K-P, Brodbeck B. 2007. Dopamine and cognitive control: the prospect of monetary gains influences the balance between flexibility and stability in a set-shifting paradigm. *Eur J Neurosci.* 26:3661–3668.

Pessoa L. 2009. How do emotion and motivation direct executive control? *Trends Cogn Sci.* 13:160–166.

Pessoa L, Engelmann JB. 2010. Embedding reward signals into perception and cognition. *Front Neurosci.* 4:1–8.

Reverberi C, Görgen K, Haynes JD. 2012a. Compositionality of rule representations in human prefrontal cortex. *Cereb Cortex.* 22:1237–1246.

Reverberi C, Görgen K, Haynes JD. 2012b. Distributed representations of rule identity and rule order in human frontal cortex and striatum. *J Neurosci.* 32:17420–17430.

Rushworth MFS, Behrens TEJ. 2008. Choice, uncertainty and value in prefrontal and cingulate cortex. *Nat Neurosci.* 11:389–397.

Rushworth MFS, Noonan MP, Boorman ED, Walton ME, Behrens TE. 2011. Frontal cortex and reward-guided learning and decision-making. *Neuron.* 70:1054–1069.

Shallice T, Burgess PW. 1991. Deficits in strategy application following frontal lobe damage in man. *Brain.* 114:727–741.

Shenhav A, Botvinick MM, Cohen JD. 2013. The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron.* 79:217–240.

Soon CS, Brass M, Heinze HJ, Haynes JD. 2008. Unconscious determinants of free decisions in the human brain. *Nat Neurosci.* 11:543–545.

Sugrue LP, Corrado GS, Newsome WT. 2004. Matching behavior and the representation of value in the parietal cortex. *Science.* 304:1782–1787.

Todd MT, Nystrom LE, Cohen JD. 2013. Confounds in multivariate pattern analysis: theory and rule representation case study. *Neuroimage.* 77:157–165.

Tusche A, Bode S, Haynes JD. 2010. Neural responses to unattended products predict later consumer choices. *J Neurosci.* 30:8024–8033.

Vandierendonck A, Demanet J, Liefhooghe B, Verbruggen F. 2011. A chain-retrieval model for voluntary task switching. *Cogn Psychol.* 65:214–283.

Venkatraman V, Rosati AG, Taren AA, Huettel SA. 2009. Resolving response, decision, and strategic control: evidence for a functional topography in dorsomedial prefrontal cortex. *J Neurosci.* 29:13158–13164.

Volz KG, Schubotz RI, von Cramon DY. 2003. Predicting events of varying probability: uncertainty investigated by fMRI. *Neuroimage.* 19:271–280.

Walton ME, Bannerman DM, Alterescu K, Rushworth MFS. 2008. Functional specialization within medial frontal cortex of the anterior cingulate for evaluating effort-related decisions. *J Neurosci.* 23:6475–6479.

Walton ME, Devlin JT, Rushworth MFS. 2004. Interaction between decision making and performance monitoring within prefrontal cortex. *Nat Neurosci.* 7:1259–1265.

Williams ZM, Bush G, Rauch SL, Cosgrove GR, Eskandar EN. 2004. Human anterior cingulate neurons and the integration of monetary reward with motor responses. *Nat Neurosci.* 7:1370–1375.

Worsley KJ, Marrett S, Neelin P, Vandal AC, Friston KJ, Evans AC. 1996. A unified statistical approach for determining significant signals in images of cerebral activation. *Hum Brain Mapp.* 4:58–73.

Wunderlich K, Rangel A, O'Doherty JP. 2009. Neural computations underlying action-based decision making in the human brain. *Proc Natl Acad Sci USA.* 106:17199–17204.

Yeung N. 2010. Bottom-up influences on voluntary task-switching: the elusive homunculus escapes. *J Exp Psychol Learn.* 36:348–362.

Zhang J, Kriegeskorte N, Carlin JD, Rowe JB. 2013. Choosing the rules: distinct and overlapping frontoparietal representations of task rules for perceptual decisions. *J Neurosci.* 33:11852–11862.