Cerebral Cortex, 2015, 1–13

doi: 10.1093/cercor/bhu327
Original Article

ORIGINAL ARTICLE

Reward Motivation Enhances Task Coding in Frontoparietal Cortex

Joset A. Etzel[1], Michael W. Cole[1,2], Jeffrey M. Zacks[1], Kendrick N. Kay[1] and Todd S. Braver[1]

[1]Washington University in St Louis, St Louis, MO, USA and [2]Rutgers University, Newark, NJ, USA

Address correspondence to Joset A. Etzel, Washington University in St Louis, 1 Brookings Drive, St Louis, MO 63130, USA. Email: jetzel@artsci.wustl.edu

Abstract

Reward motivation often enhances task performance, but the neural mechanisms underlying such cognitive enhancement remain unclear. Here, we used a multivariate pattern analysis (MVPA) approach to test the hypothesis that motivation-related enhancement of cognitive control results from improved encoding and representation of task set information. Participants underwent two fMRI sessions of cued task switching, the first under baseline conditions, and the second with randomly intermixed reward incentive and no-incentive trials. Information about the upcoming task could be successfully decoded from cue-related activation patterns in a set of frontoparietal regions typically associated with task control. More critically, MVPA classifiers trained on the baseline session had significantly higher decoding accuracy on incentive than non–incentive trials, with decoding improvement mediating reward-related enhancement of behavioral performance. These results strongly support the hypothesis that reward motivation enhances cognitive control, by improving the discriminability of task-relevant information coded and maintained in frontoparietal brain regions.

Key words: cognitive control, fMRI, motivation, MVPA, prefrontal

Introduction

In the last 10 years, there has been a resurgence of interest in a basic and perennial question for psychology and neuroscience: How do affective and cognitive processes interact in the brain? One example of this interest has been the recent explosion of research investigating how motivation influences goal-directed cognitive processing. This research has spanned many subfields, ranging from social psychology and individual differences studies that have focused on high-level motivational factors and unconscious processing of motivational cues, to work in human and animal neuroscience demonstrating that signals of reward and punishment directly modulate activation and processing in cognitive networks (Leon and Shadlen 1999; Pochon et al. 2002; Small et al. 2005; Adcock et al. 2006; Watanabe and Sakagami 2007; Kouneiher et al. 2009; Braver et al. 2014). The research has been broad-ranging, with motivational signals being shown to influence a wide variety of cognitive processes and components of behavioral performance (Botvinick and Braver 2015). Despite this variety, a converging view is that these effects could share a common cause: motivational signals modulate the coding and activation of task goals and associated control processes (Pessoa 2009; Chiew and Braver 2011; Braver 2012). This view is consistent with long-standing theoretical accounts postulating that a primary function of motivational signals is to prioritize and select behavioral goals (Miller et al. 1960; Simon 1967; Kruglanski et al. 2002). Specifically, it is hypothesized that motivation provides a control signal indicating which of the potential goals that an individual *could* pursue *should* be pursued: the goals most subjectively valuable, given the current internal and environmental states (e.g., reward contingency structure; Botvinick and Braver 2015).

The extant experimental data are consistent with this account, in that influences of reward and punishment appear to be strongest and most consistent in the brain's frontoparietal network (Pessoa 2009)—a set of regions that includes the lateral and medial prefrontal (PFC, e.g., dorsolateral PFC and anterior

cingulate, ACC) and parietal cortices—a network whose primary function appears to be supporting top-down goal-directed behaviors (i.e., goal pursuit) (Corbetta and Shulman 2002; Cole and Schneider 2007; Vincent et al. 2008; Duncan 2010). More direct evidence has come from experimental paradigms pairing motivational signals with pretrial contextual or task cues (i.e., information that can be used to prepare for upcoming response selection). In such paradigms, motivational signals influence how contextual and task-relevant information is activated, encoded, and maintained. For example, in the AX-CPT task, trials in which contextual cues are paired with reward signals are associated with enhanced maintenance and utilization of the contextual information (as indicated via both behavioral and pupillometric indices) (Chiew and Braver 2013). Similar behavioral findings have been observed in working memory and task-switching studies (Nieuwenhuis and Monsell 2002; Heitz et al. 2008; Bijleveld et al. 2009; Shen and Chun 2011; Zedelius et al. 2011). Brain imaging studies have demonstrated that these incentive effects on behavioral performance are mediated by increased frontoparietal activation, most prominently in dorsolateral PFC, during encoding and maintenance periods (Gilbert and Fiez 2004; Taylor et al. 2004; Krawczyk et al. 2007; Locke and Braver 2008; Jimura et al. 2010). We interpret these findings as indicating that incentive motivational signals enhance the activation strength and maintenance of cognitive task goals, which may be represented within dorsolateral PFC (and related lateral frontoparietal regions). A related and relevant construct is that of the "task-set," which refers to an abstract representation of the configuration of attentional, mnemonic, and motoric rules needed for successful task performance (Sakai 2008). Motivational modulations may serve as a general source of improved cognitive control, given that task-set representation and goal maintenance are thought to be core components of cognitive control (Monsell and Driver 2000; Miller and Cohen 2001; Sakai 2008).

Nevertheless, even though the prior findings are consistent with the idea that motivational signals improve cognitive control, via enhanced coding and activation of task goals, the evidence so far is indirect. Elevated *activity* in a particular brain region does not necessarily indicate that relevant *information* is more effectively encoded or represented there (Riggall and Postle 2012). Thus, stronger evidence is required to demonstrate that motivation influences the representation of task-set information in the brain. Evidence of this type is not readily provided by standard neuroimaging or neurophysiological methods, which look at the average level of activity in a brain region, but instead can be sought through approaches that examine neural representational codes, for instance as reflected in voxel-level BOLD (blood oxygenation level dependent) activation patterns. Multivariate pattern analysis (MVPA) methods can serve this purpose, by enabling the "decoding" of brain states: quantifying the degree to which a particular *brain* state can identify or predict a given *task* state (Kriegeskorte et al. 2006; Norman et al. 2006; Pereira et al. 2009). Indeed, MVPA approaches have been used to demonstrate that task state information can be decoded from the brain activity patterns detected in fMRI signals (Bode and Haynes 2009; Esterman et al. 2009; Cole et al. 2011; Woolgar, Thompson et al. 2011; Reverberi et al. 2012; Manelis and Reder 2013; Zhang et al. 2013; Waskom et al. 2014). For example, in work by Bode and Haynes (2009), it was possible to successfully decode task-related signals in lateral PFC and parietal cortex during a task in which participants were cued with which one of two simple stimulus-response rules to apply to an upcoming target stimulus. Woolgar, Thompson et al. (2011) used a similar approach, finding evidence of task-rule coding throughout frontoparietal control network regions. Those studies involved a small set of simple stimulus-response rules; recently, we used a complex task-switching experiment (involving 64 tasks and 3 cue dimensions) to show that lateral PFC was particularly involved in the coding of the task decision-rule dimension (Cole et al. 2011). However, none of these studies have tested whether task information coding is influenced by experimental manipulations of cognitively relevant variables (but see Woolgar, Hampshire et al. (2011) for findings suggesting that extensive practice modulates task-coding in frontoparietal cortex, and Waskom et al. (2014) for evidence that classification accuracy increases across repeated trials of the same task). In the current work, we directly tested the hypothesis that transient changes in motivational state may influence task information coding. Specifically, we hypothesized that when the reward value of a task trial is high, participants may code task-related information in an enhanced manner (i.e., with greater fidelity). As implemented with MVPA methods, this hypothesis predicts that task-set information will be more accurately decoded under high (compared with low) motivation conditions.

To examine this question, we manipulated motivational incentives using a standard cued task-switching paradigm. In cued task-switching, participants randomly alternate between performance of two (or more) tasks, with an advance cue specifying the task to perform on the upcoming trial. We took a rigorous and conservative methodological approach to avoid capitalizing on chance patterns in the data. First, participants completed a baseline session involving no motivational manipulation. Second, they returned to the laboratory on a later day for an incentive session, during which some of the task cues were accompanied by reward cues that signaled the availability of monetary rewards for fast and accurate performance. The baseline session was used to train a classifier to discriminate, based on fMRI brain activity patterns, which of the two tasks was to be performed on each trial. The second (reward motivation) session was used purely for hypothesis testing: Is classification accuracy on reward incentive trials different than accuracy on the intermixed nonincentive trials? Further, to test whether the effects of reward motivation on task-set coding were functional in nature, we examined MVPA decoding accuracy in relation to the effects of reward incentives on behavioral performance. Our predictions were that: 1) task-set decoding would be possible in frontoparietal brain regions; 2) reward incentives would enhance task-set decoding, leading to improved classifier accuracy on reward trials; 3) reward-related improvements in MVPA classification accuracy would statistically explain the reward-related enhancement of behavioral performance; and 4) that these effects would be directly reflected in reward-related shifts toward more well-formed and discriminable task representations (measured as voxel-level BOLD activity patterns).

