NeuroImage 192 (2019) 88–100

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Probing cortical and sub-cortical contributions to instruction-based learning: Regional specialisation and global network dynamics

Adam Hampshire [a,*], Richard E. Daws [a], Ines Das Neves [a], Eyal Soreq [a], Stefano Sandrone [a], Ines R. Violante [a,b]

[a] Computational, Cognitive and Clinical Neuroscience Laboratory, Department of Medicine, Imperial College London, London W12 0NN, UK
[b] School of Psychology, Faculty of Health and Medical Sciences, University of Surrey, Guildford GU2 7XH, UK

A B S T R A C T

Diverse cortical networks and striatal brain regions are implicated in instruction-based learning (IBL); however, their distinct contributions remain unclear. We use a modified fMRI paradigm to test two hypotheses regarding the brain mechanisms that underlie IBL. One hypothesis proposes that anterior caudate and frontoparietal regions transiently co-activate when new rules are being bound in working memory. The other proposes that they mediate the application of the rules at different stages of the consolidation process. In accordance with the former hypothesis, we report strong activation peaks within and increased connectivity between anterior caudate and frontoparietal regions when rule-instruction slides are presented. However, similar effects occur throughout a broader set of cortical and sub-cortical regions, indicating a metabolically costly reconfiguration of the global brain state. The distinct functional roles of cingulo-opercular, frontoparietal and default-mode networks are apparent from their activation throughout, early and late in the practice phase respectively. Furthermore, there is tentative evidence of a peak in anterior caudate activity mid-way through the practice stage. These results demonstrate how performance of the same simple task involves a steadily shifting balance of brain systems as learning progresses. They also highlight the importance of distinguishing between regional specialisation and global dynamics when studying the network mechanisms that underlie cognition and learning.

1. Introduction

Humans are able to learn new tasks remarkably quickly. This ability relies on a class of cognitive processes that may be dubbed 'intentional learning' because they involve the temporary formation of working-memory programmes that guide task performance prior to consolidation (Hampshire et al., 2016). There are distinct types of intentional learning. The most widely studied is reinforcement learning (RL), where the participant formulates a provisional program for performing a task, applies it and when necessary updates it based on the outcome (Boettiger and D'Esposito, 2005; Brovelli et al., 2011; Collins and Frank, 2018; Eliassen et al., 2003; Ghahremani et al., 2010; Hampshire and Owen, 2006; Law et al., 2005; Toni et al., 2001). More recently, there has been increasing interest in instruction-based learning (IBL), where the program for performing a task is formed from explicit instruction (Brass et al., 2009; Cole et al., 2013; Hampshire et al., 2016; Hartstra et al., 2011; Passingham et al., 2013; Ruge and Wolfensteller, 2010; Sliwinska et al., 2017; Wolfensteller and Ruge, 2012). This interest is motivated by the fact that IBL provides a simple analogue of important real-world behaviours; humans are often required to perform new tasks based on observation or simple instructions. Furthermore, IBL is particularly well-suited for lab-based study because, in contrast to RL, the processes of forming and applying new task programmes can be studied free from other potentially confounding processes that relate to the exploration of alternative possibilities (Levine, 1975; Liu et al., 2015), e.g., identifying candidate rules, processing positive or negative feedback events and adjusting expectancies.

It is well established that both RL and IBL involve overlapping networks of cortical and sub-cortical regions in the human brain; however, these brain regions show varied engagement as a function of learning type and learning stage (Badre et al., 2010; Brovelli et al., 2011; Hampshire et al., 2008, 2016; Mohr et al., 2016; Toni et al., 2001). The unique functional role that each brain region has remains the topic of much debate and it is unclear how they coordinate as parts of a coherent mechanism in order to support the rapid learning of new tasks. Most relevant to the current study, in a previous article in this journal, we presented a multi-way functional dissociation between the activation profiles of three frontoparietal networks (FPNs), the default mode network (DMN) and the anterior caudate, during the learning of simple binary discrimination rules (Hampshire et al., 2016). In accordance with past studies, we reported that the FPNs were active when new discrimination rules were initially applied and that this activity diminished

* Corresponding author.
E-mail address: a.hampshire@imperial.ac.uk (A. Hampshire).

https://doi.org/10.1016/j.neuroimage.2019.03.002
Received 23 October 2018; Received in revised form 28 February 2019; Accepted 3 March 2019
Available online 6 March 2019

towards a resting-state level with practice. Conversely, the DMN was initially deactivated when new discrimination rules were being applied and then tracked towards resting-state level with practice. We focused on how these activation learning curves varied for the three FPNs and examined their causal interactions. Our primary conclusions were that the FPN activation profiles approximated the sub-systems of a hierarchically organised working-memory network (Badre et al., 2010), with the performance of new discrimination rules initially requiring multiple FPNs that then disengage in a sequential manner with practice (Hampshire et al., 2016). Conversely, the DMN activation profile was accounted for as a reduction in suppression from these FPNs when the mental concentration required to support the task diminished. These differential network dynamics were common to both reinforcement and IBL variants of the discrimination task.

Our interpretation of the differential roles of the FPNs and DMN was uncontroversial. However, the activation profile of the anterior caudate proved harder to reconcile with prominent contemporary theorising. Specifically, we only observed significant caudate activity at the point of instruction during the IBL task, i.e., not when the discrimination rules were subsequently applied. Strong activation was also evident at the point of negative feedback during the RL task. Based on these results, we proposed that the observed anterior caudate activation related to the formation of new discrimination rules in working memory, which occurs at different points in time for the IBL and RL tasks (Hampshire et al., 2016). This interpretation accords with previous reports of anterior caudate activity during switching tasks when new searches are initiated and when feedback leads to the identification of the correct discrimination rule (Brovelli et al., 2011; Hampshire et al., 2012; Mattfeld and Stark, 2011).

In a subsequent commentary (Ruge and Wolfensteller, 2016), it was noted that activation in the anterior caudate has been reported in some studies to follow a steady increase across the initial set of discrimination trials after a new rule had been instructed (Ruge and Wolfensteller, 2010, 2013). Based on this pattern of results, the authors had proposed that this brain region was involved in the early stages of the consolidation process, perhaps mediating the implementation of the new rules prior to their being fully automatised (Ashby et al., 2010; Ruge and Wolfensteller, 2016; Seger and Spiering, 2011). Our task involved simple discriminations that can be learnt rapidly; therefore, they proposed that the practice-related anterior caudate activation that was observed in previous studies could have occurred more rapidly in our study. Notably, our previous study design focused on practice effects over the span of minutes not seconds and used a brief instruction timeframe. Given the slow coupling of neural activity to haemodynamic response, the anterior caudate activity that we had attributed to rule instruction could instead have related to the first few subsequent trials, i.e., when the new discrimination rules were first being applied.

Here, we report results from a modified variant of our IBL fMRI paradigm that was designed to test the predictions of these two hypotheses. Participants again learnt binary discrimination rules from simple instruction slides. However, we used a prolonged instruction period, which allowed both transient and sustained activity during instruction encoding to be differentiated from each other, and from the first few trials when the discrimination was applied. First, we used mass-univariate analyses to replicate previous findings regarding the differential sensitivities of FPNs and the DMN across the practice stages of the task. We then tested whether the anterior caudate was active when rules were being instructed vs. when they were being applied. We examined whether instruction-related activity was transient or occurred throughout the encoding stage. Next, we used timecourse analysis to determine at a finer grain whether there was evidence of heightened caudate activation during the early stages of rule application. Finally, we used generalised psychophysiological interaction models to examine the functional connectivity of the anterior caudate with the broader discrimination-related network in order to test the prediction that there should be increasing frontal-striatal coupling during early stages of rule application (Ruge and Wolfensteller, 2013, 2015).

