

The Role of the Parietal Cortex in the Representation of Task–Reward Associations

David Wisniewski,^{1,2,5} Carlo Reverberi,^{6,7} Ida Momennejad,⁸ Thorsten Kahnt,⁹ and John-Dylan Haynes^{1,2,3,4,5,10,11}

¹Bernstein Center for Computational Neuroscience Berlin, ²Berlin Center for Advanced Neuroimaging, ³Cluster of Excellence NeuroCure, and ⁴Department of Neurology, Charité University Medicine Berlin, D-10117 Berlin, Germany, ⁵Collaborative Research Centre SFB 940, Volition and Cognitive Control, Dresden University of Technology, D-01069 Dresden, Germany, ⁶Psychology Department, University of Milano–Bicocca, I-20126 Milan, Italy, ⁷Milan Center for Neuroscience, I-20052 Milan, Italy, ⁸Princeton Neuroscience Institute, Princeton University, Princeton, New Jersey 08544, ⁹Northwestern University Feinberg School of Medicine, Department of Neurology, Chicago, Illinois 60611, and ¹⁰Berlin School of Mind and Brain and ¹¹Department of Psychology, Humboldt University of Berlin, D-10099 Berlin, Germany

Rewards obtained from specific behaviors can and do change across time. To adapt to such conditions, humans need to represent and update associations between behaviors and their outcomes. Much previous work focused on how rewards affect the processing of specific tasks. However, abstract associations between multiple potential behaviors and multiple rewards are an important basis for adaptation as well. In this experiment, we directly investigated which brain areas represent associations between multiple tasks and rewards, using time-resolved multivariate pattern analysis of functional magnetic resonance imaging data. Importantly, we were able to dissociate neural signals reflecting task–reward associations from those related to task preparation and reward expectation processes, variables that were often correlated in previous research. We hypothesized that brain regions involved in processing tasks and/or rewards will be involved in processing associations between them. Candidate areas included the dorsal anterior cingulate cortex, which is involved in associating simple actions and rewards, and the parietal cortex, which has been shown to represent task rules and action values. Our results indicate that local spatial activation patterns in the inferior parietal cortex indeed represent task–reward associations. Interestingly, the parietal cortex flexibly changes its content of representation within trials. It first represents task–reward associations, later switching to process tasks and rewards directly. These findings highlight the importance of the inferior parietal cortex in associating behaviors with their outcomes and further show that it can flexibly reconfigure its function within single trials.

Key words: cognitive control; fMRI; multivariate pattern analysis; parietal cortex; reward; task-set

Significance Statement

Rewards obtained from specific behaviors rarely remain constant over time. To adapt to changing conditions, humans need to continuously update and represent the current association between behavior and its outcomes. However, little is known about the neural representation of behavior–outcome associations. Here, we used multivariate pattern analysis of functional magnetic resonance imaging data to investigate the neural correlates of such associations. Our results demonstrate that the parietal cortex plays a central role in representing associations between multiple behaviors and their outcomes. They further highlight the flexibility of the parietal cortex, because we find it to adapt its function to changing task demands within trials on relatively short timescales.

Introduction

The payoff obtained from a specific behavior often changes over time, and our cognitive system needs to adapt to such changing

environmental conditions. We often have several options to choose from, with different potential outcomes. Associating complex activities and rewards, even if this association remains counter-factual, is a key function for decision-making, motivational, and cognitive control (Ridderinkhof et al., 2004), and there is a large literature on its neural basis in both humans (O’Doherty et al., 2003; Alexander and Brown, 2011; Kovach et

Received Nov. 28, 2014; revised June 29, 2015; accepted July 16, 2015.

Author contributions: D.W., C.R., I.M., and J.-D.H. designed research; D.W. performed research; D.W., C.R., T.K., and J.-D.H. analyzed data; D.W., C.R., I.M., T.K., and J.-D.H. wrote the paper.

This work was supported by the Bernstein Computational Neuroscience Program of the German Federal Ministry of Education and Research (Grant 01GQ1001C), German Research Foundation Collaborative Research Centre SFB 940 “Volition and Cognitive Control: Mechanisms, Modulations, Dysfunctions,” German Research Foundation Grants Exc

257 NeuroCure and KFO247, and Italian Ministry of University Research Programs of Relevant National Interest Grant 2010RP5RNM_001 (C.R.). We thank Carsten Allefeld for his valuable comments on this manuscript.

al., 2012) and nonhuman primates (Hadland et al., 2003; Platt and Glimcher, 1999; Sugrue et al., 2004; Chudasama et al., 2013).

Previous work focused on how reward expectations affect task processing and execution in the brain (Knutson et al., 2000; Daw et al., 2006; Hampton and O'Doherty, 2007; Etzel et al., 2015), not directly investigating task–reward associations. Such associations are context-dependent rules linking multiple potential behaviors and rewards (“If I complete task A, I will receive reward R_A . But if I complete task B, I will receive reward R_B .”). Knowing the outcomes of different behaviors is a critical basis for successful adaptation to changing environments, yet it is difficult to distinguish knowledge about these associations from preparatory processes. If the reward outcome for a specific task is known, we associate the task and the reward, but we also prepare for task execution and expect to gain the reward.

Furthermore, past experiments focused mainly on associations of low-level actions on rewards (Ridderinkhof et al., 2004; Rushworth et al., 2004; Hayden et al., 2011), yet more complex behaviors, such as abstract task sets, can also be associated with reward values. Compared with low-level stimulus–response mappings, task sets are more generalizable and abstract (Sakai, 2008; Holroyd and Yeung, 2012), arguably making them more robust against changes in the environment to which we adapt.

Previous work using multivariate pattern analysis (MVPA) has shown that reward values are represented in striatal and orbitofrontal brain regions (Kahnt et al., 2010), whereas task sets are represented in a frontoparietal network (Sigala et al., 2008; Woolgar et al., 2011; Momennejad and Haynes, 2012; Reverberi et al., 2012a; Wisniewski et al., 2014). The parietal cortex has further been shown to represent relative action values (Sugrue et al., 2004; Kahnt et al., 2014). Based on previous findings, we hypothesized that task–reward associations might either be represented in these task and/or reward-related brain regions.

In this experiment, we modified a delayed intention paradigm (Momennejad and Haynes, 2013) to investigate task–reward associations. Subjects underwent functional magnetic resonance imaging (fMRI) while performing the task. We used time-resolved MVPA (Soon et al., 2008) to determine which brain areas contain information about task–reward associations and to assess the evolution of information across time in these brain areas. Importantly, we were able to independently assess the neural signals of task–reward associations from those related to task preparation and reward expectation processes.

Materials and Methods

Participants

Nineteen participants took part in the experiment (10 females). All subjects volunteered to participate and had normal or corrected-to-normal vision. Subjects gave written informed consent and received between 35€ and 65€ for participation, depending on their performance. The experiment was approved by the local ethics committee. No subject had a history of neurological or psychiatric disorders. Data from two subjects were discarded because of exceedingly high error rates (50% and 45%; mean error rate across all subjects, 14%). One additional subject selectively committed errors in low-reward trials so that we could not estimate the neural signal associated with low-reward trials in each run. This subject was excluded from the reward fMRI analysis.

Experimental paradigm

The experiment was implemented using MATLAB version 8.1.0 (MathWorks) and the Cogent Toolbox (<http://www.vislab.ucl.ac.uk/cogent.php>). In each trial, subjects were asked to judge either the parity or the magnitude of a number presented on screen. Before executing the task, subjects were presented with two cues, separated by a long delay. The first cue (reward mapping cue) consisted of two different symbols (Reverberi et al., 2012a; Fig. 1A) that did not indicate the task itself but determined which reward was associated with the parity and which reward was associated with the magnitude task. In half of the trials, the first cue indicated “If you are later instructed to perform the parity task, you can earn a high reward. If you are later instructed to perform the magnitude task, you can earn a low reward” (task–reward mapping 1). In the other half of the trials, the first cue indicated “If you are later instructed to perform the parity task, you can earn a low reward. If you are later instructed to perform the magnitude task, you can earn a high reward” (task–reward mapping 2). Each of the two symbols informed the subjects about the reward value associated with one of the two tasks, e.g., one symbol indicated “If parity, then low reward,” whereas the other symbol indicated “If magnitude, then high reward.” This information was not fully redundant, because we also introduced catch trials (20% of the trials) in which both tasks were associated with the same reward value or only one task was associated with a reward value. This was introduced to encourage subjects to represent the full task–reward association in each trial. One symbol was presented above and the other below the fixation point. Symbol position was pseudorandomized in each trial. Subjects likely processed the upper cue first, with the upper cue indicating the high-reward task in some trials and the low-reward task in other trials. This procedure discouraged subjects from retrieving task–reward associations with a fixed order (e.g., retrieving the highly rewarded task first). Furthermore, the same task–reward mapping could be cued by two visually different but semantically identical cues to later dissociate the neural signal of the mappings from the neural signal of the visual identity of the cues (for details, see below). Task–reward mapping cues were presented for 2000 ms, followed by a delay of 6000 ms (delay 1), during which a blank screen was shown. The delay duration was set to 1000 ms in some of the catch trials to encourage subjects to encode the task–reward associations early in the delay phase (for a similar approach, see Reverberi et al., 2012b). Catch trials were discarded from all analyses.