Materials and Methods

Participants

Twenty-four young adults were recruited from the Washington University community. Participants were screened to be right-handed, native English speakers, have normal or corrected-to-normal vision, without fMRI contraindications, history of neurological illnesses, or psychotropic medication use. The study protocol was approved by the Washington University in St Louis Institutional Review Board. Participants were compensated $10/h for the practice session and $25/h for the two fMRI scanning sessions, plus a $50 bonus. Four participants (3 female,

1 male) were excluded from the analyses, two because they did not complete both scanning sessions, and two because of equipment malfunction, leaving 20 participants for analysis (14 female, 6 male, mean age 25 years, range 19–37 years). Two of the included participants were missing data from a single scanner run due to operator error.

Task Description

Participants performed a cued task-switching paradigm in which trials randomly alternated between Face (male or female) and Word (two-syllable or not) tasks performed on composite face–word stimuli. The tasks and stimuli were adapted from Yeung et al. (2006). Figure 1 is a schematic of the relevant trial features; multiple cue and feedback stimuli were included, as described in the Supplementary Experimental Procedures and Figure 1. Participants performed the task during two scanning sessions, held at least two days apart. The first scanning session was the "baseline," and the second, the "incentive."

During the baseline session (the first fMRI scanning day), each trial began with a cue indicating whether the Face or Word task should be performed in the trial. Task cues (one verbal, one pictorial for each condition) randomly alternated with equal frequency. In most trials, the target, a word superimposed on a noisy black-and-white photograph of a human face (Supplementary Fig. 1*a*), appeared after an interval of either 1600 or 4100 ms. Eight trials in each run were "partial trials," in which the target did not appear. When the target did appear, participants were asked to press one of two response buttons to indicate which feature was present in the task-relevant dimension (response mappings were counterbalanced across participants). Following the response, visual feedback was presented, indicating whether the response was correct or not, followed by a red fixation cross indicating the end of the trial. The intertrial interval was randomly selected from the values of 2500, 5000, and 7500 ms.

During the incentive session (the second fMRI scanning day), the cue and feedback stimuli were altered to indicate the incentive status for each trial (Supplementary Fig. 1*b*). Reward incentive (Incentive) and no-incentive (NoIncentive) trials were randomly intermixed with equal frequency, and were indicated by reward cues (one symbol and one color for each incentive type) that appeared simultaneously with the task cue. The reward was obtained only when both the correct response was made and the response was made more quickly than the median of the subject's reaction time during the baseline session (first scanning day). Following the response, on Incentive trials the feedback indicated whether or not the reward was obtained (Supplementary Fig. 1*b*); on NoIncentive trials, the feedback indicated only whether the response was correct or not. Participants were told that each obtained reward would increase their monetary bonus; in fact, all participants received a $50 bonus at completion.

During both sessions, each fMRI scanning run contained 32 trials (8 of which were partial), divided into 2 blocks of 16 trials each, in a mixed blocked/event-related design (Visscher et al. 2003). The trial sequence was randomized, with the constraint that each block contain an equal number of Face and Word trials, with an equal probability of task-repeat (e.g., Word preceded by Word) and task-switch trials (e.g., Word preceded by Face). Similarly, in the incentive session, there were an equal number of Incentive and NoIncentive trials, and an equal probability of incentive-repeat (e.g., Incentive preceded by Incentive) and incentive-switch (e.g., Incentive preceded by NoIncentive). For purposes of analysis, task-switch and task-repeat trials were averaged together, not analyzed separately.

fMRI Scanning and Image Processing

Functional images were acquired on a Siemens Allegra 3T head-only scanner. BOLD images were acquired using asymmetric spin-echo echo-planar imaging, with TR = 2.5 s, TE = 25 ms, FA = 90°, 32 slices, and 4 × 4 × 4 mm voxels (the same sequence for both session days). E-Prime (Psychology Software Tools, Inc.) was used to present the stimuli (via a screen viewed with a mirror), control the stimulus timing, and record the participants' responses (via the E-Prime Fiber Optic Button Response System). Image acquisition was time-locked to the stimulus presentation at the initial fixation and target periods.

The session on each day consisted of ten functional imaging runs, each of which lasted 7.6 min (183 volumes). The first two and last two runs of each session contained only a single-task type (all trials either Face or Word), while the remaining six runs contained randomly intermixed Face and Word tasks Only the six task-switching runs (intermixed Face and Word trials) were analyzed here, since they are of primary theoretical interest (cue and delay period activity in single-task runs likely differs substantially from task-switching runs, since trial-level cue processing is not needed). Each imaging session was preprocessed in a completely independent and self-contained manner (preprocessing did not mix images across acquisition days). The volumes were realigned and spatially normalized to the MNI ICBM152 template via the EPI mask using SPM8 (Wellcome Trust Centre for Neuroimaging). The voxels were kept at the acquired size (4 × 4 × 4 mm) and the images were neither smoothed, nor slice-time corrected. The remaining procedures and analyses were performed in R (R Development Core Team) except as otherwise indicated.

To capture brain activity related to task-coding, maintenance, and preparation, rather than implementation and response execution, the analysis focused on the cue phase of each trial, rather than target and feedback (Fig. 1 and Supplementary Fig. 1). Specifically, the images from the second and third time points after the start of each trial were extracted for further processing, approximately corresponding to the cue phase (after accounting for the hemodynamic lag). While some information from the previous trial could be present in these images, the balanced task and incentive presentation order ensures that any such information "bleeding" would not positively bias classifier

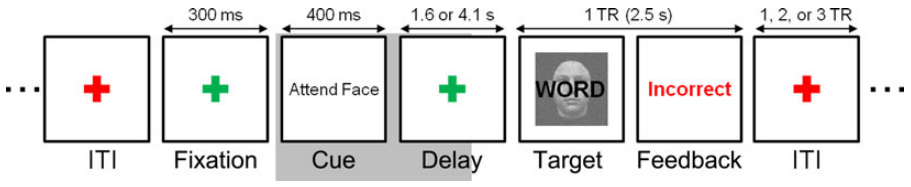

Figure 1. Schematic illustration of a baseline session trial in which the participant performed the Face task incorrectly. Details and examples of specific cue, feedback, and target stimuli are described in Supplementary Figure 1 and the Supplementary Experimental Procedures. Analysis focused on cue and delay period information, indicated here by shading. ITI, intertrial interval.

performance. Similarly, the use of these two time points does not eliminate the possibility that some target-related activity could also be captured. However, such activity should be minimal, due to the jittered delay period length, and because partial trials (trials in which the target did not appear) were also included.

The extracted data for each trial were temporally compressed (Mourao-Miranda et al. 2006) by averaging the images from the two time points together. Voxel-wise linear detrending was then performed within each of the two blocks making up each run. Finally, events of the same type were averaged, resulting in one example per person, run, task, and, for the incentive session, incentive level. Although trial averaging reduced the number of training and testing examples, it collapses across trial factors (such as cue modality) that were not of interest; moreover, preliminary analyses suggested that the averaging enhanced the robustness of the signal. One quarter of the trials in each run were partial: no target presented or response made (the cue portion of partial trials was identical to the cue portion of full trials). Both partial and incorrect trials were included in the analysis because the trial sequence was constructed to have both equal numbers of each task type, and equal probability that each task was followed by the same task, across all presented trials (i.e., including the partial trials). Since no performance or reaction time (RT) was collected for partial trials (as they lacked a target), RT-residualized regression such as advocated by Todd et al. (2013) could not be performed. Such a regression would not be appropriate in this case, regardless, since a task difficulty confound would not affect the key cross-session analyses: it is the *difference* in accuracy between incentive types within each person that is of interest, and participants would presumably find the same task (Word or Face) more difficult throughout the experiment (see also Woolgar et al. 2014).

Searchlight Analysis and Creation of the Candidate ROIs

Searchlight analysis (Kriegeskorte et al. 2006) was used to map local task information, since we did not want to make a priori assumptions regarding the functional anatomic boundaries and extent of regions involved with task-set representation. The searchlight analysis used two-voxel radius searchlights (8 mm, containing at most 93 voxels), fitting a classifier to each subject individually. Classification was of task (Face or Word), using linear support vector machines (SVMs), the e1071 R interface to libsvm, the parameter at $c = 1$, and partitioning on the runs (leave-one-run-out; 6-fold cross-validation), with accuracy computed by averaging across the folds. The group-level information map was constructed with voxel-wise *t*-tests, testing if each voxel's accuracy was significantly greater than chance (0.5) across the twenty subjects. The participant's maps were not smoothed prior to the *t*-test, nor warped, since the searchlight analysis was performed after spatially normalizing each participant's images.