2. Methods

2.1. Participants

20 healthy participants (10 female) aged 20–39 completed this study. All participants were right-handed, English speakers with normal or corrected to normal eyesight. Volunteers were excluded if they had a history of neurological or psychiatric illness, were taking psychoactive medications or did not meet MRI safety criteria. The Imperial College Research Ethics Committee approved this study. Participants gave informed consent prior to entering the scanner.

2.2. Behavioural paradigm

Participants learnt four binary discrimination rules over contiguous 2-min learning blocks. Each block started with a slide showing one of the discrimination rules (Fig. 1). Participants had 16 s to encode the rule (e.g. brown shapes = left button response and green shapes = right button response) followed by a 2 s fixation cross, subsequent to which the 'practice' phase began. During the practice phase, participants applied the rule to sort a sequence of coloured shapes by making left or right button presses. There were 4 compound stimuli per rule, constructed from 2 exemplars per dimension (i.e. shape or colour). There was no feedback post response. Stimuli were presented in pseudo-randomised order at a rate of 1 per 2.1 s with 1/3 of trials showing fixation as opposed to a stimulus, which allowed activation during discriminations to be estimated relative to fixation. After 2 min a new rule slide was presented; this design ensured activations related to rule learning were not confounded by the total time spent in the scanner or on task. Rules always changed across dimensions; i.e. if one rule related to shape then the next related to colour, and all exemplars were replaced when the rules changed. This design ensured that the previously learned stimulus-response behaviour did not have to be overridden when a new rule was applied. 36 discriminations were presented for each of the four rule slides. Participants read a written description of the task and had the opportunity to ask questions before the start of the experiment. To ensure that the instructions had been fully understood, participants undertook a demo version of the task with just one rule prior to entering the MRI scanner.

2.3. Data acquisition and preprocessing

Tasks were programmed in MATLAB using Psychophysics Toolbox extensions (Brainard, 1997). Stimuli were projected on a screen, visible via a mirror, at the end of the scanner bore. All participants were scanned using a 3-T Siemens Magnetom Verio MRI scanner with syngo MR B17 software (Siemens, Erlangen, Germany). A T2-weighted echo planar image depicting blood oxygenation level dependent (BOLD) contrast was acquired every 2 s. Images consisted of 32*3 mm slices, with a 64X64 matrix, 192 × 192 mm field of view, 30 ms TE, 2 s TR, 80° flip angle, 0.61 ms echo spacing, and 2232 Hz/Px bandwidth. A 1 mm resolution MPRAGE structural scan was also collected for each individual for preprocessing. For each participant, functional and anatomical raw DICOMS were first converted to NIFTI format using MRtrix 3.0 "mrconvert" function. They were preprocessed using our standard pipeline, which depends on SPM12 functions (Statistical Parametric Mapping Welcome Department of Imaging Neuroscience), and MATLAB 2016b. Specifically, functional images were slice-time and motion corrected and co-registered to the participant's structural scan. Then structural scans from all participants were used to generate a custom DARTEL study template. Using the non-linear study template, functional scans were warped onto the standard Montreal Neurological Institute coordinate system, resampled to isotropic 2 mm cubed voxel size, and spatially smoothed with an 8 mm3 full width at half maximum Gaussian kernel. The data were high-pass filtered to remove low-frequency drifts. Responses were recorded using a pair of

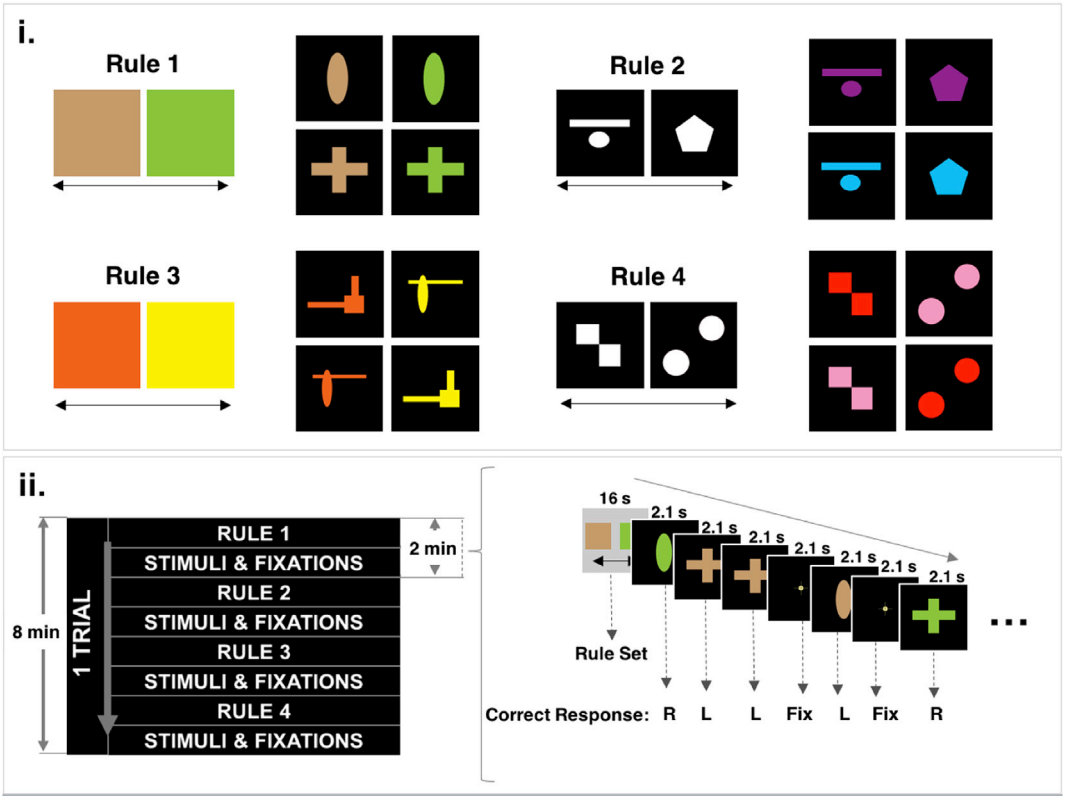

Fig. 1. Outline of the instructed learning task.
(i) The four binary discrimination rule slides and stimulus sets. (ii) Task structure, comprising performance of four discrimination rules over 2 min each including instruction (16s) and practice at a rate of 1 discrimination every 2.1s.

MRI-compatible response grips (ResponseGrip, NordicNeuroLab AS, Bergen, Norway); one held in each hand. Left or right responses were made by pressing the response grip button with the index finger of the corresponding hand.