After the delay, a task cue was presented for 1500 ms, instructing subjects whether to judge the parity or the magnitude of the digit subsequently presented on screen. Again, two different abstract cues were associated with each task to decorrelate task signals from the visual identity of the cues. Half of the trials in which subjects were presented with task–reward mapping 1 were followed by parity cues, whereas the other half were followed by magnitude cues. The same was true for trials in which subjects were presented with task–reward mapping 2. This ensured that task–reward mappings, tasks, and reward values were independent in the experiment. Importantly, to know how much reward could be obtained in the current trial, subjects had to combine the information provided by the task–reward mapping and task cues, because the current reward condition was never explicitly or separately cued. Before the onset of the task cue, subjects could neither prepare the execution of the instructed task nor could they expect a high or a low reward, leaving only the mapping information to be represented before the task cue onset. Using time-resolved MVPA methods, we were able to test this hypothesis (for details, see below).

The task cue was followed by another delay of 1500 ms (delay 2), after which the task stimuli were presented. A single digit (one to nine, excluding five) was presented centrally on screen. Subjects performed the instructed task on each numerical stimulus, judging either whether the digit was odd or even (parity task) or whether it was larger or smaller than five (magnitude task). Four response symbols (Momennejad and Haynes, 2013) were presented together with the digit, and subjects were instructed to respond as accurately and quickly as possible. The response mapping was pseudorandomized in each trial to avoid motor preparation of responses before the onset of the task screen. Each stimulus was presented for 2000 ms, regardless of the reaction time (RT), and was followed by a brief delay (300 ms) and a reward feedback (500 ms). The

The authors declare no competing financial interests.

Correspondence should be addressed to David Wisniewski, Charité University Medicine Berlin, Bernstein Center for Computational Neuroscience, Haus 6, Philippstrasse 13, D-10115 Berlin, Germany. E-mail: david.wisniewski@bccn-berlin.de.

DOI:10.1523/JNEUROSCI.4882-14.2015

Copyright © 2015 the authors 0270-6474/15/3512356-11\$15.00/0

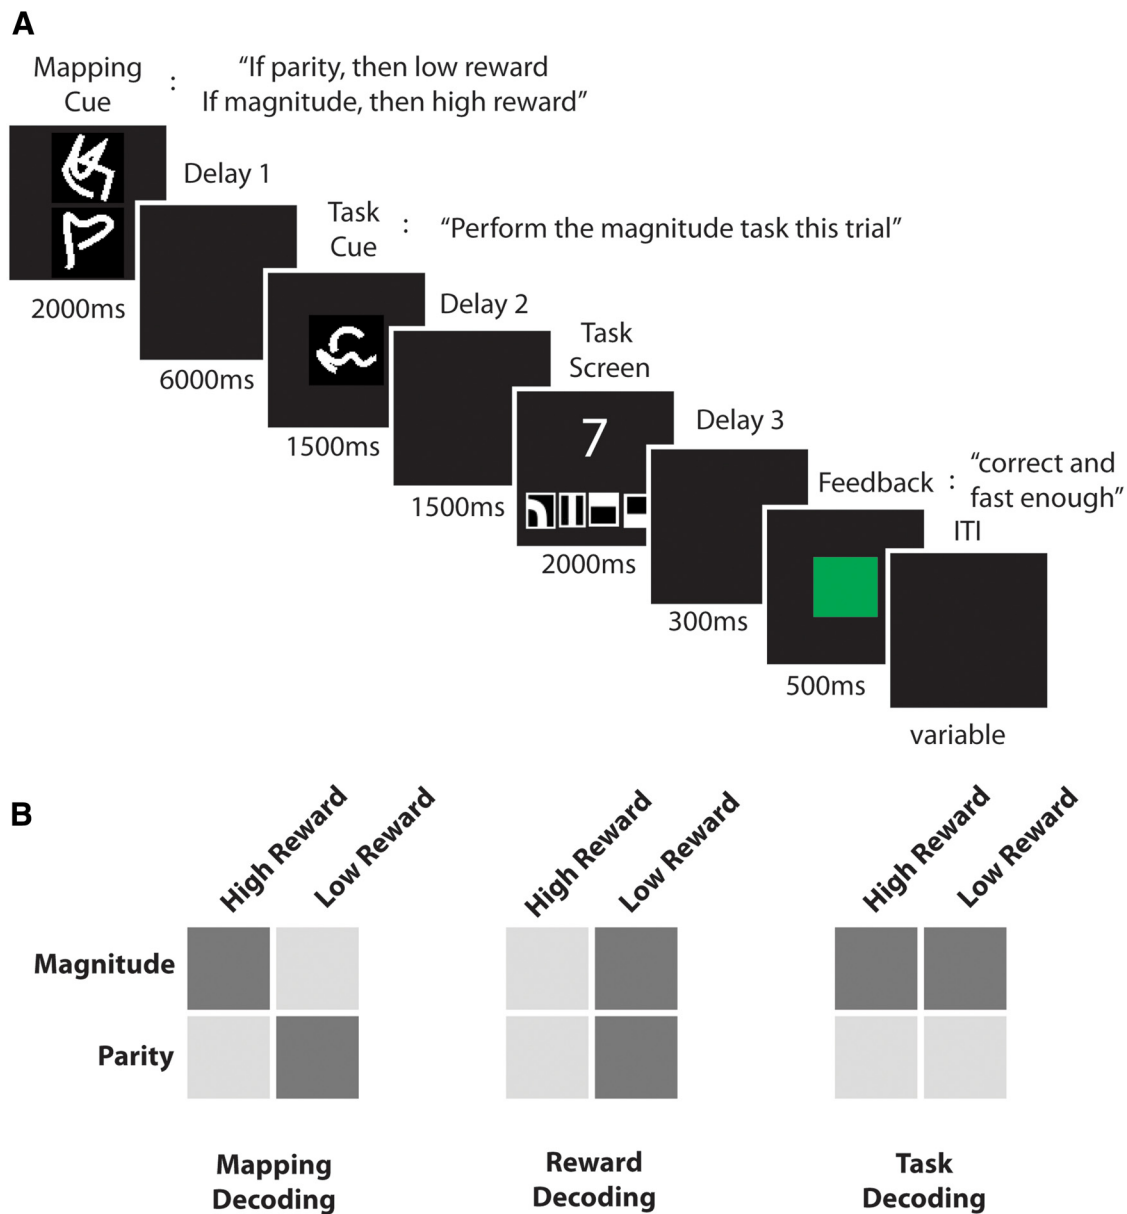


Figure 1. *A*, Experimental paradigm. One example trial is shown. First, subjects were presented with a mapping cue (2000 ms) that indicated the current association of tasks and rewards. In this trial, parity is associated with a low reward and magnitude is associated with a high reward. After a long delay (delay 1, 6000 ms) subjects were presented with a task cue (1500 ms) indicating which task is to be performed this trial. Note that only by combining information from both cues could subjects determine the current reward condition, in this case a high reward. Before the onset of the task cue, subjects could neither prepare the execution of the instructed task nor could they expect a high or a low reward, leaving only the mapping information to be represented before the task cue onset. After a second delay (delay 2, 1500 ms) the task screen was presented (2000 ms). A single digit was shown together with four response symbols (from left to right: odd, even, larger than 5, and smaller than 5), and subjects were instructed to press the correct button as quickly and accurately as possible. After a third short delay (delay 3, 300 ms), subjects received a reward feedback (500 ms), indicating whether they were correct and fast enough to receive a reward (green for high reward, yellow for low reward), were correct but too slow to receive a reward (magenta), were wrong (red), or did not press any button (red). After a variable intertrial interval (ITI, 2000–7000 ms), the next trial started. *B*, Analysis overview. Depicted are the three main MVPAs performed in this experiment. In the task–reward mapping analysis (left), task–reward mapping 1 (light gray) was contrasted with task–reward mapping 2 (dark gray). In the reward analysis (middle), the high-reward (light gray) and low-reward (dark gray) conditions were contrasted. In the task analysis (right), the parity task (light gray) was contrasted with the magnitude task (dark gray). In all analyses, we used cross-classification across visual cues to decorrelate the conditions of interest from the visual identity of the cues presented to the subjects.

reward feedback informed the subjects about the actual reward outcome of the trial (no reward, low reward, or high reward; for additional details, see below). The next trial started after a variable intertrial interval of 2000–7000 ms (mean, 3000 ms). Each participant performed a total of 280 trials across seven runs inside the scanner.