The group-level searchlight information map was thresholded and clustered, creating a set of candidate ROIs. We used ROIs for the incentive effect analyses, since they tend to more robustly summarize the information present in a specified group of voxels, and are also less affected by variations in the distribution of information between individuals (Etzel et al. 2009, 2013). There is no standard technique for transforming an information map into a set of clusters, so we chose a simple procedure: the AFNI (Cox 1996) program Clusterize (within NeuroDebian 6.0.5), with a *t*-statistic threshold of 4.5, minimum cluster size of 10 voxels, and requiring cluster voxels to have at least one touching face. This threshold was somewhat conservative (corresponding to an image-wise family-wise error rate of 0.015; see the FWE section of the Supplementary Experimental Procedures), and was employed to separate the data into anatomically distinct focal clusters small enough to be suitable for MVPA. To ensure that our results were not restricted to this particular choice of threshold and group-level statistic, sensitivity analyses were performed, including creating the group information map by the proportion of subjects significant according to the binomial distribution (e.g., Pereira and Botvinick 2011). Although the binomial and *t*-based group-level statistics have quite different assumptions and properties, the results were compatible between the two methods (Supplementary Fig. 2), indicating the robustness of the primary findings to the clustering procedure.

Significance Testing

Both permutation and parametric statistical approaches were used to characterize the effect of incentive on cross-session classification accuracy. Two relabeling schemes were used for the permutation tests, one targeting the classification accuracy in Incentive and NoIncentive trials separately, and the other targeting the difference in classification accuracy between Incentive and NoIncentive trials; see the Permutation Testing section of the Supplementary Experimental Procedures. In all cases, the classification accuracy on all possible label rearrangements within each participant were computed, and the *P*-value calculated as the proportion of permuted-label datasets with a higher across-subjects statistic (mean accuracy or *t*-value) than the true across-subjects statistic. Relationships across the ROIs between classification accuracy and incentive were evaluated with a repeated-measures ANOVA (R command aov(accuracy ~ incentive.ID * ROI.ID + Error(sub.ID)/(incentive.ID * ROI.ID)); hierarchical linear models produced nearly identical results.

Results

The cued task-switching paradigm (Fig. 1 and Supplementary Fig. 1) presented targets consisting of faces with words superimposed. A cue presented at the beginning of each trial indicated whether the participant should respond to the gender of the face (male or female; Face task) or the number of syllables in the word (two-syllables or not; Word task). The paradigm was performed on two separate days: the baseline session, in which no incentives were available (or known about), and the incentive session, which included rewards. In this latter session, the task was identical but trials randomly varied between incentive (Incentive) and nonincentive (NoIncentive), with equal frequency for each.

Although trials also randomly varied in terms of task-switch versus task-repeat (relative to the prior trial), we collapsed across this variable for purposes of analysis. The primary reason for this choice was methodological (i.e., to maximize signal while maintaining trial randomization), but it also reflects that our theoretical interest was in the effects of task *cueing* rather than task *switching*. Indeed, considerable controversy remains as to the source of behavioral task-switch costs, and to theoretical interpretations regarding whether task-set preparation and updating is switch-specific, or instead occurs similarly on task-switch and task-repeat trials (Kiesel et al. 2010; Ruge et al. 2013).

We utilized MVPA methods to determine whether the task to be performed (Word or Face) could be decoded from fMRI activity patterns occurring during cue presentation and subsequent preparation for the upcoming target. Thus, we extracted the MR scans that corresponded to the cue and early delay period of each trial. Both the task (Word/Face) and incentive (Incentive/

No-Incentive) conditions were each specified by two randomly varying cues (8 cues total; see the Supplementary Experimental Procedures). Following prior MVPA studies (Vindiola and Wolmetz 2011; Woolgar, Thompson et al. 2011; Reverberi et al. 2012), we took this step to prevent the development of one-to-one associations between the perceptual features of the cue and the task or incentive condition, which could confound the interpretation of MVPA decoding effects.

Behavioral

The behavioral data indicated that the cued task-switching paradigm was challenging, but that participants ($N = 20$) could perform successfully: 90.6% accuracy (0.048 standard deviation [SD]) and 1123 ms mean RT (158 SD) was obtained in the baseline session. Although the Face and Word tasks differed in both accuracy (Face: 88.6%, Word: 92.6%, $t_{19} = 1.86$, $P = 0.08$) and RT (Face: 1083 ms, Word: 1163 ms, $t_{19} = 5.43$, $P < 0.001$), the pattern of performance (Word more accurate than Face, but Face faster than Word) suggests that the two tasks were not unequally difficult. Critically, the data from the incentive session suggested that reward motivation enhanced performance: When comparing incentive with nonincentive trials, participants were both faster (Incentive: 850, NoIncentive: 950; $t_{19} = 4.262$, $P < 0.001$) and more accurate (Incentive: 90.7%, NoIncentive: 87.3%, $t_{19} = 3.23$, $P = 0.002$). Further, neither the effect of incentive on RT nor performance interacted with task (F's < 1.2), suggesting that both Word and Face task performance were equally enhanced. Thus, it is unlikely that task difficulty confounded the interpretation of MVPA results related to task-set coding or incentive effects (Todd et al. 2013).

ROI Creation and Validation

We first identified candidate regions of interest (ROIs), groups of voxels possibly containing task information, by thresholding and then clustering the group-level information map resulting from a whole-brain searchlight analysis classifying task (Word or Face), using cue/delay-related fMRI BOLD activation images from the baseline session only. Fourteen candidate ROIs were identified (Supplementary Fig. 3 and Table 1), and then subjected to two validation tests. The first validation test was to determine whether each candidate ROI could classify task using only the *baseline session* images. Such a validation is important, given known issues regarding whether ROIs constructed from group-level searchlight analyses are actually informative themselves (Etzel et al. 2013). Twelve of the fourteen candidate ROIs did significantly ($P < 0.05$, uncorrected) classify task in the baseline session images, and so progressed to the second validation test, which was to determine whether each candidate ROI could classify task using only the *incentive session* images. Seven of the candidate ROIs showed significant task classification in this second validation step (Fig. 2, Table 1, and Supplementary Table 1), and so were considered "fully validated" and retained for our primary analyses.

The validation tests used images averaged from either the baseline (first validation test) or the incentive (second validation test) sessions only (not both), using leave-one-run-out cross-validation to compute classification accuracies for each participant

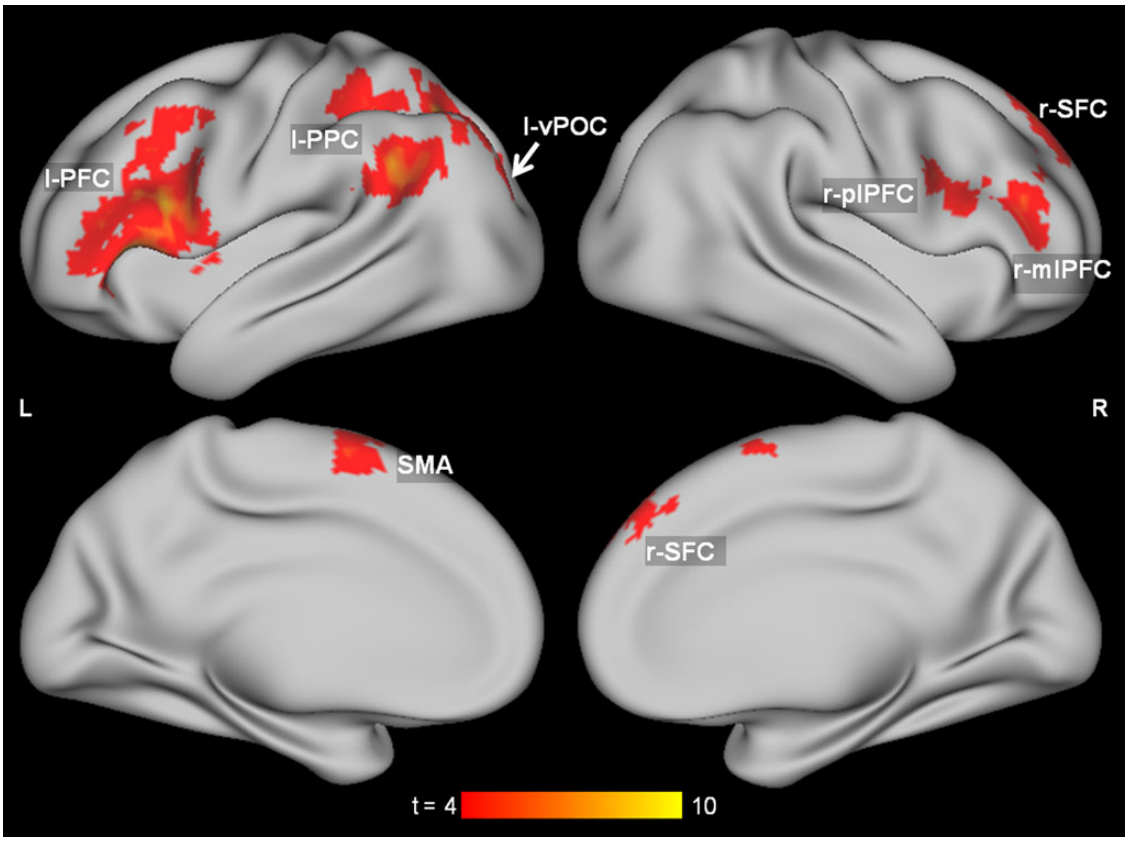

Figure 2. The seven validated ROIs, which together constitute the aggregate ROI. These ROIs significantly classify task on both the baseline and incentive sessions. Colors give *t*-values from the group-level searchlight analysis of task in the baseline session. Supplementary Figure 3 shows all candidate ROIs.