2.4. FMRI analyses

2.4.1. Single subject models

FMRI data were modelled at the individual participant level in SPM12 using general linear models (GLMs). Discrimination trials were modelled using four predictor functions, each consisting of event timings convolved with the canonical haemodynamic response function. These included the onsets and durations of all discrimination trials broken down into 4 x 'practice stages' comprising ~9 discriminations arranged contiguously to estimate neural activation at a coarse grain as rules transitioned from novel to familiar. For example, the first predictor included the first 9 discriminations after definition of each of the four rules. The second predictor captured the next 9 discriminations, and so forth. Rule encoding was captured by two further predictors, one capturing transient activation at the onset of the rule slides, and the other capturing sustained activation throughout the entire 16 s rule encoding phase. Six additional predictors were included to capture noise due to head movements. These were the translations and rotations in the x, y and z planes.

2.4.2. Group level analyses of regional brain activity

Whole brain maps depicting parameter estimates for the experimental predictor functions were exported for group level random effects analyses. Analyses of nodes within the discrimination-related network were conducted using 5 mm radius spherical regions of interest (ROIs) with the MarsBaR toolbox (Brett et al., 2002), which calculates the average value from all voxels within the ROI. ROIs (Fig. 2) were predefined by placing them at peak coordinates from two previous studies (Hampshire et al., 2016; Ruge and Wolfensteller, 2010). Supplementary voxel-wise group level analyses were carried out in SPM12 and, unless reported otherwise, used cluster correction with initial voxelwise thresholding at $p < 0.01$ uncorrected followed by family wise error (FWE) cluster correction for the whole brain mass at $p < 0.05$. Finite impulse response functions (FIR) were calculated for ROIs by extracting mean activation values for each image, i.e., one value every 2 s, yoked to the onset of the rule instruction slide, and averaged across the four rules. Supplemental FIR analyses were also conducted on a previously published data-driven functional parcellation of the entire striatum (Tziortzi et al., 2014). This included seven ROIs that were defined by their differential functional connectivity to limbic, executive, caudal motor, rostral motor, parietal, occipital and temporal areas of the cortex.

2.4.3. Analysis of functional connectivity

Anterior caudate connectivity was analysed using generalised psychophysiological interaction models. Specifically, to examine how different stages of the instructed based learning task affected functional connectivity between ROI pairs, a generalised variant of the psychophysiological interactions (PPI) analysis was conducted (Friston et al., 1997; McLaren et al., 2012) using the following model:

$$Y^T = \beta_0 + [Y^s, H(X), E]\ \beta_G + [Y^s * H(X)]\ \beta j + e$$

where Y^T and Y^S are the target and source time series, respectively; H is the HRF, which is convolved with X, the matrix of experimental onsets; E corresponds to the nuisance regressors; β_j and β_G are weights of interest and of no interest respectively; β_0 is the intercept and e is the estimate residual. PPI parameter estimates were exported for each individual and examined at the group level as detailed in the results section.

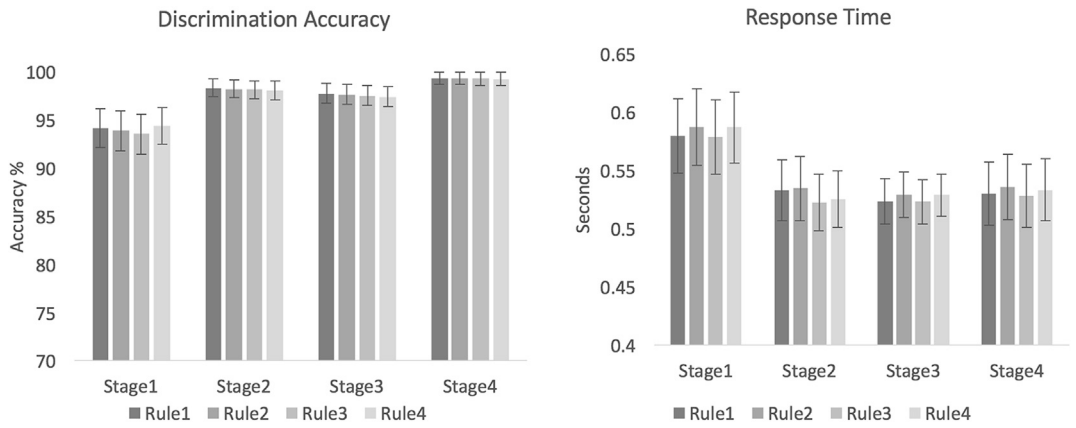

Fig. 2. Effects of rule and practice on discrimination performance.
Mean accuracy (left) and response times (right) for the rules and practice stages of the task. Error bars report the standard error of the mean.

3. Results

3.1. Behavioural results

We first examined the effects of practice on discrimination performance. At the single participant level, percentage discrimination accuracies (Fig. 2a) were calculated separately for each rule (Rule 1–4) within four practice stages, which comprised ~9 responses each (Stage 1–4). These data were examined in a 4*4 (Rule * Stage) repeated measures analysis of variance (ANOVA). As intended, overall discrimination accuracy was close to ceiling (>97%). There was a sub-threshold effect of Stage ($F(3,57) = 2.395$ $p = 0.096$) no main effect of Rule ($F(3,57) = 2.138$ $p = 0.118$) and a sub-threshold interaction ($F(9,171) = 2.259$ $p = 0.067$), (all Greehouse-Geisser corrected for non-sphericity).

Mean discrimination response times were calculated and examined in a 4*4 Rule * Stage repeated measures ANOVA of the same structure (Fig. 2b). There were robust main effects of Rule ($F(3,57) = 6.762$ $p = 0.003$) and Stage ($F(3,57) = 4.945$, $p = 0.012$) with no significant interaction ($F(9,171) = 0.4$,$p = 0.842$). Further examination with paired t-tests showed that response times for Stage 1 were slower than for Stages 2–4 (all $p < 0.01$) with no significant differences between Stages 2–4. Colour discriminations showed a small but significant response time cost relative to shape discriminations ($t = 3.288$, $p = 0.004$). Thus, the task produced the expected effects of practice on response time whilst ensuring a high overall level of accuracy.

3.2. Brain activation at different stages of the learning curve: ROI analyses

We confirmed the reproducibility of practice-related dissociation between the FPNs and the DMN from our previous study. Specifically, 28 5 mm radius regions of interest (ROIs) were defined at the peak activation coordinates for three contrasts of interest in the previous study (Fig. 4 & Table 1). They included:

(1) Brain regions that had reducing discrimination-related activity with practice (*Practice-*), including dorsolateral prefrontal cortex, lateral frontopolar cortex, the pre-supplementary motor area, parietal cortex and areas in the lateral occipital cortex bilaterally.
(2) Brain regions that had sustained responsiveness through all practice stages (*Sustained*), including within occipital cortex, inferior parietal cortex, anterior insular and frontal operculum bilaterally, the left motor cortex and thalamus.
(3) Brain regions that had increasing activation with practice (*Practice +*), including the medial orbitofrontal cortex, medial occipital cortex extending towards the precuneus, and temporal lobes bilaterally.

Table 1
Predefined regions of interest.