Subjects performed a separate training session a few days before the fMRI session. They were given 1.5 h to learn the meaning of all cues and familiarize themselves with the two tasks and the task screen. At the end of the training session, they performed two runs of the task to minimize learning effects during the fMRI session.

Reward manipulation and feedback

In each trial, subjects could either earn a high reward (10€) or a low reward (0.1€). To receive a reward, they had to perform the task quickly and accurately. We used an adaptive RT threshold to keep subjects constantly motivated. Specifically, we took the RTs from the last 25 trials, extracted the RT distribution, and used the 60th percentile of this distribution as the RT threshold in each trial. This would lead to subjects being too slow to earn a reward in ~40% of the trials, even if they pressed the correct button. For the first 25 trials of the fMRI session, we used the last training run to determine the RT threshold. Subjects were instructed that

they would have to react correctly and quickly to receive rewards, although they did not know the details of the thresholding procedure. They were also instructed that six random trials will be picked at the end of the experiment and that the rewards earned in these trials will be paid to them additionally to a performance-independent reimbursement. Subjects received a detailed reward feedback at the end of each trial using a color code, indicating whether their response was correct and fast enough (green in high-reward trials and yellow in low-reward trials), correct but too slow (magenta), wrong (red), or a miss (red).

Image acquisition

Functional imaging was conducted on a 3 tesla Siemens Trio scanner with a 12-channel head coil. For each run, 322 T2*-weighted echo-planar images were acquired (TR, 2000 ms; TE, 30 ms; flip angle, 90°). Each volume consisted of 33 slices, separated by a gap of 0.6 mm. Matrix size was 64 × 64, and field of view was 192 mm, which resulted in a voxel size of 3 × 3 × 3 mm³. The first three images of each run were discarded.

Data analysis: behavior

For each subject, task performance was assessed by calculating the mean RT and mean error rate across all runs. Errors were subdivided into the following: correct button but too slow, wrong button, or miss. We then refined the analysis by investigating RTs and error rates separately for different mappings, tasks, reward values, and visual cues using paired-sample *t* tests. We expected rewards to have a strong effect on both measures but did not expect the other variables to affect them.

It has been reported previously that switching between trials leads to an increase in RTs, which is partly attributable to task-set reconfiguration processes (Monsell, 2003). The so-called switch costs have also been shown to interact with the reward condition of the preceding trial (Braem et al., 2012). Therefore, we compared the average RTs of task switch and repeat trials and further tested whether these switch costs interacted with the reward manipulation in our task. To do so, we performed a three-factorial repeated-measures ANOVA on RTs and error rates, entering switch/repeat, the reward condition of the current trial, and the reward condition of the previous trial as factors.

Data analysis: fMRI

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 (Cox and Savoy, 2003; Haynes and Rees, 2006; Kriegeskorte et al., 2006; Haxby, 2012; Tong and Pratte, 2012), we performed three independent analyses to investigate the neural encoding of task–reward associations, reward effects, and tasks, respectively.

Analysis 1: MVPA of task–reward mappings. To examine the neural encoding of mappings, first a GLM was estimated for each subject. The following regressors were estimated in each run: (1) mapping 1 cued by visual cue 1; (2) mapping 1 cued by visual cue 2; (3) mapping 2 cued by visual cue 3; and (4) mapping 2 cued by visual cue 4. Movement parameters were added to the GLM to account for potential head movement during scanning. We also added RTs as a parametric modulation to the model. This was done to ensure that potential differences in difficulty (as measured by RTs) between the mappings were modeled explicitly and regressed out of the data (Todd et al., 2013; but see Woolgar et al., 2014). We used the finite impulse response (FIR) function as the basis function to get time-resolved estimates of the neural activity underlying task–reward mappings. Each regressor was modeled using 11 time bins with a duration of 2 s each. The first time bin was locked to the onset of the reward mapping cue in each trial. This allowed us to cover the whole trial and still be able to correct for the hemodynamic lag by shifting the analysis by two bins (or 4 s).

In the next step, a decoding analysis was performed on the parameter estimates of the GLM. One support vector classifier (SVC) was applied to each time bin separately. This approach allowed us to investigate the time course of information in each voxel (Soon et al., 2008; Bode and Haynes, 2009). The SVC applied a linear kernel using a fixed regularization parameter ($C = 1$) on the parameter estimates of the GLM of the time bin (Cox and Savoy, 2003; Mitchell et al., 2004; Kamitani and Tong, 2005;

Haynes and Rees, 2006), as implemented in LIBSVM (Chang and Lin, 2011; <http://www.csie.ntu.edu.tw/~cjlin/libsvm>). We applied a search-light 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 three voxels around each measured voxel in the acquired volumes. For each condition (task–reward mapping 1 cued by visual cue 1, task–reward mapping 1 cued by visual cue 2, task–reward mapping 2 cued by visual cue 3, task–reward mapping 2 cued by visual cue 4), we created a pattern vector by extracting parameter estimates for each voxel in the sphere. This was done for each run independently. The SVC was trained to discriminate between the two task–reward mappings, cued by visual cues 1 and 2 (Fig. 1B). Classification performance was tested on the independent contrast between task–reward mappings cued by visual cue 3 and 4. A second SVC was trained to discriminate between the two task–reward mappings cued by visual cues 3 and 4 and tested on the independent contrast of task–reward mappings cued by visual cues 1 and 2. This resulted in two accuracy values for each time bin. This cross-classification allowed us to disentangle the mapping signal from signals encoding the visual shape of the cues presented on screen (Reverber et al., 2012a). This procedure further mimicked the way subjects learned to associate cues and mapping conditions during the training session. Subjects were taught that cues 1 and 2 were one “cue set,” whereas cues 3 and 4 were a different “cue set.” The cross-classification did not mix data from different cue sets, training on one such cue set and testing on the other. Furthermore, training the SVC on one condition and testing it on an independent condition was necessary to control for potential problems of overfitting. The classification was repeated for every sphere in the measured brain volume, resulting in two three-dimensional classification accuracy maps for each bin in each subject. The resulting accuracy maps were normalized to a standard brain [Montreal Neurological Institute (MNI) EPI template as implemented in SPM8] and resampled to an isotropic resolution of 3 × 3 × 3 mm. Normalized images were smoothed with a Gaussian kernel with 6 mm FWHM to account for differences in localization across subjects.

Then, a random-effects group analysis was performed on the accuracy maps, using a 2 (visual cue training set) × 11 (time bin) ANOVA. The two levels of the visual cue training set factor were as follows: training on visual cues 1/2 and testing on visual cues 3/4 versus training on visual cues 3/4 and testing on visual cues 1/2. In this ANOVA, two contrasts were computed: (1) one testing for above-chance accuracies in the time bins encompassing the mapping cue and delay 1 (time bins 3–6, called “delay 1”) across both visual cue training sets; and (2) the other testing for above-chance accuracies in all succeeding phases, i.e., task cue, delay 2, task execution, and feedback (time bins 7–9, called “task execution”) across both visual cue training sets. In each condition, the first two time bins were not analyzed to account for the hemodynamic lag in the data. Because the SVC was performed on two conditions, the chance level was 50%. We applied a statistical threshold of $p < 0.05$ (FWE corrected at the cluster level, initial threshold $p < 0.001$).