Table 1 Cross-session accuracy: classifiers trained on the baseline session and then tested on the incentive session, Incentive (Test on Inc) and NoIncentive (Test on NoI) trials separately. The location and size of each ROI, as well as the validation test results, are given in Supplementary Table 1.

ROI	Description	Cross-session accuracy		Difference	t-Value
		Test on Inc	Test on NoI	Inc – NoI	(diff >0)
	Aggregate ROI	0.67**	0.61**	0.061‘	2.16*
l-PFC	Left mid-lateral prefrontal cortex (BA 44/45/47)	0.67**	0.6**	0.067‘	1.36‘
SMA	Supplementary motor area (BA 6)	0.64*	0.56	0.075	1.75
l-vPOC	Left ventral parieto-occipital cortex (BA 19)	0.58**	0.5	0.083*	1.83*
r-SFC	Right superior frontal cortex (BA 9)	0.57*	0.52	0.045	0.996
r-plPFC	Right postero-lateral prefrontal cortex (BA 44)	0.57*	0.56*	0.005	0.11
r-mlPFC	Right mid-lateral prefrontal cortex (BA 45/46)	0.55‘	0.53	0.025	0.62
l-PPC	Left posterior parietal cortex (BA 7/40)	0.57*	0.58*	−0.01	−0.21

P-values are from permutation tests, uncorrected.
**$P < 0.01$.
*$P < 0.05$.
‘$P < 0.1$.

and ROI. These validation tests are distinct from the cross-session analyses used to test for reward incentive effects: first, because they consider data from each session individually; and second, because the incentive session images were prepared differently. Specifically, for the second validation test, the data were collapsed across incentive types (averaged by task only), in order to make the examples comparable with the baseline data in terms of signal; equal numbers of trials were averaged in each. Additionally, it is worth noting that the second validation test can be used to estimate the reliability and robustness of decoding, since classification was of training and testing data from the second acquisition day (incentive session), but with ROIs defined from the first acquisition day (baseline session). Thus, although the classification accuracy was lower in this second validation test compared with baseline (Supplementary Table 1), it likely reflects a more appropriate (i.e., generalizable) lower bound estimate of task decoding.

Incentive session task classification accuracy in the seven validated ROIs ranged from 62 to 72% (mean = 0.66, SD = 0.037, Supplementary Table 1), comparable with MVPA accuracies previously reported for task-decoding analyses (e.g., Bode and Haynes 2009; Reverberi et al. 2012). The ROIs were primarily located within the frontal and parietal cortices, including prominent bilateral dorsolateral PFC regions (Fig. 2). Many (but not all) of these regions match well to the locations of regions within the frontoparietal control network (Power et al. 2011; Yeo et al. 2011), and are comparable with the frontoparietal areas found in previous task-decoding and cognitive control network analyses (e.g., Waskom et al. 2014). Based on their demonstrated robustness for task-set decoding and anatomic plausibility, further analyses were conducted exclusively with these seven validated ROIs.

Testing for Reward Incentive Effects

The critical analysis was to test whether the validated ROIs exhibited reward incentive effects: Is the to-be-performed task more accurately classified on reward incentive or nonincentive trials? This was accomplished with *cross-session* analysis: a linear SVM classifier was trained once, on all baseline session data for each participant and ROI, and then tested twice, on the incentive session Incentive and NoIncentive trials separately. Using a single classifier allows direct comparison of accuracy (and other statistics) between Incentive and NoIncentive. This approach is somewhat similar to the functional localizer approach used in univariate analyses based on the general linear model (GLM; Saxe et al. 2006), in that we used independent datasets to create and test ROIs for theoretically relevant effects.

Task was more accurately classified on Incentive trials, with a range of 55–65% across the validated ROIs for Incentive, but only 50–60% for NoIncentive (Table 1 and Fig. 3); six of the seven ROIs had greater accuracy on Incentive trials. When each ROI was considered separately, permutation analysis indicated significant ($P < 0.05$) classification on Incentive trials in six of the seven ROIs, with the remaining ROI at $P < 0.1$. In contrast, only three ROIs had significant classification of NoIncentive trials. A separate, stringent permutation analysis evaluated the *difference* in accuracy between Incentive and NoIncentive in each ROI, finding significant effects in the left ventral parieto-occipital cortex ROI (l-vPOC, difference-based test, $P = 0.02$; *t*-based test, $P = 0.03$), and a trend-level effect in the left mid-lateral prefrontal cortex ROI (l-PFC, difference-based test, $P = 0.09$; *t*-based test, $P = 0.1$). Given this consistent observation of greater classification accuracy in Incentive than NoIncentive trials across the ROIs, we used a repeated-measure ANOVA to formally test whether the magnitude of the incentive effect interacted with ROI. A significant main effect of trial type (Incentive or NoIncentive) was observed in this analysis ($F = 4.6$, $P = 0.045$) and a marginal effect of ROI ($F = 1.5$, $P = 0.18$), but the trial type by ROI interaction was not even close to significant ($F = 0.67$, $P = 0.67$). These results suggest that the seven validated ROIs consistently showed a pattern of better task discrimination on Incentive than NoIncentive.

Given this common pattern across the ROIs, we conducted a follow-up analysis in which we treated the validated ROIs as a single ROI, which hereafter we call the “aggregate ROI.” Specifically, we trained a new classifier on all voxels from the validated ROIs, and then conducted the same analyses on it as on the ROIs individually. For this aggregate ROI, Incentive classification accuracy was 67% while NoIncentive accuracy was 61%, both significantly greater than chance (Table 1 and Fig. 3). The permutation analyses on the Incentive–NoIncentive difference were also statistically significant (difference-based test, $P = 0.07$; *t*-based test, $P = 0.01$): task decoding was more accurate on Incentive trials. Finally, we examined classifier performance separately for each task type, determining that the enhancement in classifier accuracy was present for both Word and Face tasks (Supplementary Fig. 4). Moreover, an ANOVA indicated that the incentive

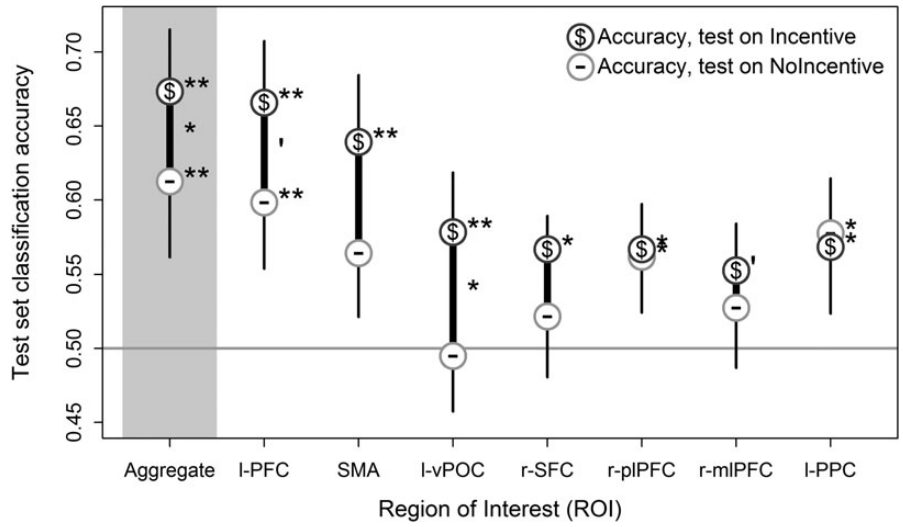

Figure 3. The cross-session task classification results for the validated ROIs and the aggregate ROI, testing on Incentive ($) or NoIncentive (-). Narrow lines give the standard error of the mean (SEM) for each accuracy. Significance is given by *P*-values from the permutation tests, $^{**}P < 0.01$ $^{*}P < 0.05$ $^{'}P < 0.1$, as in Table 1. Symbols beside each point indicate if the corresponding accuracy is greater than chance, while symbols between the points indicate if the difference (wide lines) is significant (*t*-based permutation test).

by task interaction was not significant ($F < 1$). Thus, the incentive decoding effect was not driven by task asymmetries or change in classifier bias, further arguing against potential confounds due to the classifier detecting an asymmetric change in task difficulty under reward incentive conditions (Todd et al. 2013).