Contrast	ROI	MNI coordinates		
		X	Y	Z
1 Practice-	Left DLPFC	−48	34	22
	Left LOFC	−38	48	−8
	Left posterior DLPFC	−44	6	32
	Right DLPFC	42	40	20
	Right LOFC	42	52	−12
	Right posterior DLPFC	44	8	38
	Pre SMA	0	30	42
	Left PC	−44	−44	46
	Left LOC	−44	−56	−10
	Right PC	34	−48	48
	Right LOC	36	−46	−16
2 Sustained	Left occipital	−16	−82	−8
	Right occipital	−16	−96	14
	Left anterior insular	−38	14	8
	Right anterior insular	48	12	4
	left inferior PC	−36	−28	44
	Right inferior PC	52	−26	44
	Left frontal operculum	−42	−2	10
	Right frontal operculum	42	0	12
	SMA	0	0	48
	Left thalamus	−14	−20	6
	Left motor cortex	−38	−20	56
	Right motor cortex	32	−4	54
3 Practice+	Left temporal cortex	−50	−18	−14
	Right temporal cortex	54	−8	−22
	Left occipital cortex	−10	−72	18
	Right occipital cortex	10	−76	20
	MOFC	−4	48	−10
4 Caudate	Left anterior caudate	−12	20	4
	Right anterior caudate	12	20	4

DLPFC: dorsolateral prefrontal cortex. LOFC: lateral orbitofrontal cortex. PC: parietal cortex. SMA: supplementary motor area. MOFC: medial orbitofrontal cortex.

(4) Additionally, two ROIs were placed within the anterior caudate, one at the right coordinates in the study of instruction-based learning reported by Ruge et al. (Ruge and Wolfensteller, 2010), the other at the same coordinates mirrored in the left anterior caudate.

Parameter estimates for the current study were averaged across all voxels within each of these ROIs separately for each of the four practice stages. The same three contrasts of interest were then examined at the group level. First, we examined the mean responses of each of the four ROI sets described above in a 4 * 4 repeated measures ANOVA with the conditions ROI Set and Practice Stage. As expected, there was no main effect of Practice Stage ($F(3,57) = 0.73$ $p = 0.115$), a significant main

effect of ROI set ($F(3,57) = 22.585$ $p < 0.001$) and a robust interaction of ROI Set * Practice Stage ($F(9,171) = 9.357$, $p < 0.001$). In accordance with the previous study, ROI Set 1, which includes dorsal frontoparietal areas that are associated with working memory, was significantly active at stage 1 followed by a subsequent trend of decreasing activity (Fig. 3). ROI Set 2, which includes visual and motor areas, and the anterior insula and frontal operculum, showed sustained activity through all stages. ROI Set 3, which includes regions proximal to the DMN, showed significant deactivation at stage 1 followed by progressively diminishing deactivation with practice. ROI Set 4, which included the anterior caudate ROIs based on the work of (Ruge and Wolfensteller, 2010) was not significantly active at any of the practice stages, although a sub-threshold trend towards activity was evident for Stage 2 ($t = 1.88$ $p = 0.076$). For completeness, supplemental analyses of the individual ROIs conducted with a liberal uncorrected threshold further confirmed the expected differential sensitivity of brain regions to the practice stages of the task (Table 2).

A further possibility raised during review, was that the anterior striatum might be active at an intermediate as opposed to early stage of the practice phase. To test this possibility, repeated measures ANOVA was applied to the caudate ROI data across the four practice stages. There were significant main effects for both ROIs (left anterior caudate: $F = 2.91$, $p = 0.037$; right anterior caudate: $F = 3.08$, $p = 0.030$). Plotting mean parameter estimates across the practice stages showed that the numerically greatest activation level occurred within the second stage; however, a second increase was evident in the last stage (Fig. 5a).

3.3. Brain activation at different stages of the learning curve: voxelwise analyses

Focused ROI analyses were complemented by voxelwise analyses conducted unconstrained within the whole brain. All results are reported with family wise error (FWE) cluster correction at $p < 0.05$ for multiple comparisons after voxelwise thresholding the statistical maps at $p < 0.01$. The contrast of reducing activity across the practice stages (contrast = 3 1 -1 -3) showed significant clusters within the dorsolateral and lateral orbitofrontal cortex, parietal cortex and lateral occipital cortex bilaterally (Fig. 4a). The contrast of sustained activity across the practice stages (contrast = 1 1 1 1) was significant for clusters within the anterior insular, frontal operculum and temporal parietal junction bilaterally, the supplementary motor area, early visual areas, motor cortex and the thalamus (Fig. 4b). The contrast of increasing activity across the practice stages (contrast = −3 -1 1 3) showed significant clusters within the medial orbitofrontal, temporal and medial occipital cortices (Fig. 4c). The first two contrasts were in close concordance with our previous study. The third contrast produced an activation map that was largely non-overlapping with the previous study.

Notably, there were no significant voxels within the anterior caudate for any of these contrasts. We further tested whether there was involvement of the anterior caudate transiently in the first block, i.e., due to the simple nature of the IBL task (Ruge and Wolfensteller, 2016), or midway through the practice phase as suggested during review. To do this, we applied ANOVA across the four learning stages. We then repeated the analysis with data from a supplemental model that broke the practice phase down into 8 finer grained contiguous stages. In accordance with the above results, widespread patterns of activation were evident in both models. However, these also included a single cluster within the striatum volume ($p < 0.01$ FWE cluster corrected for the whole brain mass. Fig. 5b), which had a sub-peak in the left anterior caudate proximal to the ROI coordinates defined previously and extending medially. The cluster was evident in both analyses (4-stage model: $x = -14$, $y = 14$, $z = 4$, $F = 3.22$, voxelwise $p < 0.001$ uncorrected; 8-stage model: $x = -14$, $y = 18$, $z = 4$, $F = 3.46$, voxelwise $p < 0.001$ uncorrected). Plotting parameter estimates for that peak voxel showed greatest activity in the early-mid stage. However, the scale of the stage effect was much smaller than activity at the onset of the rule slide (Fig. 5c&d).