To investigate the temporal development of mapping information, we extracted voxel coordinates from all significant clusters. Then, we determined the mean decoding accuracy across subjects in these clusters in each time bin. We also extracted the accuracy time course in the same task–reward mapping ROIs for the task and reward MVPA analyses, respectively (see below). This allows us to compare information about different aspects of the tasks throughout the trial within the same ROI. fMRI is not a technique ideally suited to resolve decoding time courses with a high temporal resolution. Therefore, one should be careful interpreting the exact onsets and offsets of accuracy time courses down to a single time bin. For instance, the temporal interpolation during slice-time correction might shift the onset or offset of an accuracy time course by half a time bin. However, for our purposes, it is enough to distinguish more broadly between the first part and the second part of the trial. In the first part, before the task cue onset, subjects only had information about the mapping of both tasks to the reward conditions. They could not prepare for the execution of a specific task and could not expect to gain a high or low reward yet. In the second part, after the task cue onset, subjects knew which task was to be executed and which reward they could expect and could therefore prepare for the upcoming execution of the

task. As an additional test of the validity of the ROI analysis, we applied an additional shape constraint in the two contrasts used to define significant regions. To add greater physiological plausibility we fitted a gamma function (as implemented in SPM) to the contrast vectors. We then extracted the significant clusters and computed the accuracy time courses in these clusters, as before. We found no apparent differences between the results.

As stated above, there might be some effects of the preceding trial on performance of the current trial (Braem et al., 2012), and therefore, we decided to conservatively control for the influence of the preceding trial conditions (previous mapping, previous reward, and previous task) in the neural data of the current trial by explicitly modeling them in our GLM analysis (as was done with the current RT). Controlling data carefully is important especially for sensitive multivariate analyses (Todd et al., 2013; Woolgar et al., 2014). Therefore, we repeated the whole analysis described above an additional three times, only changing the parametric modulator from the current RT to the previous mapping, previous reward, or previous task, accounting for and regressing out any variance in the data attributable to these conditions. After performing the decoding and group-level analyses that were identical to the ones described above, we had four group-level contrast maps for the mapping decoding analysis. We computed the overlap of those four maps, and only areas in which all four models showed significant above-chance decoding accuracies are interpreted. In contrast to simply averaging the resulting accuracy maps, this procedure ensured that the influence of the previous trial conditions and the current RT was minimized. If for instance the current RT were to affect the decoding, by averaging, we would aggregate one analysis in which this was controlled for with three analyses in which this was not controlled for, and the resulting average map would still represent results that would be partly correlated with RTs. Computing the overlap avoids this pitfall. We also want to point out that the reason for running four independent analyses instead of one analysis incorporating all control regressors is that using the FIR basis function leads to the estimation of a very large number of parameters in each run. Adding all regressors in one analysis would likely lead to a model that cannot be estimated. Therefore, our approach ensures a balance between high quality of the results and being able to estimate all parameters.

Analysis II: MVPA of reward effects. This analysis was analogous to the analysis of task–reward mappings. For each subject, we estimated a GLM with the following regressors: (1) high reward cued by visual cue 1; (2) high reward cued by visual cue 2; (3) low reward cued by visual cue 3; and (4) low reward cued by visual cue 4. Movement parameters were added, and RTs were explicitly modeled as parametric modulations. All regressors were locked to the onset of the mapping cue. However, the actual reward condition of each trial was only revealed to the subject after a delay, with the presentation of the task cue. Before that point in time, subjects could not expect to gain a high or a low reward. Note that we used the potential reward outcome, not the actually received reward in this analysis, because the latter is obviously strongly dependent in the performance of the trial. It should also be mentioned that results in this analysis could be driven by either the reward value or the processes associated with the value signal, such as attentional preparation (Kahnt et al., 2014). Again, we used the FIR as the basis function with the same parameters as in the mapping analysis. A searchlight decoding analysis was performed on the results of the GLMs, as in the mapping analysis (Fig. 1B). The resulting accuracy maps were entered into a random-effects group analysis, and statistical testing was performed using a two-factorial ANOVA (factor 1, time bins; factor 2, visual cue training set). The same two contrasts were computed for delay 1 and task execution. We applied a statistical threshold of $p < 0.05$ (FWE corrected at the cluster level, initial threshold $p < 0.001$). Also, as in the task–reward mapping analysis, we repeated the whole analysis procedure three more times, adding the previous trial conditions (mapping, reward, and task) as control regressors to the GLM, respectively. Again, the overlap of all four analyses was computed, and only regions surviving this control procedure are interpreted.

Analysis III: MVPA of tasks. For the task analysis, a GLM was estimated for each subject using the following regressors: (1) parity cued by visual cue 1; (2) parity cued by visual cue 2; (3) magnitude cued by visual cue 3; and (4) magnitude cued by visual cue 4. Movement parameters were

added, and RTs were explicitly modeled as parametric modulations. All regressors were locked to the onset of the mapping cue. As in the reward-effect analysis, the task to be performed in each trial was not revealed until later in the trial, when the task cue was presented. Before that point in time, subjects could not prepare for the execution of a specific task. Again, time-resolved searchlight decoding was applied to the parameter estimates as in the mapping analysis. The resulting accuracy maps were entered into a random-effects group analysis, and two contrasts were computed (delay 1 and task execution). The maps were thresholded at $p < 0.05$ (FWE corrected at the cluster level, initial threshold $p < 0.001$). As for the task decoding, we repeated the whole analysis, controlling for the influence of the previous trial conditions (mapping, reward, and task). We computed the overlap of all four maps and only interpret regions surviving this conservative control procedure.

Results

Behavioral

The mean \pm SE RT across all trials was 963 ± 29 ms. Subjects were correct and fast enough in $56 \pm 0.8\%$ of the trials, correct but too slow to meet the RT threshold in $34 \pm 0.9\%$ of the trials, wrong in $9 \pm 1.3\%$ of the trials, and missed the response window in $1.1 \pm 0.8\%$ of the trials. The fact that subjects were correct and fast enough in $\sim 60\%$ of the trials validates our RT thresholding procedure. For all further analyses, trials in which the correct button was pressed were considered correct, regardless of the RT. Trials in which the wrong or no button was pressed are considered error trials and are excluded from all further analyses.

We then analyzed whether there were any differences in RTs between mappings, tasks, visual cues, or rewards. We found no difference in RTs between mappings ($t_{(16)} = 1.02$, $p = 0.32$) nor between tasks ($t_{(16)} = 1.9$, $p = 0.07$). For the latter test, the observed p value was very close to the significance threshold. To determine how task effects on RTs could be resolved in a future experiment, we calculated the sample size needed to detect this effect with a probability (power) of 0.80, given the current α level of 0.05. Results indicate that such a study would need 39 subjects to resolve this effect. With the current sample, the results remain somewhat inconclusive, because they provide no strong evidence for either the existence or non-existence of an effect.

We also found no RT difference between the two visual cue conditions ($t_{(16)} = 0.2$, $p = 0.81$). This demonstrates that there were no strong differences in the difficulty between these conditions. As expected, we found a significant difference in RTs between the reward conditions ($t_{(16)} = 3.4$, $p = 0.003$). High-reward trials were on average 74 ms faster than low-reward trials, showing that subjects invested more effort into high-reward trials. This is an important finding because the only source of information about the reward condition was the mapping cue. By demonstrating a behavioral effect of reward, we therefore also demonstrate that subjects attended to the mapping cue and used the information provided by it. Note that this effect cannot be explained by different RT thresholds between the two reward conditions, because this difference was negligible (6 ms). No difference was found between mappings, tasks, visual cues, or rewards in the error rates (all $t_{(16)}$ values < 1.52 , all p values > 0.15).