An additional question is whether these effects were driven by mean activation differences between Incentive and NoIncentive or Word and Face conditions. The ROI creation procedures did not include an across-voxels normalization step, and so mean activity in the ROIs could have varied by incentive or task condition. However, a subsequent analysis revealed no differences in mean activity (Supplementary Fig. 5). Moreover, a control analysis which did include such normalization (across the aggregate ROI voxels, each classified example separately) actually enhanced rather than attenuated the finding of improved accuracy on Incentive trials: Incentive classification accuracy was 69% (*t*-test $P < 0.001$) while NoIncentive accuracy was 61% (*t*-test $P = 0.029$), with a *t*-test of the difference significant at $P = 0.036$. As an additional investigation into activation differences between Incentive and NoIncentive trials, we performed a univariate GLM-style analysis at each voxel (Supplementary Fig. 6), which revealed increased incentive-related activity within the expected regions of the reward/valuation network, including dorsal and ventral striatum, anterior insula, and ventromedial PFC (Pessoa and Engelmann 2010). Moreover, there was little overlap between the univariate (Incentive vs. NoIncentive) and MVPA-based (Word vs. Face) maps: only 15 of the 438 voxels in the validated ROIs also appeared in the univariate map, even at the liberal threshold of $P < 0.001$ uncorrected. This lack of overlap further suggests that univariate incentive-related activation differences are unlikely to have driven the current results.

Relationship with Behavioral Performance

We found both enhanced behavioral performance and MVPA classification accuracy on Incentive trials. Our theoretical hypothesis was that these two effects are linked: that the enhanced behavioral performance arises from more effective neural coding of task information when reward motivation is high. Further, we hypothesized that classification accuracy is a measure of task decoding based on neural activity patterns (as detected in BOLD changes at the voxel level), and that higher classification accuracy thus also reflects better encoding of task information. Putting these two hypotheses together, incentive-related performance enhancements should be mediated by more effective neural encoding, quantified here by the proxy measure of MVPA classification accuracy.

We tested for mediation using the stepwise regression-based approach recommended in Judd et al. (2001) for within-subject designs. The full set of tests is presented in the Mediation Analysis section of the Supplementary Experimental Procedures. The final test fits a regression equation to determine whether the difference in aggregate ROI MVPA classification accuracy (more accurate when tested on Incentive trials) predicts the incentive-related difference in behavioral performance (also more accurate on Incentive trials). This regression (Fig. 4) was significant: adjusted $R^2 = 0.236$, $F_{2,17} = 3.94$, $P = 0.039$. Examining the coefficients from the fit model, the classification accuracy difference estimate is significant ($t_{17} = 2.73$, $P = 0.014$), and so classification accuracy is a mediator of behavioral performance. The classification accuracy sum is not ($t_{17} = -1.53$, $P = 0.144$), however, and so classification accuracy is not a moderator of behavioral performance. Put simply, this mediation effect indicates that participants with a larger incentive-related increase in MVPA classification accuracy also had a larger incentive-related increase in behavioral performance accuracy. A similar analysis conducted on RT failed to reach significance; however, the pattern was consistent with that observed for behavioral performance (faster RTs associated with better MVPA classification; Supplementary Figure 7, Mediation Analysis section of the Supplementary Experimental Procedures).

The analysis did not exclude error trials, and so it is possible that both the increased MVPA classification accuracy and correlational effects may have been influenced by the incentive-related reduction in errors. However, this would not be problematic: since we claim only that classifier accuracy reflects task-coding effectiveness, if a failure of task encoding on some trials contributes to both a classifier and a behavioral error, it is consistent

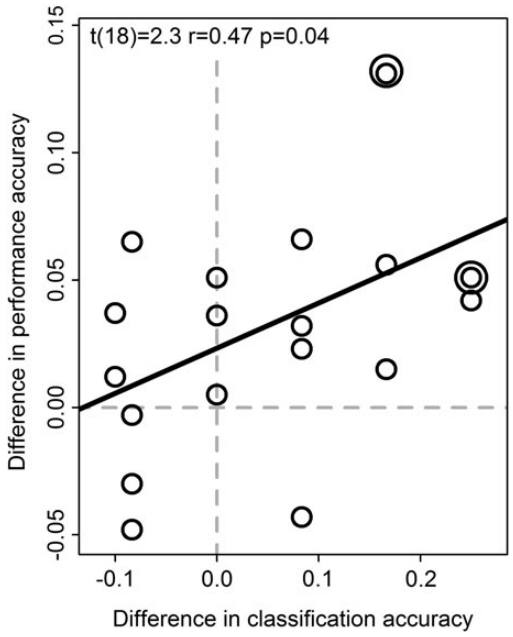

Figure 4. Participants with larger incentive-related increases in MVPA classification accuracy tended to also have a larger incentive-related increase in behavioral performance accuracy. Both differences in behavioral performance and classification accuracy were calculated as Incentive – NoIncentive. The scatterplot and regression line indicate the statistical significance of the zero-order correlation on these difference scores. Note that these are illustrated purely for descriptive purposes, and are distinct from the procedures and statistical values used to evaluate the mediation effect. Two pairs of subjects had similar differences, shown here as larger circles to avoid overplotting. See the Mediation Analysis section of the Supplementary Experimental Procedures for further details, and Supplementary Figure 7 for MVPA and behavioral performance correlations plotted separately for Incentive and NoIncentive conditions.

with our suggestion that such effects should be less likely to occur on Incentive trials.

Exploring Dataset Structure: Why Was Task Decoding More Accurate with Incentive?

We found that task decoding was significantly more accurate on reward incentive trials, compared with nonincentive trials. Necessarily, *decoding* partly depends on how well information was *encoded*, in this case as voxel-level BOLD patterns. Thus, we suggest that the incentive manipulation enhanced the effectiveness of task encoding and maintenance within a distributed set of frontoparietal regions (the aggregate ROI); cue-related activity patterns more clearly specified the upcoming task on reward incentive trials. As such, we predicted that task-set representations would be more distinct and less noisy on Incentive trials, leading to more effective biasing of on-going behavior. Two additional analyses were conducted to explore these predictions by probing the dataset structure in greater detail.

The first analysis examined Incentive and NoIncentive dataset structure by projecting the data into the one-dimensional space defined by the linear SVM decision boundary (typically referred to as the "hyperplane"). The distance of each test example to the SVM hyperplane can be interpreted as reflecting classifier "confidence": examples further from the hyperplane more truly belong to the class than examples near the hyperplane (Ramaswamy et al. 2001; Jaeger et al. 2005). Thus, if task representations are more distinct with incentive, then the Word and Face examples should be further apart in the Incentive than the NoIncentive dataset: further from the hyperplane, on opposite sides. Note that even when MVPA classification accuracy is higher in one condition than another, it does not guarantee greater distance to the hyperplane, since the classification decision is based solely on *which side* of the hyperplane each test case falls, not on its distance.

We calculated the distance of each cross-session test example to the participant's hyperplane. Since there is only one hyperplane per participant (using the aggregate ROI, trained on all baseline session examples), and equal numbers of Incentive, NoIncentive, Word, and Face test examples, we can create distributions of the distances in all participants combined (Fig. 5). These histograms show that the means of the distance distributions are further apart on Incentive (Face mean – Word mean = 2.21) than NoIncentive (Face mean – Word mean = 1.42) trials. Using Cohen's *d* to describe this effect size in standardized units, we found a value of 0.51 for Incentive and 0.30 for NoIncentive, a difference of 0.21. Permutation testing indicates this difference is significant ($P = 0.032$), such that the mean distance to the hyperplane was significantly greater (a larger effect size) on Incentive trials. In other words, the Word and Face task activity patterns were more distinct and more discriminable under incentive conditions, leading to higher confidence in the classifiers' decisions.

Complementing the exploration of spatial relationships between test examples using the hyperplane as a one-dimensional projection of the data, another approach is to probe these spatial relationships directly in the full high-dimensional data space (here, a 438-dimensional space described by the 438 voxels in the aggregate ROI). An appropriate statistic for quantifying the noise and structure present in high-dimensional datasets is the *likelihood of distance concentration*. When distances are strongly concentrated in high-dimensional space, all of the points become essentially equidistant; thus, attempts to meaningfully compare distances (or, relatedly, similarity) using standard metrics (e.g., Euclidean distance) will fail, and so will clustering (or pattern-similarity) methods that rely on nearest-neighbor calculations (François et al. 2007; Durrant and Kabán 2009; Kabán 2011). A dataset in which distances are more strongly concentrated has higher intrinsic dimensionality and less structure, and so is "noisier." Approaches for quantifying and compensating for distance concentration have been rapidly gaining attention within the machine learning community (Radovanović et al. 2010; Zimek et al. 2012; Winkler et al. 2013); however, we believe that this is the first use of distance concentration metrics to examine the structure of neuroimaging activation patterns (though see Schurger et al. (2010) for a different approach, but similarly motivated analysis). We tested the prediction that Incentive trials would be associated with a reduced likelihood of distance concentration (relative to NoIncentive trials), reflecting greater intrinsic structure in the activity patterns.