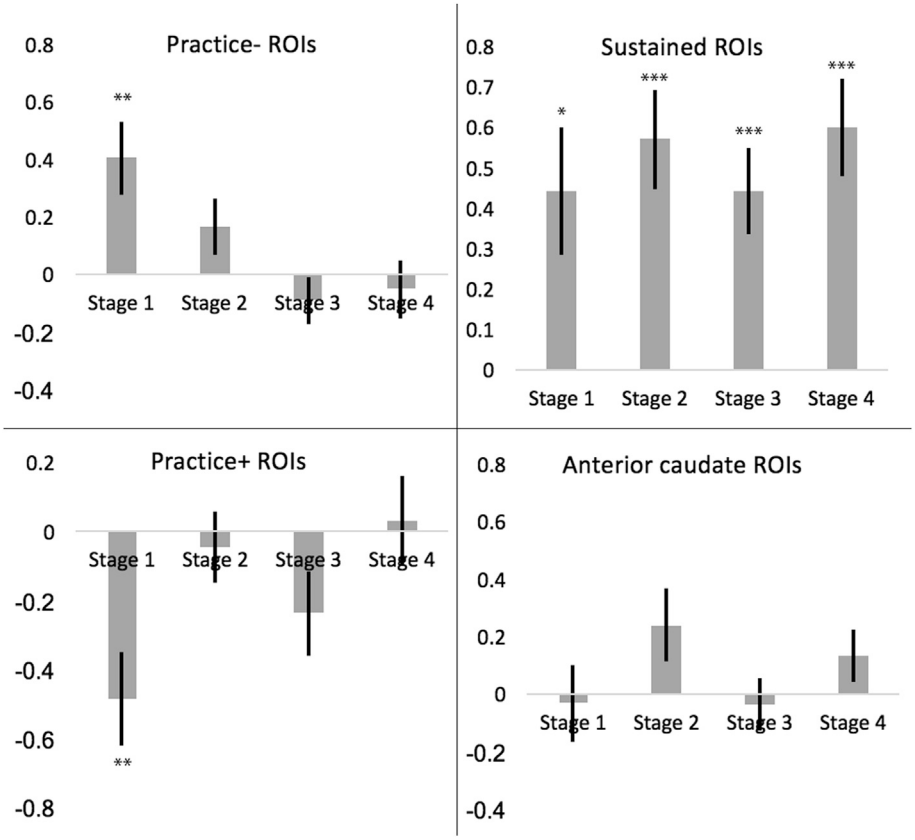

Fig. 3. Mean ROI activation across the four practice stages.
Analysis of discrimination related activity for ROIs defined in contrasts of Hampshire et al. (2016) and Ruge et al. (2010). ROIs were grouped and averaged by contrast then examined across the four practice stages. Top left: frontoparietal ROIs showed decreasing activation with practice. Top right: visual, motor insular and frontal operculum areas showed sustained activation at all practice stages. Bottom left: areas including the DMN showed progressively less deactivation with practice. Bottom right, there were no significant effects for the anterior caudate, although practice stage 2 showed a sub-threshold trend towards heighted activation ($p = 0.076$). Error bars report standard error of the mean. Significant results (two tailed) *$p < 0.05$, **$p < 0.005$, ***$p < 0.001$.

Table 2
Supplemental ROI results for contrasts of practice-related brain activation.

Contrast	ROI	Practice (3 1 -1 -3)		Practice (-3–1 1 3)		All stages (1 1 1 1)	
		t	p	t	p	t	p
Practice-	Left DLPFC	2.510	0.014	2.410	0.017	−0.730	0.467
	Left LOFC	3.100	0.002	2.850	0.005	−0.600	0.548
	Left posterior DLPFC	2.780	0.006	3.160	0.002	1.220	0.224
	Right DLPFC	3.610	<0.001	4.180	<0.001	0.510	0.614
	Right LOFC	1.350	0.181	1.230	0.220	−0.180	0.856
	Right posterior DLPFC	2.670	0.009	2.430	0.017	0.280	0.777
	Frontal midline	1.940	0.055	1.340	0.183	−0.830	0.408
	Left PC	3.820	<0.001	4.810	<0.001	3.250	0.002
	Left LOC	3.180	0.002	3.340	0.001	2.030	0.045
	Right PC	3.380	0.001	3.990	<0.001	2.890	0.005
	Right LOC	3.790	<0.001	3.870	<0.001	2.700	0.008
Sustained	Left occipital	−2.180	0.032	−3.130	0.002	7.580	<0.001
	Right occipital	−3.030	0.003	−3.620	<0.001	6.010	<0.001
	Left anterior insular	0.220	0.824	0.120	0.901	1.830	0.070
	Right anterior insular	−0.410	0.681	−0.380	0.701	2.360	0.020
	left inferior PC	0.010	0.994	−0.150	0.880	4.190	<0.001
	Right inferior PC	1.710	0.090	2.050	0.043	4.490	<0.001
	Left frontal operculum	−1.240	1.784	−0.790	0.431	6.200	<0.001
	Right frontal operculum	−0.240	1.190	−0.930	0.355	3.520	0.001
	Frontal midline	−0.670	0.507	−1.100	0.275	3.770	<0.001
	Left thalamus	0.250	0.800	0.310	0.757	2.850	0.005
	Left motor cortex	0.150	0.885	−0.130	0.898	4.190	<0.001
	Right motor cortex	−0.090	0.931	−0.830	0.408	3.080	0.003
Practice+	Left temporal cortex	−1.060	0.293	−2.670	0.009	−2.330	0.022
	Right temporal cortex	−1.230	0.219	−1.840	0.068	−1.420	0.159
	Left occipital	−3.130	0.002	−3.840	<0.001	−1.900	0.060
	Right occipital	−3.690	<0.001	−3.840	<0.001	1.370	0.174
	MOFC	−1.380	0.170	−2.800	0.006	−2.000	0.048
Striatum	Left anterior caudate	0.160	0.874	−0.900	0.370	0.860	0.391
	Right anterior caudate	−0.440	0.662	−1.710	0.090	0.570	0.571

All p values reported two tailed and uncorrected.

3.4. Brain activation during instruction: ROI analysis

Next, the responses to the instruction slide were examined for the four predefined ROI sets. A one-way repeated measures ANOVA examined the transient responses at the onset of the instruction slides. There was a significant main effect of ROI Set ($F(3,57) = 5.730$ $p = 0.002$) and a significant positive effect of condition ($F(1,19) = 49.609$ $p < 0.001$). T-tests showed robust transient activity of various strengths for all ROI sets (Set 1: $t = 6.478$ $p < 0.001$; Set 2: $t = 5.730$ $p < 0.001$; Set 3: $t = 8.222$ $p < 0.001$; Set 4: $t = 4.827$ $p < 0.001$). Next, a one-way ANOVA examined the sustained response throughout the 16 s period of time when the instruction slides were on the screen. There was a significant main effect of ROI Set ($F(3,57) = 18.279$ $p < 0.001$). The positive effect of condition was non-significant ($F(1,19) = 0.674$ $p = 0.422$). T-tests showed a threshold level effect of sustained activity within ROI Set 2 only (Set 1: $t = 1.612$ $p = 0.124$; Set 2: $t = 2.065$ $p = 0.053$; Set 3: $t = 0.159$ $p = 0.875$; Set 4: $t = -1.336$ $p = 0.197$). For completeness, supplementary analysis of individual ROIs at liberal uncorrected threshold (Table 3) showed significant transient activation in all but four ROIs when analysed individually. Sustained activation was evident primarily within occipital and parietal ROIs, and weakly within the left frontal cortex. The medial OFC and left temporal cortex ROIs showed sustained deactivation.

3.5. Brain activation during instruction: voxelwise analysis

Voxelwise analysis of transient activation relative to the implicit baseline rendered a widespread pattern of cortical and subcortical brain regions (Fig. 6a). In accordance with the ROI results, this included brain areas corresponding to the default mode network and the frontoparietal networks, which is in close agreement with the effects observed in previous learning studies (Chein and Schneider, 2005; Hampshire et al., 2016; Mohr et al., 2016). Critically, the peak striatum coordinates for this contrast were located within the anterior caudate bilaterally (left: $x = -16$, $y = 10$, $z = 2$, $t = 9.04$; right: $x = 16$, $y = 6$, $z = 6$, $t = 8.49$, both $p < 0.001$ FWE voxelwise corrected). There were no significant areas of deactivation at the cluster corrected threshold. By contrast, sustained activation was primarily confined to the visual and parietal cortices, with a single frontal lobe area surviving cluster correction located towards the posterior extent of the left middle frontal gyrus (Fig. 6b). Interestingly, areas consistent with the DMN, i.e., more akin to the *Practice* + contrast of the previous study, were deactivated in a sustained manner during this period of time at the cluster corrected threshold (Fig. 6c).