We then assessed the task switch costs by subtracting the average RTs in task repeat trials from the average RT in task switch trials. There was a significant switch cost ($t_{(16)} = 7.07$, $p < 0.001$; the average \pm SE switching cost was 56 ± 8 ms), which likely reflects residual switch costs that are still detectable after long preparatory periods (Monsell, 2003). To assess task switching costs and their interaction with reward values in more detail, we performed a three-factorial ANOVA on RTs and error rates, using the following factors: (1) task switch/repeat; (2) reward value

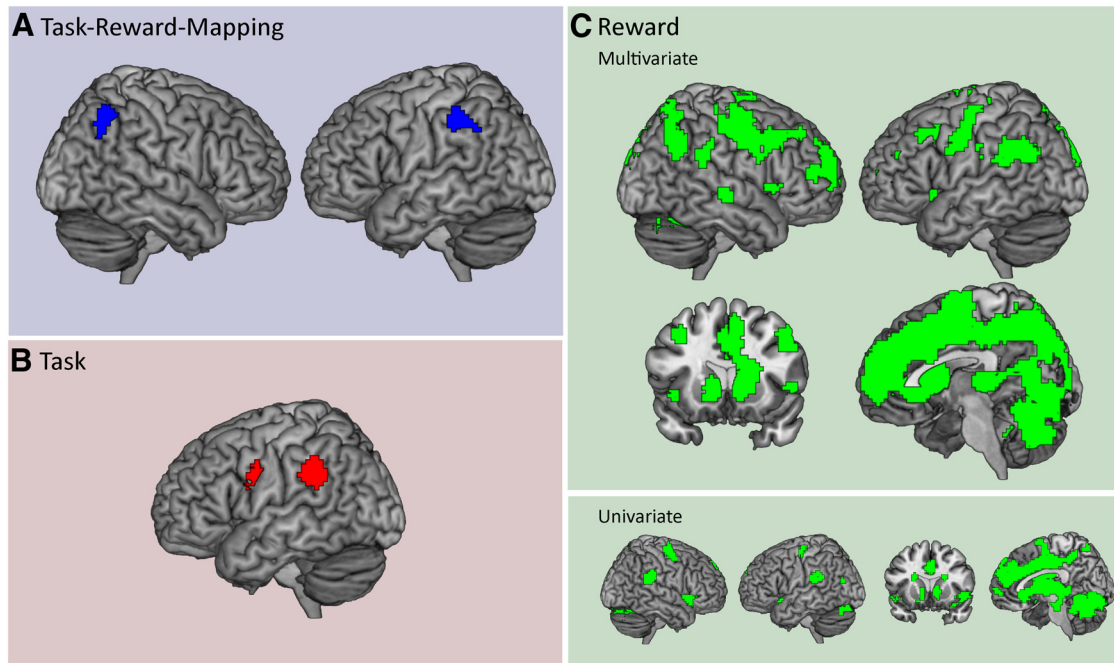


Figure 2. Depicted are the results from all decoding analyses. **A**, Task–reward mapping decoding. It can be seen that two clusters in the bilateral inferior parietal cortex represented task–reward mappings during the presentation of mapping cues and delay 1. **B**, Task decoding. The left inferior parietal cortex and left premotor cortex represented tasks during the presentation of the task cue, delay 2, task execution, and reward feedback. **C**, Above, results from the reward-effects decoding analysis are shown. It can be seen that a large cortical and subcortical network shows reward effects during the presentation of the task cue, delay 2, task execution, and reward feedback. This network included striatal, medial prefrontal, and parietal regions. Below, brain regions showing a univariate signal increase in high reward trials during the same period in the trial are shown for comparison.

of the current trial; and (3) reward value of the previous trial. For this analysis, trials following error trials were also excluded. Given previous findings (Braem et al., 2012), one might expect that task switch costs in our experiment interacted with the reward value of the previous trial. Results showed a significant main effect of task switching ($F_{(1,16)} = 52.9, p < 0.001$) and a significant main effect of the current reward value ($F_{(1,16)} = 19.1, p < 0.001$). Neither the effect of the previous reward value nor any interaction reached significance (all $F_{(1,16)}$ values < 2.53 , all p values > 0.13). Thus, our findings do not replicate previous results on the interaction of rewards and task switch costs with our experimental setup (Braem et al., 2012). An ANOVA of the error rates yielded no significant results (all $F_{(1,16)}$ values < 4.13 , all p values > 0.06). Again, we calculated the sample size needed to resolve this effect in a future experiment with a probability (power) of 0.80 and an α level of 0.05. Results indicate that such a study would need 72 subjects to resolve this effect. With the current sample, the results remain somewhat inconclusive, because they provide no strong evidence for either the existence or non-existence of an effect. Future research will investigate the relationship of task switch costs and rewards more directly.

Analysis I: representation of task–reward associations

During task–reward mapping cue presentation and delay 1, distributed activation patterns in the bilateral inferior parietal cortex (Brodmann area 40) contained information about the task–reward mappings (Fig. 2A, Table 1). During task cue presentation, delay 2, and task execution, no significant information was found. In a next step, we repeated the analysis, using a simpler model that did not explicitly control for the influence of the previous trial conditions and current RT. This analysis was performed to test our control procedure. In fact, results were highly similar to the original analysis, with the bilateral inferior parietal

Table 1. Summary of MPPA results

Brain region	Side	Cluster size	MNI coordinates (peak voxel)		
			x	y	z
Task–reward associations					
Left inferior parietal cortex	Left	67	−48	−52	37
Right inferior parietal cortex	Right	73	57	−64	34
Tasks					
Ventral premotor cortex	Left	80	−42	2	19
Inferior parietal cortex	Left	134	−54	−43	22
Rewards					
Inferior parietal cortex	Left	300	−45	−43	19
Inferior parietal cortex	Right	56	57	−40	28
Medial cortex	Bilateral	7736	6	−4	67
Precentral gyrus	Left	305	−30	−19	49
Midfrontal gyrus	Left	74	−36	14	43
Superior temporal gyrus	Right	31	60	−19	−5
Anterior midfrontal gyrus	Right	35	42	47	13
Anterior midfrontal gyrus	Left	59	−27	44	22
Inferior frontal gyrus, pars triangularis	Left	28	−30	35	7
Inferior frontal gyrus, operculum	Right	26	48	14	1
Cerebellum	Right	231	27	−52	−32

Results are shown for a statistical threshold of $p < 0.001$, corrected for multiple comparisons at the cluster level ($p < 0.05$). The regions shown for the task–reward association results encoded these associations before the onset of the task cue. The regions shown for the task and reward results encoded these variables after the onset of the task cue. All clusters larger than 20 voxels are reported.

cortex representing task–reward associations ($p < 0.05$, FWE corrected at the cluster level, initial threshold $p < 0.001$). We then extracted time-resolved decoding accuracies from both parietal clusters for the task and reward-effects analyses (Fig. 3A, for details, see below). As soon as the task cue was presented and the current task and reward conditions were revealed to the subjects, they could start preparing for the task execution and could expect to gain either a high or a low reward. At this point in time, both