Distance concentration was estimated in the Incentive and NoIncentive datasets using a fixed set of epsilons, a parameter that indicates the relative contrast between the largest distance and the smallest distance of data points from a reference point (Kabán 2012); larger epsilons correspond to more liberal criteria for concluding that distances are concentrated. Estimating the lower bound on the probability of distance concentration across a range of epsilon values is thus akin to plotting a receiver operating curve, and a useful quantification metric is the area under the curve (AUC). We found that the distances are more likely to be concentrated on the NoIncentive examples, with AUC being 23.09 for NoIncentive and 21.62 for Incentive, a difference of 1.47

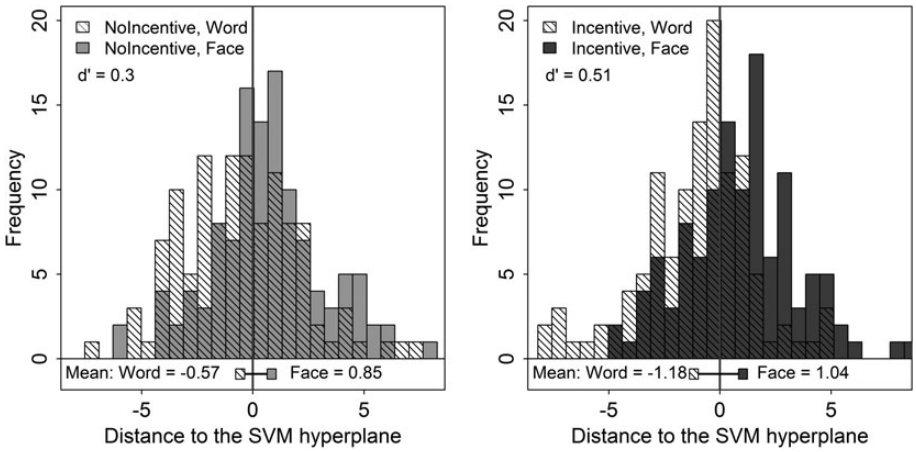

Figure 5. The Word and Face distributions are more distinct with incentive. Each histogram has the distance of each example to the SVM hyperplane for the aggregate ROI and all participants.

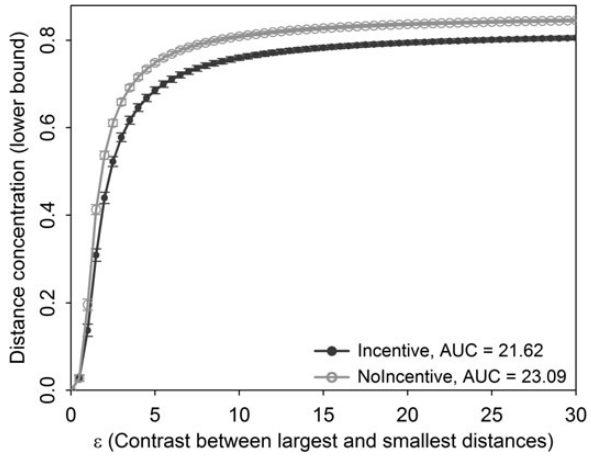

Figure 6. Distance concentration at various epsilons, following Figure 5 (Kabán 2012); larger epsilons correspond to more liberal criteria for concluding that distance is concentrated. Calculated with all participants' data combined, and the aggregate ROI.

(Fig. 6). Permutation testing indicates this difference is significant ($P = 0.007$), consistent with a lower probability of distance concentration in the Incentive dataset. In other words, this result confirms our prediction that under incentive the Word and Face activity patterns exhibited more intrinsic structure, and thus can be considered less noisy.

Discussion

This study examined whether the decoding of task-set information from distributed patterns of brain activity is modulated by trial-by-trial changes in reward motivation. Specifically, using MVPA methods, we tested the hypothesis that cues indicating the prospect of reward on the current trial would facilitate encoding and maintenance of the upcoming task for that trial. Replicating previous results, we found that: 1) the upcoming task can be decoded from voxel-level activity patterns, particularly within frontoparietal cognitive control regions (Bode and Haynes 2009; Cole et al. 2011; Woolgar, Hampshire et al. 2011; Reverberi et al. 2012; Zhang et al. 2013; Waskom et al. 2014); and 2) both brain activity and behavioral performance are modulated by reward motivation (Pessoa and Engelmann 2010). Most critically, reward motivation did not merely increase activity levels in the validated ROIs (primarily frontoparietal brain regions), but actually enhanced how task-relevant information was coded. This novel finding has implications for theories of cognitive control and motivation.

Task decoding was significantly more accurate on reward incentive trials compared with nonincentive trials. This is a highly selective effect: The two trial types were randomly intermixed, and the classifier was never exposed to incentive trials during training. Moreover, in the regions that showed task selectivity, there was no difference in mean activation between Word and Face conditions, or between Incentive and NoIncentive trials (Supplementary Figs 5 and 6). Thus, the task-decoding effects found here are likely different than the comparatively widespread activation differences related to task-switching or reward motivation detected by standard (mass-univariate) GLM-based analyses (the latter of which were observed in the current dataset within classic reward and valuation regions, such as ventral striatum, Supplementary Fig. 6). Necessarily, task decoding accuracy (in this case, by a classification algorithm) partly depends on how well the task information is encoded in the fMRI BOLD activity patterns. Thus, the incentive manipulation appears to have enhanced the effectiveness of task encoding and maintenance within the aggregate ROI. Restated, we propose that following task cue presentation, the activity patterns more clearly specified the upcoming task on reward incentive trials. That is, task-set representations were more distinct and less noisy, leading to more effective biasing of on-going behavior. Evidence consistent with this assertion was provided by follow-up analyses that investigated how representational coding was modulated on incentive trials. These indicated that on incentive trials there was: 1) increased discriminability of individual trial representations, reflected in greater distance of trial representations from the SVM hyperplane; and 2) more meaningful structure in the distribution of data points in high-dimensional space (lower probability of distance concentration), an effect that was independent of the classifier.

The relationship between classification accuracy and behavioral performance provides further support for the idea that enhancement of task representation may be an important mechanism by which reward motivation modulates on-going cognitive processing. Specifically, we found that the improved behavioral performance exhibited by participants on reward incentive trials was statistically mediated by the higher task decoding

(MVPA classification accuracy) occurring on these trials. This finding is consistent with decoding accuracy being a functionally relevant variable, reflecting the effectiveness of encoding, representation, and maintenance of task-set information. Task-set representation and maintenance serve as central components of cognitive control in many theoretical accounts. In particular, these representations are thought to be critical for biasing attention toward task-relevant processing pathways, and for resisting interference from task-irrelevant perceptual information (Monsell and Driver 2000; Miller and Cohen 2001; Kane and Engle 2002; Sakai 2008). Thus, we suggest that the activity patterns indicate the efficacy of task-set encoding, representation, and maintenance (during the delay period), which in turn contributes to successful performance during cued task-switching.

Why might task-set representation and maintenance be impacted by cues signaling reward incentives? Although the current data do not bear on this directly, we speculate that they are consistent with a neurocomputational link between task-set (or goal) updating and interactions between the midbrain dopamine (DA) system and lateral PFC. Specifically, phasic DA signals in PFC may serve an important neuromodulatory function that "gates" afferent inputs to lateral PFC signaling relevant task-set information, thus enabling successful updating and maintenance of this new information (Braver and Cohen 2000; Cohen et al. 2002; O'reilly 2006; Durstewitz and Seamans 2008; D'ardenne et al. 2012). This account is consistent with neurophysiological data demonstrating that DA modulation of PFC neuronal activity leads to a sharpening of PFC representations (by increasing signal-to-noise ratio) (Vijayraghavan et al. 2007; Thurley et al. 2008). Likewise, it is now well established that phasic DA firing occurs in response to reward-predictive cues (Schultz et al. 1997; Bayer and Glimcher 2005). Putting these findings together, we hypothesize that, on reward incentive trials, reward cues trigger phasic DA responses that modulate and potentiate task-set afferent signals to lateral PFC, thus facilitating the updating/encoding of these task-set representations within lateral PFC, and their subsequent active maintenance within the wider frontoparietal cognitive control network. Of course, further work will be needed to provide evidence of a three-way link between reward motivation, DA release in lateral PFC, and enhanced encoding and maintenance of task-set representations. Such work will most likely require a multimethod approach, such as integrating fMRI MVPA with pharmacological and/or radioligand PET methods.