3.6. Analysis of ROI timecourses

The above results showed an unexpectedly strong and widespread pattern of transient brain activation at the point in time when the rule slide was initially displayed. To examine this effect in more detail, we plotted BOLD activation as a finite impulse response function (FIR) anchored to the point in time when the rule was first displayed and extending for the entire duration of the four practice stages. Examination of the ROI timecourses after simple de-trending showed pronounced activation peaks at all four points in time when the rule slides were initially displayed (Fig. 7 i&ii); ruling out the possibility that the strong effects were dependent on time in scanner or MRI-related artefacts. Averaging the timecourses across the four rules showed a pronounced transient peak in activation at the fifth time-bin (10th second) of the FIR function (Fig. 7 iii).

Examining the timecourses for the anterior caudate ROIs showed the same pattern of results, with a major peak in activation at around the 5th time bin (Fig. 7 iv). Given the temporally delayed nature of the haemodynamic response function and that the rule slide was displayed for 16 s this result is best accounted for by a transient spike in activity at the onset of rule definition, i.e., as opposed to sustained activation throughout rule presentation or through the first few trials when the rule is applied. A second minor peak occurred 10 s (5 time points) after the onset of the first discrimination trial and was followed by an upwards trend in

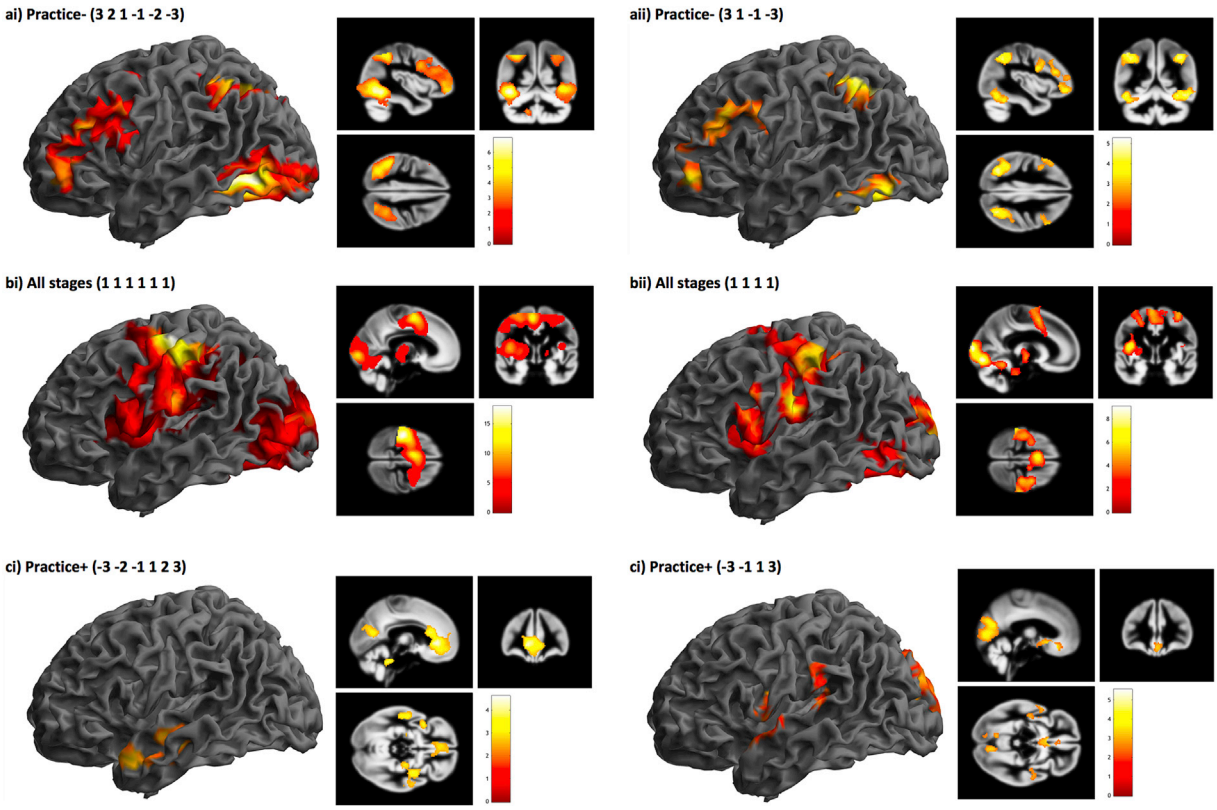

Fig. 4. Voxelwise analyses of practice effects for current and previous study. Discrimination-related activation. a
i & aii
: Reducing activity with practice. b
i & bii
: Sustained activation for all practice stages. c
i & cii: Increasing activity with practice. Left images from the previous study (Hampshire et al., 2016). Right images from the current study. All rendered with $p < 0.01$ voxelwise and $k = 100$ extent threshold.

activation over the subsequent 1 min, equivalent to two learning stages. In a post hoc analysis, a repeated measures ANOVA was run for the mean FIR time course for the anterior caudate ROIs focused on 40 s after the onset of the first discrimination. This showed a significant main effect of time ($F(19,361) = 2.462$ $p < 0.001$). This effect remained significant when analysing just the first ten time bins, capturing the secondary initial onset peak ($F(9,171) = 4.679$ $p < 0.001$) and when focusing on just the last ten time bins, capturing the upwards activation trend with practice ($F(9,171) = 3.031$ $p = 0.005$). For completeness, FIR functions were also calculated for seven functional subdivisions of the striatum (Fig. 7 v). A similar pattern of results was evident in all seven ROIs, with major and minor peaks at the onset of rule definition and rule implementation phases of the task, followed by a slow upward trend in activity throughout practice stages 1 and 2. Notably, this pattern, with primary and secondary peaks, followed by an upward trend was also evident for the mean timecourse as calculated across the whole network of ROIs.

3.7. Connectivity analysis

Finally, it has been proposed that increased activation of the anterior caudate with practice is paralleled by increased functional coupling with the lateral prefrontal cortex (Ruge and Wolfensteller, 2013, 2015), reflecting an increasingly prominent role in mediating performance of the task during the early consolidation phase. Consequently, we calculated generalised psychophysiological interaction models (PPIs) between these regions to capture connectivity (1) transiently at the onset of the instruction, (2) sustained through the instruction encoding phase, and (3) during each of the practice stages. Here, because practice effects on activation in the anterior caudate were primarily evident across stages 1 and 2 we report the PPI contrast of stage 1 minus stage 2 specifically.

We observed no significant changes across the practice stages in the coupling of either of the anterior caudate ROIs with the right or the left dorsolateral prefrontal cortex even at a liberal $p < 0.05$ uncorrected threshold (Fig. 8a). However, visual inspection showed a sub-threshold trend towards greater PPI strength in practice stage 1 than 2 for the right DLPFC. Examining the mean connectivity of the network with the anterior caudate ROIs (averaged across all connections) also showed no significant effect ($t = 0.38$ $p = 0.71$). Examining individual connections at the exploratory uncorrected threshold primarily showed practice-related decreases for the left anterior caudate across stages 1 and 2, including for connections with the right preSMA, right occipital cortex, right frontal operculum and left thalamus ROIs (Fig. 8b i).