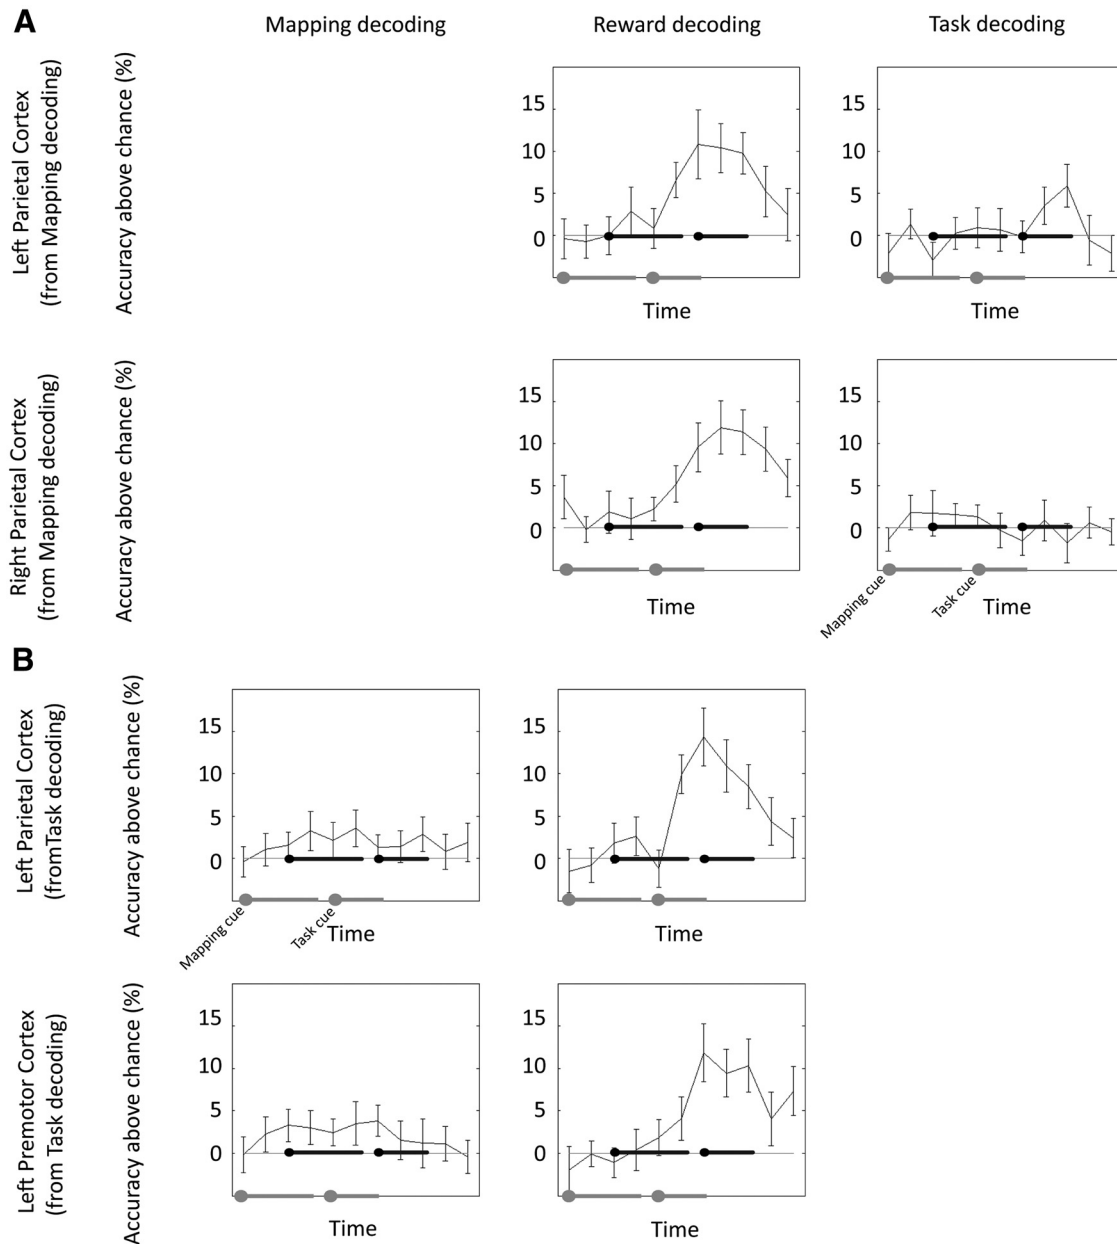


Figure 3. *A*, Time course of information for the two brain regions identified in the task–reward mapping analysis (left parietal cortex and right parietal cortex). The left gray line shows the onset of the mapping cue and the duration of the delay 1. The right gray line shows the onset of the task cue, and the duration of the delay 2, task execution, and reward feedback periods. The two black lines represent the same events with a hemodynamic lag of 4 s factored in. Error bars represent SEs. Four independent decoding analyses were performed, each correcting for a different possible confounding variable. The plot shows the minimum accuracy value of those four analyses at each time point. Because the mapping analysis has been used to define the ROIs, data from that analysis are not shown to avoid circularity. Task and reward accuracy time courses are independent of the data used to define the ROIs. *B*, Time course of information for the two brain regions identified in the task decoding analysis (left parietal cortex, left premotor cortex). Because the task analysis has been used to define the ROIs, data from that analysis are not shown to avoid circularity. Mapping and reward accuracy time courses are independent of the data used to define the ROIs.

parietal regions changed from representing mappings to representing rewards or their effects on preparatory processes. Additionally, the left parietal cortex also represented the task, while it was being performed. This is also supported by the overlap of mapping, task, and reward effects-related brain regions depicted in Figure 4. To ensure that this overlap is not merely induced by spatial smoothing of the accuracy maps, we repeated the group-level analysis using unsmoothed data. In this analysis, we found an overlap of mapping and task-related brain regions as well, indicating that spatial smoothing did not have a strong effect on the results depicted in Figure 4. However, this analysis does not

test directly whether there is an overlap in single-subject space. Although unlikely, we cannot fully rule out that the parietal regions identified are neighboring but non-overlapping in individual subjects. Our results indicate that the inferior parietal cortex can flexibly reconfigure its function within a trial to meet changing task demands (Rao et al., 1997; Sigala et al., 2008; Stokes et al., 2013).

One possible explanation of this result is that subjects had a bias toward preparing the task associated with a high reward in each trial instead of representing the whole task–reward mapping. This would help them being prepared for the highly reward-

ing task if it was cued later in the trial. If this were the case, the neural pattern that encodes mappings should be similar to the pattern encoding tasks. Therefore, we trained a classifier to distinguish between the task–reward mapping conditions early in the trial before the task cue was presented. We then used this classifier to distinguish task patterns (parity vs magnitude) later in the trial after the task cue was presented. We found no significant results ($p < 0.05$, FWE corrected at the cluster level, initial threshold $p < 0.001$), suggesting that subjects did not simply prepare the highly rewarded task in each trial. As an additional confound control, we tested whether the cue set training direction showed any significant results. We found no significant results for the cue set factor ($p > 0.05$, FWE corrected at the cluster level, initial threshold $p > 0.001$), further validating our cross-classification procedure.

Analysis II: representation of tasks

As expected, we found no information about tasks during mapping cue presentation and delay 1, before the onset of the task cue. During task cue presentation, delay 2, and task execution, local spatial activation patterns in the left inferior parietal cortex (Brodmann area 40) and left premotor cortex (Brodmann area 6) contained information about the currently performed task (Fig. 2*B*). We then extracted the voxel coordinates for both clusters and investigated the temporal development of decoding accuracies in these clusters for the mapping and reward-effects decoding analyses (Fig. 3*B*). It can be seen that both clusters also represented reward effects in the second half of the trial, while simultaneously representing tasks. Again, this finding demonstrates the flexibility of the parietal cortex to flexibly adapt to changing task demands.

Analysis III: reward effects

As expected, the MVPA found no effect of reward during mapping cue presentation and delay 1. As soon as the task cue was presented, local spatial activation patterns predicted the reward, or its effects, of the trial in a widespread network, including subcortical striatal areas, as well as prefrontal, parietal, temporal, and occipital brain regions (Fig. 2*C*). The reward MVPA results can be partly explained by the presence of a clear signal difference between high-reward and low-reward trials (Fig. 2*C*) at the single-voxel level (univariate analysis, $p < 0.05$ voxelwise FWE corrected). However, it is unlikely the MVPA results reflect a “pure” reward value signal. The reward manipulation likely led to widespread preparatory processes in, for example, attentional or motor systems. Furthermore, it might increase the salience of the stimuli (cf. Kahnt et al., 2014). Such effects are reflected in the MVPA results as well. We call these findings “reward” effect simply because this was the underlying experimental manipulation that elicited this effect.

Discussion

Adaptation and decision-making often rely on the abstract knowledge on how a number of complex behaviors (e.g., jobs) would be rewarded under different circumstances (e.g., salaries under different employers). In this experiment, we investigated




Figure 4. Overlap of all three MVPAs. Regions informative about the mapping during delay 1 are shown in red, regions informative about the task during task execution are shown in blue, and regions informative about the reward during task execution are shown in green. The informative regions of all three analyses overlap in the left inferior parietal cortex, whereas regions informative about mappings and rewards overlap in the right inferior parietal cortex. This shows that brain areas that represent an association between tasks and rewards early in the trial represent actual tasks and rewards or their effects at a later stage in the trial.

the representation of task–reward associations using fMRI and multivariate pattern classification. We used time-resolved searchlight decoding to identify which brain areas represented task–reward associations and tasks and which brain areas processed rewards across each trial. Importantly, we were able to separate patterns of activity underlying task–reward associations from those underlying task preparation and reward effects. We hypothesized that parietal and/or medial prefrontal brain areas play a role in processing task–reward associations. We found them to be encoded in local spatial activation patterns in the bilateral inferior parietal cortex, whereas tasks were processed in the left inferior parietal and left premotor cortices. Reward effects, such as attentional preparation, were found in a widespread cortical and subcortical network, including striatal, medial prefrontal, and parietal brain regions. Interestingly, we found that the inferior parietal cortex is highly flexible in adapting to the current task demands within a trial. It first represented task–reward mappings and later switched to processing rewards as soon as information about them became available. The left inferior parietal cortex additionally encoded which of two tasks were performed during their execution.