The present results in some ways resemble effects observed in single neurons, in which selective attention enhances the neural population coding of attended features and dimensions, by improving the signal-to-noise ratio (Serences et al. 2009; Cohen and Maunsell 2010). This raises the question of whether the reward motivation conditions utilized in the current study can be re-construed primarily as attentional manipulations. Dissociating reward motivation effects from those of selective attention is notoriously difficult (Maunsell 2004), and since it was not the focus of the current study, we are reluctant to make strong claims here. Regardless of the degree of overlap or dissociation between the two constructs, it seems uncontroversial to suggest that reward motivation and selective attention might share common neural pathways and mechanisms to produce the types of behavioral facilitation effects observed here and also in the attentional literature (Pessoa and Engelmann 2010). One potential organizing framework to encapsulate such ideas is "attentional episode" account recently put forth by Duncan (2013), which suggests that the core function of the frontoparietal control network is to construct the formation of subgoals or subtasks, by linking the relevant perceptual inputs, behavioral actions, and desired outcomes. Our results suggest the possibility that reward motivation might facilitate this process by highlighting and enhancing the motivational value associated with the formation of a particular attentional episode representing the upcoming task. Additionally, this suggests a similarity to theoretical accounts of event representations, since these have also been postulated to involve the frontoparietal network and be updated through dopaminergic reward prediction mechanisms (Zacks et al. 2007, 2011). Thus, an important future direction might be to directly compare the similarities and distinctions between reward motivation and selective attention, in terms of their effects on task and event representations.

Finally, these results highlight the utility of MVPA methods for examining the influence of cognitively relevant experimental manipulations on the neural coding of task information. Both the current and prior results (Bode and Haynes 2009; Esterman et al. 2009; Cole et al. 2011; Woolgar, Thompson et al. 2011; Reverberi et al. 2012; Zhang et al. 2013) suggest that MVPA might be more appropriate than univariate GLM approaches for examining task-set representations, given that the latter depend on finding large contiguous voxel clusters that exhibit uniformly increased activity during Task A compared with Task B (or vice versa). Such uniform increases in task-related activity do not seem to be a prominent feature of task-set representation in frontoparietal regions. Instead, as we and others have found, task-coding in these regions appears to be primarily expressed in individual-specific and spatially intermixed activity patterns, in which weak task sensitivity is present at the voxel-wise level (Supplementary Fig. 8; Esterman et al. 2009). Such activity patterns are perhaps best detected with multivariate analyses, which can detect these as distinct profiles or "neural signatures" of task representation. Thus, MVPA methods might be the most effective approach to examine how task representations are impacted by relevant cognitive (and other) variables. For example, Woolgar, Hampshire et al. (2011) demonstrated that increased practice in learning about and switching between task rules (which presumably shifted the nature of task representation) led to reduced rule classification accuracy in frontoparietal regions. Recently, Manelis and Reder (2013) showed that classification of the upcoming working memory task condition (1, 2, or 3-back load in the N-back task) during the instruction period predicted individual differences in behavioral performance.

To our knowledge, ours is the first study to combine MVPA methods with both experimental manipulation and brain–behavior correlations to demonstrate that effects of a cognitive manipulation on behavioral performance might be mediated by a shift in task-coding properties. Additionally, the current findings suggest that reward incentives may be an especially effective method by which to enhance the fidelity of task-coding, providing important new insights regarding the neural mechanisms that underlie motivation–cognition interaction effects.

Supplementary Material

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

Funding

This work was supported by the National Institutes of Health (grant number R37MH066078) to T.B. and 1R01AG031150 to J.M.Z.

Notes

We thank Ata Kabán for discussions, code, and assistance with interpreting the distance concentration analyses. Computations were performed using the facilities of the Washington University Center for High-Performance Computing, which were partially provided through grant NCRR 1S10RR022984-01A1. *Conflict of Interest*: None declared.

References

Adcock RA, Thangavel A, Whitfield-Gabrieli S, Knutson B, Gabrieli JDE. 2006. Reward-motivated learning: mesolimbic activation precedes memory formation. Neuron. 50:507–517.

Bayer HM, Glimcher PW. 2005. Midbrain dopamine neurons encode a quantitative reward prediction error signal. Neuron. 47:129–141.

Bijleveld E, Custers R, Aarts H. 2009. The unconscious eye opener: pupil dilation reveals strategic recruitment of resources upon presentation of subliminal reward cues. Psychol Sci. 20:1313–1315.

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

Botvinick M, Braver T. 2015. Motivation and cognitive control: from behavior to neural mechanism. Annu Rev Psychol. 66:83–113.

Braver TS. 2012. The variable nature of cognitive control: a dual mechanisms framework. Trends Cogn Sci. 16:106–113.

Braver TS, Cohen JD. 2000. On the control of control: the role of dopamine in regulating prefrontal function and working memory. In: Monsell S, Driver J, editors. Attention and performance XVIII. Cambridge (MA): MIT Press. p. 713–737.

Braver TS, Krug MK, Chiew KS, Kool W, Westbrook JA, Clement NJ, Adcock RA, Barch DM, Botvinick MM, Carver CS, et al. 2014. Mechanisms of motivation-cognition interaction: challenges and opportunities. Cogn Affect Behav Neurosci. 14:443–472.

Chiew KS, Braver TS. 2011. Positive affect versus reward: emotional and motivational influences on cognitive control. Front Psychol. 2:1–10, Article 279.

Chiew KS, Braver TS. 2013. Temporal dynamics of motivation-cognitive control interactions revealed by high-resolution pupillometry. Front Psychol. 4:1–15, Article 15.

Cohen JD, Braver TS, Brown JW. 2002. Computational perspectives on dopamine function in prefrontal cortex. Curr Opin Neurobiol. 12:223–229.

Cohen MR, Maunsell JHR. 2010. A neuronal population measure of attention predicts behavioral performance on individual trials. J Neurosci. 30:15241–15253.

Cole MW, Etzel JA, Zacks JM, Schneider W, Braver TS. 2011. Rapid transfer of abstract rules to novel contexts in human lateral prefrontal cortex. Front Hum Neurosci. 5:1–13, Article 142.

Cole MW, Schneider W. 2007. The cognitive control network: integrated cortical regions with dissociable functions. Neuroimage. 37:343–360.

Corbetta M, Shulman GL. 2002. Control of goal-directed and stimulus-driven attention in the brain. Nat Rev Neurosci. 3:201–215.

Cox RW. 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Comput Biomed Res. 29:162–173.

D'Ardenne K, Eshel N, Luka J, Lenartowicz A, Nystrom LE, Cohen JD. 2012. Role of prefrontal cortex and the midbrain dopamine system in working memory updating. Proc Natl Acad Sci USA. 109:19900–19909.

Duncan J. 2010. The multiple-demand (MD) system of the primate brain: mental programs for intelligent behaviour. Trends Cogn Sci. 14:172–179.

Duncan J. 2013. The structure of cognition: attentional episodes in mind and brain. Neuron. 80:35–50.

Durrant RJ, Kabán A. 2009. When is "nearest neighbour" meaningful: a converse theorem and implications. J Complexity. 25:385–397.

Durstewitz D, Seamans JK. 2008. The dual-state theory of prefrontal cortex dopamine function with relevance to catechol-O-methyltransferase genotypes and schizophrenia. Biol Psychiatry. 64:739–749.

Esterman M, Chiu Y-C, Tamber-Rosenau BJ, Yantis S. 2009. Decoding cognitive control in human parietal cortex. Proc Natl Acad Sci USA. 106:17974–17979.

Etzel JA, Gazzola V, Keysers C. 2009. An introduction to anatomical ROI-based fMRI classification analysis. Brain Res. 1282:114–125.

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

François D, Wertz V, Verleysen M. 2007. The concentration of fractional distances. IEEE Trans Knowl Data Eng. 19:873–886.

Gilbert AM, Fiez JA. 2004. Integrating rewards and cognition in the frontal cortex. Cogn Affect Behav Neurosci. 4:540–552.

Heitz RP, Schrock JC, Payne TW, Engle RW. 2008. Effects of incentive on working memory capacity: behavioral and pupillometric data. Psychophysiology. 45:119–129.

Jaeger S, Huanfeng M, Doermann D. 2005. Identifying script on word-level with informational confidence. In: Document Analysis and Recognition, 2005. Proceedings. Eighth International Conference on. Vol. 411, p 416–420.

Jimura K, Locke HS, Braver TS. 2010. Prefrontal cortex mediation of cognitive enhancement in rewarding motivational contexts. Proc Natl Acad Sci USA. 107:8871–8876.

Judd CM, Kenny DA, McClelland GH. 2001. Estimating and testing mediation and moderation in within-subject designs. Psychol Methods. 6:115–134.

Kabán A. 2011. On the distance concentration awareness of certain data reduction techniques. Pattern Recognit. 44:265–277.

Kabán A. 2012. Non-parametric detection of meaningless distances in high dimensional data. Stat Comput. 22:375–385.

Kane MJ, Engle RW. 2002. The role of prefrontal cortex in working-memory capacity, executive attention, and general fluid intelligence: an individual-differences perspective. Psychon Bull Rev. 9:637–671.

Kiesel A, Steinhauser M, Wendt M, Falkenstein M, Jost K, Philipp AM, Koch I. 2010. Control and interference in task switching—a review. Psychol Bull. 136:849–874.