In contrast, there were significant increases in connectivity between the right caudate and the left DLPFC at the onset of the rule slide with the left caudate showing a trend in the same direction (Fig. 8a). Similarly, there was increased connectivity between the left caudate and the right DLPFC during rule encoding with a similar trend evident for the right caudate. More globally, mean connectivity of the anterior caudate ROIs with the rest of the network was significantly increased at the onset of instruction ($t = 3.15$ $p = 0.005$) and throughout the instruction encoding phase ($t = 3.59$ $p = 0.002$). Examining the basis of this effect at the exploratory individual connection level showed increased connectivity for both the left and right caudate ROIs with multiple areas of the network at the onset of, and throughout, the instruction phase (Fig. 8b ii

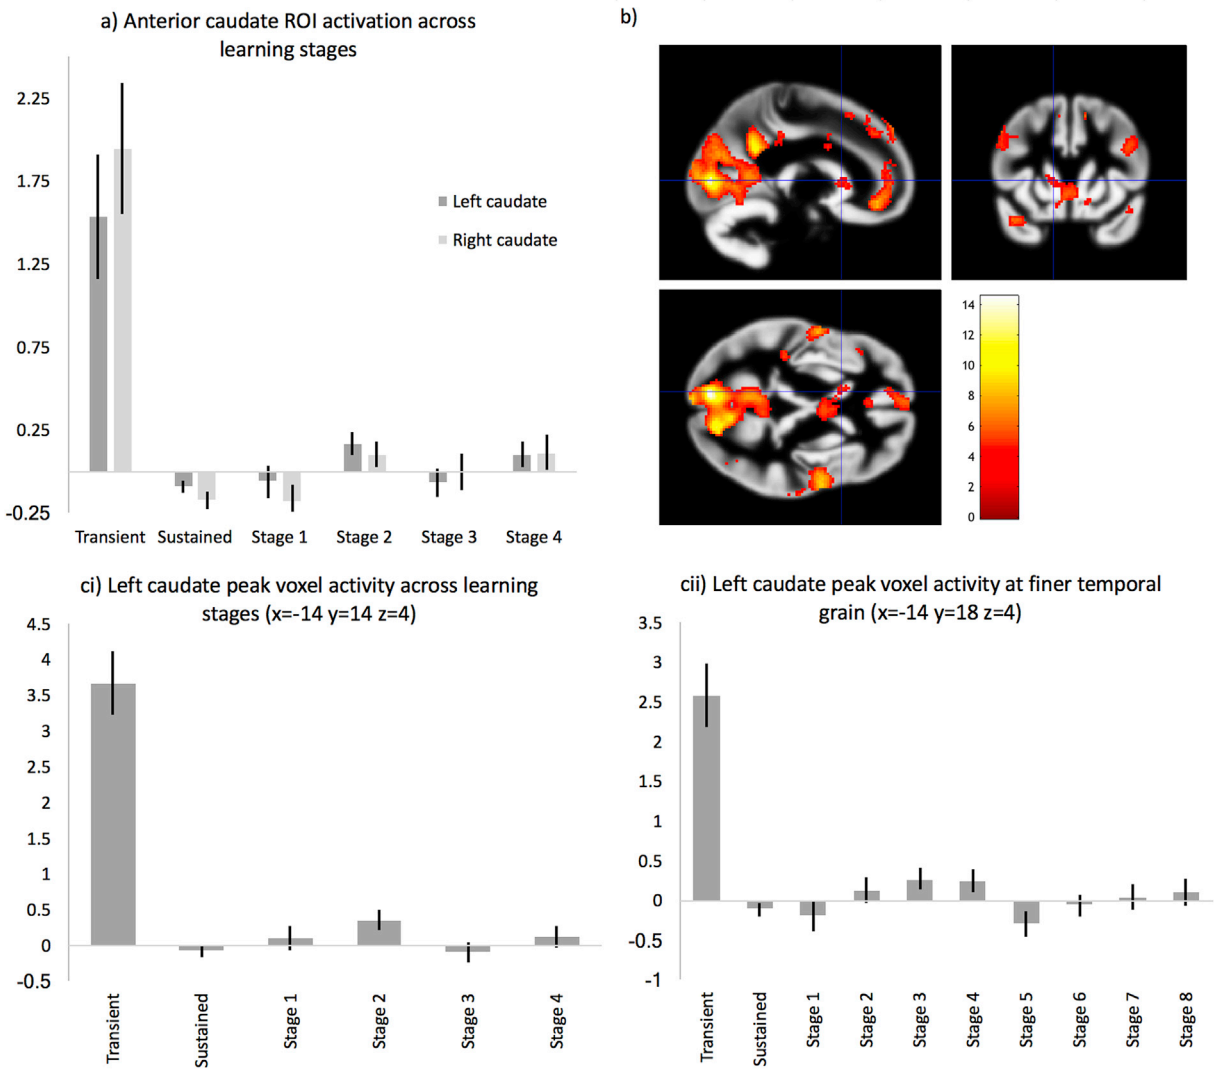

Fig. 5. ANOVA focused on anterior caudate activity across implementation stages.
a) Anterior caudate ROI activity at different practice stages. There was a significant main effect across the four stages (right hand bars) characterised by heightened activity at stage 2 and a further increase at stage 4. For comparison, transient activity and sustained activity are plotted for the same ROIs (left 2 bars). Error bars report the standard error of the mean. **b)** Voxelwise ANOVA across the four practice stages rendered widespread clusters including one with a sub-peak close to the expected coordinates in the left anterior caudate ($p < 0.05$ FWE cluster correction after initial voxelwise thresholding at $p < 0.01$). **c)** Plots of mean peak voxel parameter estimates across the encoding and practice stages of the task. **i)** There was a peak in activity during practice stage 2. This was dwarfed by the transient activity when the rule slides were initially presented. Note, analyses identifying the peak voxel were conducted across the practice stages only. Error bars report the standard error of the mean. **cii)** Conducting ANOVA with a model that broke the practice phase down into 8 finer-grained stages showed a smooth increase in activity towards the early-mid stage of the practice phase.

& iii). These results accord with heightened coupling of the caudate with the discrimination-related network when rules are being encoded as opposed to applied.

4. Discussion

The results presented here confirm those of our previous study, where we reported activation within the anterior caudate when new discrimination rules were presented. Using a prolonged instruction period, we extended these results by demonstrating that there was a large peak in the BOLD timecourse corresponding to the onset times of the rule instruction slides. The event-related model indicated that this peak was best explained by a transient spike in activity corresponding to the onset of the instruction slide, i.e., as opposed to sustained encoding activity throughout the duration of the slide presentation. Furthermore, the FIR model showed that the BOLD timecourse had entered the post-peak trough prior to presentation of the first discrimination trial.

These results accord with our hypothesis that the anterior caudate and frontoparietal networks co-activate when new rules are being updated in working memory. However, it would be erroneous to infer a unique role for the anterior caudate in rule encoding based on these results. For example, the transient activation contrasts with more sustained activity in other brain regions including occipital and lateral prefrontal cortex areas when the rule slide was displayed. This may accord better with a reorienting as opposed to encoding mechanism. Furthermore, taking a more holistic view, the activation spike observed at instruction onset was evident throughout much of the cortex. In fact, it was observed for regions of the brain corresponding to both the DMN and the FPNs concurrently. Although these networks are often anti-correlated with respect to their sensitivities to task demands (Gao and Lin, 2012; Raichle, 2015; Spreng et al., 2010), they have been observed to positively co-activate under certain conditions, most relevantly, during task preparation (Koshino et al., 2014). Furthermore, the instruction-related spike was evident for all seven of the functional striatum ROIs in the

Table 3
Transient and sustained ROI activation during instruction.