Task–reward associations in the parietal cortex

We found the bilateral inferior parietal cortex to represent associations between tasks and rewards in our experiment. Converging behavioral results demonstrated a strong effect of the reward condition on both RTs and behavioral accuracies. Given that the task–reward association cue was the only source of information about the current reward condition, both results suggest that subjects attended to and processed the task–reward association cues in the beginning of each trial. However, it can be difficult to distinguish the knowledge of rules linking multiple behaviors with rewards from task preparation and reward expectation processes. To separate the neural signals arising from task–reward associations from task and reward-related signals, we orthogonalized these variables in our design. Furthermore, we separated these variables in time as well, using a double-cueing procedure. Therefore, our MVPA results in the parietal cortex represent the

abstract knowledge of which tasks are associated with which rewards only.

The parietal cortex plays a role both in the processing of rewards and tasks, which will be discussed in turn, and is therefore a likely candidate to represent associations between them as well. Previous research demonstrated that neural activity in the parietal cortex in nonhuman primates correlates with the relative subjective desirability of actions in strategic games (Dorris and Glimcher, 2004) and during foraging (Sugrue et al., 2004) and also represents expected outcomes and their probability (Platt and Glimcher, 1999). It is also important for the processing of future rewards in humans (Tanaka et al., 2004), playing an important role in guiding decisions based on reward values. Our results are mostly in line with these findings, because the parietal cortex represented the potential outcomes of two alternative tasks, while also being involved in processing current rewards directly. Note that our experimental manipulation of rewards likely led to widespread preparatory processes, for example, recruiting attentional resources in high-reward trials. We do not claim that the reward regions in our experiment represent a pure value signal. In fact, a recent theory suggests that rewards have three different effects on behavior (Peck et al., 2009): (1) biasing attention; (2) triggering conditioned responses; and (3) motivating to increase performance. Rewards in our experiment likely had all three effects, but the reward effect in the inferior parietal cortex is best explained with changing the attentional bias (Peck et al., 2009; Kahnt et al., 2014).

We have further shown that the left lateral parietal cortex represents which of the two alternative numerical tasks is performed in the current trial. This finding supports previous evidence that highlights the role of the lateral parietal cortex in task processing (Bode and Haynes, 2009; Reverberi et al., 2012a, Woolgar et al., 2015). Given that these studies used very different tasks, the parietal cortex might play a more general role in representing task sets.

In summary, we have shown that the inferior parietal cortex has a key role in directly representing associations between tasks and rewards. This extends previous findings focusing on the role of the inferior parietal cortex in processing tasks or rewards. More broadly, we demonstrate that the parietal cortex directly links cognitive control and motivational functions in the brain, which further extends previous findings on its role in cognitive control (Brass et al., 2005; Yeung et al., 2006; Esterman et al., 2009; Shomstein, 2012; Power and Petersen, 2013). Although we demonstrate that task–reward associations are partly represented in local spatial activity patterns, these associations need not be exclusively represented locally. It is possible that the same information is also contained in distributed connectivity patterns between brain regions as well. Future research will address this question more directly.

Flexible reconfiguration of parietal cortex function

In each trial, subjects were first provided with information about the current task–reward association condition, specifying how both tasks could be rewarded. Only after a delay was the actual task specified, and only at this point in time could subjects have prepared for executing a specific task and expected a specific reward. This allowed us to dissociate neural signals arising from task–reward associations from those related to task preparation and reward processing, using time-resolved searchlight decoding. We demonstrated that the parietal cortex first represented task–reward associations, switching to reward processing once

the reward condition was specified. This highlights the ability of the parietal cortex to flexibly reconfigure its functions on relatively short timescales.

It has been argued that the parietal cortex is part of a larger network that flexibly reconfigures its function, depending on the current task demands (Dehaene et al., 1998; Duncan, 2010; Fedorenko et al., 2013). Previous studies in nonhuman primates showed that neurons, mainly in prefrontal brain regions, can even flexibly reconfigure their function within a trial, as task demands change from one trial phase to the next (Rao et al., 1997; Sigala et al., 2008; Stokes et al., 2013). Our results corroborate the idea of flexible reconfiguration within trials by demonstrating that the inferior parietal cortex behaves similarly in humans, as shown using time-resolved MVPA. It represents the relevant information for the subject at each point in time, first the abstract task–reward association and later the actual task and reward conditions, rapidly and flexibly changing the content of representation.

Brain regions related to task–reward associations represent task sets

More generally, one might speculate that task–reward associations are represented in brain areas closely related to the processing of the specific task currently performed (for a similar argument, see Reverberi et al., 2012b). As shown above, the parietal cortex is generally well suited to represent and process task sets (Bode and Haynes, 2009; Reverberi et al., 2012a), and this might be the reason why we found it to represent task–reward associations as well. In fact, the left parietal cluster identified in the mapping decoding analysis also seems to represent task information directly in our experiment, which is compatible with this idea. Converging evidence for this hypothesis also comes from nonhuman primates, in which action outcomes (juice rewards) in an oculomotor task were represented in a brain area closely linked to eye movements (Platt and Glimcher, 1999). Furthermore, research on action–reward associations highlighted the role of the dorsal anterior cingulate cortex (dACC; Kennerley et al., 2006; Alexander and Brown, 2011; Hayden et al., 2011; Shenhav et al., 2013), which is also an important region for goal-directed action selection (Ridderinkhof et al., 2004) and therefore well suited to process actions. Interestingly, we did not find the dACC to represent task–reward associations. We can only speculate why we did not see the dACC, but it might have to do with the tasks we used in our design, which relied more strongly on parietal cortex functioning (Fig. 2) and not as much on dACC. Future research will have to show whether the parietal cortex also represents task–reward associations in task relying on brain regions different from the parietal cortex.

Conclusion

In this experiment, we investigated the neural processing of tasks, rewards, and task–reward associations. We found all three variables to rely on inferior parietal cortex functioning. Importantly, we orthogonalized these three variables and further dissociated them in time to investigate their neural correlates independently. Interestingly, the inferior parietal cortex flexibly changed its content of representation on a short timescale within trials, from task–reward associations early in the trial to reward effects and tasks later in the trial. This demonstrates that the parietal cortex is a key area to link cognitive control and motivational functions in the brain, while showing a remarkable flexibility in adapting its function to current demands.