Kouneiher F, Charron S, Koechlin E. 2009. Motivation and cognitive control in the human prefrontal cortex. Nat Neurosci. 12:939–945.

Krawczyk DC, Gazzaley A, D'Esposito M. 2007. Reward modulation of prefrontal and visual association cortex during an incentive working memory task. Brain Res. 1141:168–177.

Kriegeskorte N, Goebel R, Bandettini P. 2006. Information-based functional brain mapping. PNAS. 103:3863–3868.

Kruglanski AW, Shah JY, Fishbach A, Friedman R, Chun WY, Sleeth-Keppler D. 2002. A theory of goal systems. In: Zanna MP, editor. Advances in experimental social psychology. San Diego (CA): Academic Press. p. 331–378.

Leon MI, Shadlen MN. 1999. Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. Neuron. 24:415–425.

Locke HS, Braver TS. 2008. Motivational influences on cognitive control: behavior, brain activation, and individual differences. Cogn Affect Behav Neurosci. 8:99–112.

Manelis A, Reder LM. 2013. He who is well prepared has half won the battle: an fMRI study of task preparation. Cereb Cortex. doi:10.1093/cercor/bht262.

Maunsell JHR. 2004. Neuronal representations of cognitive state: reward or attention? Trends Cogn Sci. 8:261–265.

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

Miller GA, Galanter E, Pribram KH. 1960. Images and plans. In: Plans and the structure of behavior. New York (NY): Henry Holt and Co. p. 5–19.

Monsell S, Driver J. 2000. Control of cognitive processes: attention and performance XVIII. Cambridge, MA: MIT Press.

Mourao-Miranda J, Reynaud E, McGlone F, Calvert G, Brammer M. 2006. The impact of temporal compression and space selection on SVM analysis of single-subject and multi-subject fMRI data. Neuroimage. 33:1055–1065.

Nieuwenhuis S, Monsell S. 2002. Residual costs in task switching: testing the failure-to-engage hypothesis. Psychon Bull Rev. 9:86–92.

Norman KA, Polyn SM, Detre GJ, Haxby JV. 2006. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. Trends Cogn Sci. 10:424–430.

O'Reilly RC. 2006. Biologically based computational models of high-level cognition. Science. 314:91–94.

Pereira F, Botvinick M. 2011. Information mapping with pattern classifiers: a comparative study. Neuroimage. 56:476–496.

Pereira F, Mitchell T, Botvinick M. 2009. Machine learning classifiers and fMRI: a tutorial overview. Neuroimage. 45:S199–S209.

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, Article 17.

Pochon JB, Levy R, Fossati P, Lehericy S, Poline JB, Pillon B, Le Bihan D, Dubois B. 2002. The neural system that bridges reward and cognition in humans: an fMRI study. Proc Natl Acad Sci USA. 99:5669–5674.

Power JD, Cohen Alexander L, Nelson Steven M, Wig Gagan S, Barnes Kelly A, Church Jessica A, Vogel Alecia C, Laumann Timothy O, Miezin Fran M, Schlaggar Bradley L, et al. 2011. Functional network organization of the human brain. Neuron. 72:665–678.

Radovanović M, Nanopoulos A, Ivanović M. 2010. Time-series classification in many intrinsic dimensions. In: Parthasarathy S, Liu B, Goethals B, Pei J, Kamath C, editors. Proceedings of the 2010 SIAM International Conference on Data Mining, p. 677–688, Philadelphia, PA.

Ramaswamy S, Tamayo P, Rifkin R, Mukherjee S, Yeang C-H, Angelo M, Ladd C, Reich M, Latulippe E, Mesirov JP, et al. 2001. Multiclass cancer diagnosis using tumor gene expression signatures. Proc Natl Acad Sci USA. 98:15149–15154.

Reverberi C, Görgen K, Haynes J-D. 2012. Compositionality of rule representations in human prefrontal cortex. Cereb Cortex. 22:1237–1246.

Riggall AC, Postle BR. 2012. The relationship between working memory storage and elevated activity as measured with functional magnetic resonance imaging. J Neurosci. 32:12990–12998.

Ruge H, Jamadar S, Zimmermann U, Karayanidis F. 2013. The many faces of preparatory control in task switching: reviewing a decade of fMRI research. Hum Brain Mapp. 34:12–35.

Sakai K. 2008. Task set and prefrontal cortex. Annu Rev Neurosci. 31:219–245.

Saxe R, Brett M, Kanwisher N. 2006. Divide and conquer: a defense of functional localizers. Neuroimage. 30: 1088–1096.

Schultz W, Dayan P, Montague PR. 1997. A neural substrate of prediction and reward. Science. 275:1593–1599.

Schurger A, Pereira F, Treisman A, Cohen JD. 2010. Reproducibility distinguishes conscious from nonconscious neural representations. Science. 327:97–99.

Serences JT, Saproo S, Scolari M, Ho T, Muftuler LT. 2009. Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. Neuroimage. 44:223–231.

Shen YJ, Chun MM. 2011. Increases in rewards promote flexible behavior. Atten Percept Psychophys. 73:938–952.

Simon HA. 1967. Motivational and emotional controls of cognition. Psychol Rev. 74:29–39.

Small DM, Gitelman D, Simmons K, Bloise SM, Parrish T, Mesulam MM. 2005. Monetary incentives enhance processing in brain regions mediating top-down control of attention. Cereb Cortex. 15:1855–1865.

Taylor SF, Welsh RC, Wager TD, Phan KL, Fitzgerald KD, Gehring WJ. 2004. A functional neuroimaging study of motivation and executive function. Neuroimage. 21: 1045–1054.

Thurley K, Senn W, Lüscher H-R. 2008. Dopamine increases the gain of the input-output response of rat prefrontal pyramidal neurons. J Neurophysiol. 99:2985–2997.

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

Vijayraghavan S, Wang M, Birnbaum SG, Williams GV, Arnsten AFT. 2007. Inverted-U dopamine D1 receptor actions on prefrontal neurons engaged in working memory. Nat Neurosci. 10:376–384.

Vincent JL, Kahn I, Snyder AZ, Raichle ME, Buckner RL. 2008. Evidence for a frontoparietal control system revealed by intrinsic functional connectivity. J Neurophysiol. 100: 3328–3342.

Vindiola M, Wolmetz M. 2011. Mental encoding and neural decoding of abstract cognitive categories: a commentary and simulation. Neuroimage. 54:2822–2827.

Visscher KM, Miezin FM, Kelly JE, Buckner RL, Donaldson DI, McAvoy MP, Bhalodia VM, Petersen SE. 2003. Mixed blocked/event-related designs separate transient and sustained activity in fMRI. Neuroimage. 19:1694–1708.

Waskom ML, Kumaran D, Gordon AM, Rissman J, Wagner AD. 2014. Frontoparietal representations of task context support the flexible control of goal-directed cognition. J Neurosci. 34:10743–10755.

Watanabe M, Sakagami M. 2007. Integration of cognitive and motivational context information in the primate prefrontal cortex. Cereb Cortex. 17:i101–i109.

Winkler R, Klawonn F, Kruse R. 2013. A new distance function for prototype based clustering algorithms in high dimensional spaces. In: Giudici P, Ingrassia S, Maurizio V, editors. Statistical models for data analysis. Springer International Publishing, p. 371–378.

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 & Cohen 2013. Neuroimage. 98:506–512.

Woolgar A, Hampshire A, Thompson R, Duncan J. 2011. Adaptive coding of task-relevant information in human frontoparietal cortex. J Neurosci. 31:14592–14599.

Woolgar A, Thompson R, Bor D, Duncan J. 2011. Multi-voxel coding of stimuli, rules, and responses in human frontoparietal cortex. Neuroimage. 56:744–752.

Yeo BTT, Krienen FM, Sepulcre J, Sabuncu MR, Lashkari D, Hollinshead M, Roffman JL, Smoller JW, Zöllei L, Polimeni JR, et al. 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. J Neurophysiol. 106:1125–1165.

Yeung N, Nystrom LE, Aronson JA, Cohen JD. 2006. Between-task competition and cognitive control in task switching. J Neurosci. 26:1429–1438.

Zacks JM, Kurby CA, Eisenberg ML, Haroutunian N. 2011. Prediction error associated with the perceptual segmentation of naturalistic events. J Cogn Neurosci. 23:4057–4066.

Zacks JM, Speer NK, Swallow KM, Braver TS, Reynolds JR. 2007. Event perception: a mind-brain perspective. Psychol Bull. 133:273–293.

Zedelius CM, Veling H, Aarts H. 2011. Boosting or choking—how conscious and unconscious reward processing modulate the active maintenance of goal-relevant information. Conscious Cogn. 20:355–362.

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.

Zimek A, Schubert E, Kriegel H-P. 2012. A survey on unsupervised outlier detection in high-dimensional numerical data. Stat Anal Data Min. 5:363–387.