References

- Alexander WH, Brown JW (2011) Medial prefrontal cortex as an action–outcome predictor. *Nat Neurosci* 14:1338–1344. [CrossRef Medline](#)
- Bode S, Haynes JD (2009) Decoding sequential stages of task preparation in the human brain. *Neuroimage* 45:606–613. [CrossRef Medline](#)
- Braem S, Verguts T, Roggeman C, Notebaert W (2012) Reward modulates adaptations to conflict. *Cognition* 125:324–332. [CrossRef Medline](#)
- Brass M, Ullsperger M, Knoesche TR, von Cramon DY, Phillips NA (2005) Who comes first? The role of the prefrontal and parietal cortex in cognitive control. *J Cogn Neurosci* 17:1367–1375. [CrossRef Medline](#)
- Chang CC, Lin CJ (2011) LIBSVM: a library for support vector machines. *ACM Trans Intell Syst Technol* 2:1–27.
- Chudasama Y, Daniels TE, Gorrin DP, Rhodes SEV, Rudebeck PH, Murray EA (2013) The role of the anterior cingulate cortex in choices based on reward value and reward contingency. *Cereb Cortex* 23:2884–2898. [CrossRef Medline](#)
- 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. [CrossRef Medline](#)
- Daw ND, O’Doherty JP, Dayan P, Seymour B, Dolan RJ (2006) Cortical substrates for exploratory decisions in humans. *Nature* 441:876–879. [CrossRef Medline](#)
- Dehaene S, Kerszberg M, Changeux JP (1998) A neuronal model of a global workspace in effortful cognitive tasks. *Proc Natl Acad Sci U S A* 95:14529–14534. [CrossRef Medline](#)
- Dorris MC, Glimcher PW (2004) Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* 44:365–378. [CrossRef Medline](#)
- Duncan J (2010) The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. *Trends Cogn Sci* 14:172–179. [CrossRef Medline](#)
- Esterman M, Chiu YC, Tamber-Rosenau BJ, Yantis S (2009) Decoding cognitive control in human parietal cortex. *Proc Natl Acad Sci U S A* 106:17974–17979. [CrossRef Medline](#)
- Etzel JA, Cole MW, Zacks JM, Kay KN, Braver TS (2015) Reward motivation enhances task coding in frontoparietal cortex. *Cereb Cortex*. Advance online publication. Retrieved July 23, 2015. doi:10.1093/cercor/bhu327. [CrossRef Medline](#)
- Fedorenko E, Duncan J, Kanwisher N (2013) Broad domain generality in focal regions of frontal and parietal cortex. *Proc Natl Acad Sci U S A* 110:16616–16621. [CrossRef Medline](#)
- 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. [CrossRef](#)
- Hadland KA, Rushworth MFS, Gaffan D, Passingham RE (2003) The anterior cingulate and reward-guided selection of actions. *J Neurophysiol* 89:1161–1164. [CrossRef Medline](#)
- Hampton AN, O’Doherty JP (2007) Decoding the neural substrates of reward-related decision making with functional MRI. *Proc Natl Acad Sci U S A* 104:1377–1382. [CrossRef Medline](#)
- Haxby JV (2012) Multivariate pattern analysis of fMRI: the early beginnings. *Neuroimage* 62:852–855. [CrossRef Medline](#)
- 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. [CrossRef Medline](#)
- Haynes JD, Rees G (2006) Decoding mental states from brain activity in humans. *Nat Rev Neurosci* 7:523–534. [CrossRef Medline](#)
- 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. [CrossRef Medline](#)
- Holroyd CB, Yeung N (2012) Motivation of extended behaviors by anterior cingulate cortex. *Trends Cogn Sci* 16:122–128. [CrossRef Medline](#)
- Kahnt T, Heinzle J, Park SQ, Haynes JD (2010) The neural code of reward anticipation in human orbitofrontal cortex. *Proc Natl Acad Sci U S A* 107:6010–6015. [CrossRef Medline](#)
- Kahnt T, Park SQ, Haynes JD, Tobler PN (2014) Disentangling neural representations of value and salience in the human brain. *Proc Natl Acad Sci U S A* 111:5000–5005. [CrossRef Medline](#)
- Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8:679–685. [CrossRef Medline](#)
- Kennerley SW, Walton ME, Behrens TEJ, Buckley MJ, Rushworth MFS (2006) Optimal decision making and the anterior cingulate cortex. *Nat Neurosci* 9:940–947. [CrossRef Medline](#)
- Knutson B, Westdorp A, Kaiser E, Hommer D (2000) fMRI visualization of brain activity during a monetary incentive delay task. *Neuroimage* 12:20–27. [CrossRef Medline](#)
- Kovach CK, Daw ND, Rudrauf D, Tranel D, O’Doherty JP, Adolphs R (2012) Anterior prefrontal cortex contributes to action selection through tracking of recent reward trends. *J Neurosci* 32:8434–8442. [CrossRef Medline](#)
- Kriegeskorte N, Goebel R, Bandettini P (2006) Information-based functional brain mapping. *Proc Natl Acad Sci U S A* 103:3863–3868. [CrossRef Medline](#)
- Mitchell TM, Hutchinson R, Niculescu RS, Pereira F, Wang X (2004) Learning to decode cognitive states from brain images. *Mach Learn* 57:145–175. [CrossRef](#)
- Momennejad I, Haynes JD (2012) Human anterior prefrontal cortex encodes the “what” and “when” of future intentions. *Neuroimage* 61:139–148. [CrossRef Medline](#)
- Momennejad I, Haynes JD (2013) Encoding of prospective tasks in the human prefrontal cortex under varying task loads. *J Neurosci* 33:17342–17349. [CrossRef Medline](#)
- Monsell S (2003) Task switching. *Trends Cogn Sci* 7:134–140. [CrossRef Medline](#)
- O’Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ (2003) Temporal difference models and reward-related learning in the human brain. *Neuron* 28:329–337. [CrossRef Medline](#)
- Peck CJ, Jangraw DC, Suzuki M, Efem R, Gottlieb J (2009) Reward modulates attention independently of action value in posterior parietal cortex. *J Neurosci* 29:11182–11191. [CrossRef Medline](#)
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238. [CrossRef Medline](#)
- Power JD, Petersen SE (2013) Control-related systems in the human brain. *Curr Opin Neurobiol* 23:223–228. [CrossRef Medline](#)
- Rao SC, Rainer G, Miller EK (1997) Integration of what and where in the primate prefrontal cortex. *Science* 276:821–824. [CrossRef Medline](#)
- Reverber C, G6rgen K, Haynes JD (2012a) Compositionality of rule representations in human prefrontal cortex. *Cereb Cortex* 22:1237–1246. [CrossRef Medline](#)
- Reverber C, G6rgen K, Haynes JD (2012b) Distributed representation of rule identity and rule order in human frontal cortex and striatum. *J Neurosci* 32:17420–17430. [CrossRef Medline](#)
- Ridderinkhof KR, van den Wildenberg WPM, Segalowitz SJ, Carter CS (2004) Neurocognitive mechanisms of cognitive control: the role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain Cogn* 56:129–140. [CrossRef Medline](#)
- Rushworth MFS, Walton ME, Kennerley SW, Bannerman DM (2004) Action sets and decision in the medial frontal cortex. *Trends Cogn Sci* 8:410–417. [CrossRef Medline](#)
- Sakai K (2008) Task set and prefrontal cortex. *Annu Rev Neurosci* 31:219–245. [CrossRef Medline](#)
- Shenhav A, Botvinick MM, Cohen JD (2013) The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron* 79:217–240. [CrossRef Medline](#)
- Shomstein S (2012) Cognitive functions of the posterior parietal cortex: top-down and bottom-up attentional control. *Front Integr Neurosci* 6:38. [CrossRef Medline](#)
- Sigala N, Kusunoki M, Nimmo-Smith I, Gaffan D, Duncan J (2008) Hierarchical coding for sequential task events in the monkey prefrontal cortex. *Proc Natl Acad Sci U S A* 105:11969–11974. [CrossRef Medline](#)
- Soon CS, Brass M, Heinze HJ, Haynes JD (2008) Unconscious determinants of free decisions in the human brain. *Nat Neurosci* 11:543–545. [CrossRef Medline](#)
- Stokes MG, Kusunoki M, Sigala N, Nili H, Gaffan D, Duncan J (2013) Dynamic coding for cognitive control in prefrontal cortex. *Neuron* 78:364–375. [CrossRef Medline](#)
- Sugrue LP, Corrado GS, Newsome WT (2004) Matching behavior and the representation of value in the parietal cortex. *Science* 304:1782–1787. [CrossRef Medline](#)

- Tanaka SC, Doya K, Okada G, Ueda K, Okamoto Y, Yamawaki S (2004) Prediction of immediate and future rewards differential recruits cortico-basal ganglia loops. *Nat Neurosci* 7:887–893. [CrossRef Medline](#)
- Todd MT, Nystrom LE, Cohen JD (2013) Confounds in multivariate pattern analysis: theory and rule representation case study. *Neuroimage* 77:157–165. [CrossRef Medline](#)
- Tong F, Pratte MS (2012) Decoding patterns of human brain activity. *Annu Rev Psychol* 63:483–509. [CrossRef Medline](#)
- Wisniewski D, Reverberi C, Tusche A, Haynes JD (2014) The neural representation of voluntary task-set selection in dynamic environments. *Cereb Cortex*. Advance online publication. Retrieved July 23, 2015. doi:10.1093/cercor/bhu155. [CrossRef Medline](#)
- Woolgar A, Hampshire A, Thompson R, Duncan J (2011) Adaptive coding of task-relevant information in human frontoparietal cortex. *J Neurosci* 31:14592–14599. [CrossRef Medline](#)
- Woolgar A, Golland P, Bode S (2014) Coping with confounds in multivoxel pattern analysis: what should we do about reaction time differences? A comment on Todd, Nystrom and Cohen 2013. *Neuroimage* 98:506–512. [CrossRef Medline](#)
- Woolgar A, Afshar S, Williams MA, Rich AN (2015) Flexible coding of task rules in frontoparietal cortex: an adaptive system for flexible cognitive control. *J Cogn Neurosci*. Advance online publication. Retrieved July 23, 2015. doi:10.1162/jocn_a_00827. [CrossRef Medline](#)
- Yeung N, Nystrom LE, Aronson JA, Cohen JD (2006) Between-task competition and cognitive control in task switching. *J Neurosci* 26:1429–1438. [CrossRef Medline](#)