Contrast	ROI	Transient		Sustained	
		t	p	t	p
Practice-	Left DLPFC	5.280	<0.001	1.85	0.067
	Left LOFC	1.290	0.200	1.68	0.096
	Left posterior DLPFC	7.780	<0.001	2.04	0.044
	Right DLPFC	3.870	<0.001	1.41	0.161
	Right LOFC	1.400	0.165	0.39	0.700
	Right posterior DLPFC	6.300	<0.001	−0.33	0.741
	Frontal midline	4.830	<0.001	0.03	0.976
	Left PC	6.420	<0.001	3.49	0.001
	Left LOC	5.500	<0.001	2.78	0.006
	Right PC	7.340	<0.001	2.77	0.006
	Right LOC	7.810	<0.001	1.93	0.057
Sustained	Left occipital	6.970	<0.001	9.85	<0.001
	Right occipital	4.510	<0.001	9.29	<0.001
	Left anterior insular	3.360	0.001	−0.17	0.868
	Right anterior insular	2.880	0.005	−0.97	0.335
	left inferior PC	4.460	<0.001	0.58	0.564
	Right inferior PC	4.820	<0.001	0.11	0.911
	Left frontal operculum	1.140	0.259	−0.63	0.530
	Right frontal operculum	1.000	0.318	−0.9	0.371
	Frontal midline	4.090	<0.001	−0.59	0.554
	Left thalamus	5.530	<0.001	1.18	0.239
	Left motor cortex	4.100	<0.001	0.79	0.431
	Right motor cortex	6.920	<0.001	1.16	0.250
Practice+	Left temporal cortex	4.470	<0.001	−2.53	0.013
	Right temporal cortex	4.670	<0.001	−1.25	0.215
	Left occipital	8.340	<0.001	1.99	0.049
	Right occipital	8.150	<0.001	6.81	<0.001
	MOFC	2.770	0.007	−2.73	0.007
Striatum	Left anterior caudate	4.150	<0.001	−0.78	0.440
	Right anterior caudate	4.890	<0.001	−1.41	0.160

All p values reported two tailed and uncorrected.

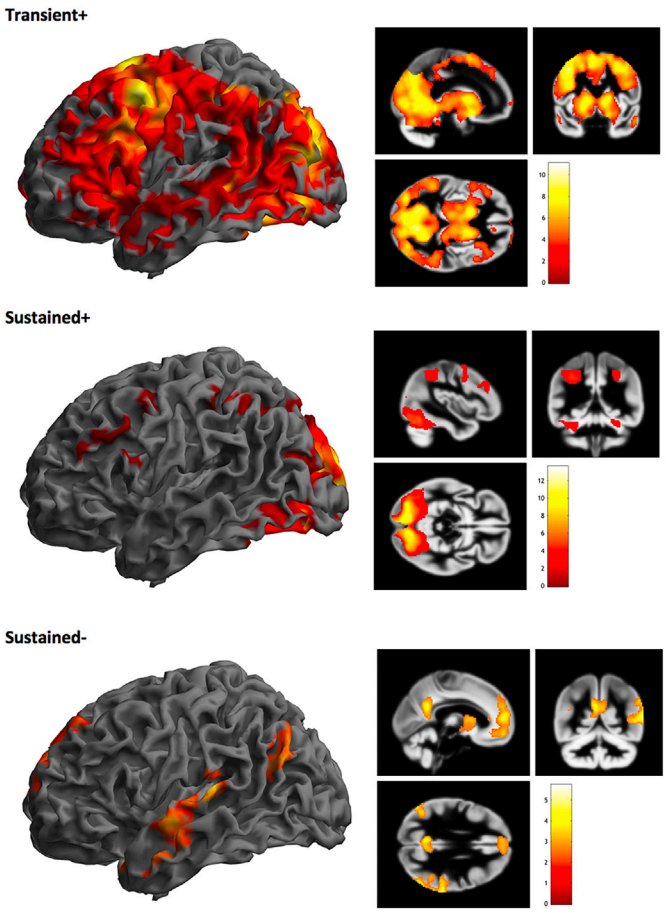

Fig. 6. Transient and sustained brain activation during instruction.
Cluster corrected voxelwise analysis of transient and sustained activation when the rule slide is presented. Top: transient activity at the onset of the rule slide. Middle: sustained activation throughout rule encoding. Bottom: sustained deactivation during rule encoding.

supplemental FIR analysis. Based on this widespread and transient pattern of activity, we suggest that the instruction of new rules leads to a widespread and metabolically costly adjustment of the dynamic brain networks that transiently coalesce to support cognitive tasks. This is in accordance with previous research showing system-wide alterations in network topology associated with switch related activation costs (Hampshire and Owen, 2006) or more effective behavioural performance (Kitzbichler et al., 2011), although we note that the effects observed here were unexpectedly strong and widespread. It logically follows that the instruction-related activation spike should not be viewed as indicative of the distinct contribution that the anterior caudate, or any other brain region, has in learning. Instead this should be interpreted as a global network effect that is definitive of how the brain switches between or updates task programmes.

The above results highlight the importance of distinguishing between global and local phenomena when seeking to understand the dynamic network mechanisms that underlie human cognition. Although many brain regions show similar sensitivities to rule presentation, they could still be differentiated with respect to other demands of the task, most notably the activation profiles of the anterior caudate, FPNs and DMN were dissociated by their sensitivities to the encoding period and across the practice stages. This provided a replication of results from our and others previous studies (Hampshire et al., 2016; Mohr et al., 2016; Toni et al., 2001). More specifically, decreasing frontoparietal network activation with practice was evident in both this and our previous study (Hampshire et al., 2016). This accords well with a working-memory role for this network (Constantinidis and Klingberg, 2016; Pessoa et al., 2002), i.e., supporting temporary programmes for performing novel tasks prior to consolidation (Hampshire et al., 2016; Mohr et al., 2016; Ruge and Wolfensteller, 2010, 2013). In the current study, the practice-related decline in frontoparietal activity appeared to be somewhat faster than our previous study. We believe that this is due to the prolonged instruction period of 16 s, which allowed for preparatory processes such as mental simulation of the task. Indeed, alongside instruction and reinforcement, mental simulation may be considered a sub-type of intentional learning (Taylor et al., 1998). In accordance with this view, there was sustained visual and frontoparietal activity during that time. As per the previous study, we also replicated the observation of event-related activation throughout all practice stages within a set of visual and motor areas, alongside areas within the cingulate, insula cortex and frontal operculum bilaterally. This is consistent with a role for these brain regions in the basic visual, motor and attentional demands of the task.

Interestingly, there was less correspondence between our current and previous studies in terms of the spatial pattern of increasing activation with practice. On the one hand, the previous study rendered areas of the brain corresponding to the DMN for this contrast. On the other, the current study rendered brain regions that were proximal to the DMN, but that were non-overlapping, although the ROI analysis did at least show a similar direction of effect. Here, peak areas of activation were situated more posteriorly in the medial orbitofrontal cortex, more dorsally in the temporal lobes, and ventral to the posterior cingulate within occipital cortex. Notably, brain areas corresponding more closely to the DMN were also deactivated in the current study, but this was most reliably evident during the prolonged instruction phase as opposed to the practice phase. As above, these results could be a consequence of additional mental simulation during the extended 16 s instruction period prior to applying the discrimination rules.

The correspondences between the caudate activation profile observed here and the predictions of Ruge et al. (Ruge and Wolfensteller, 2016)

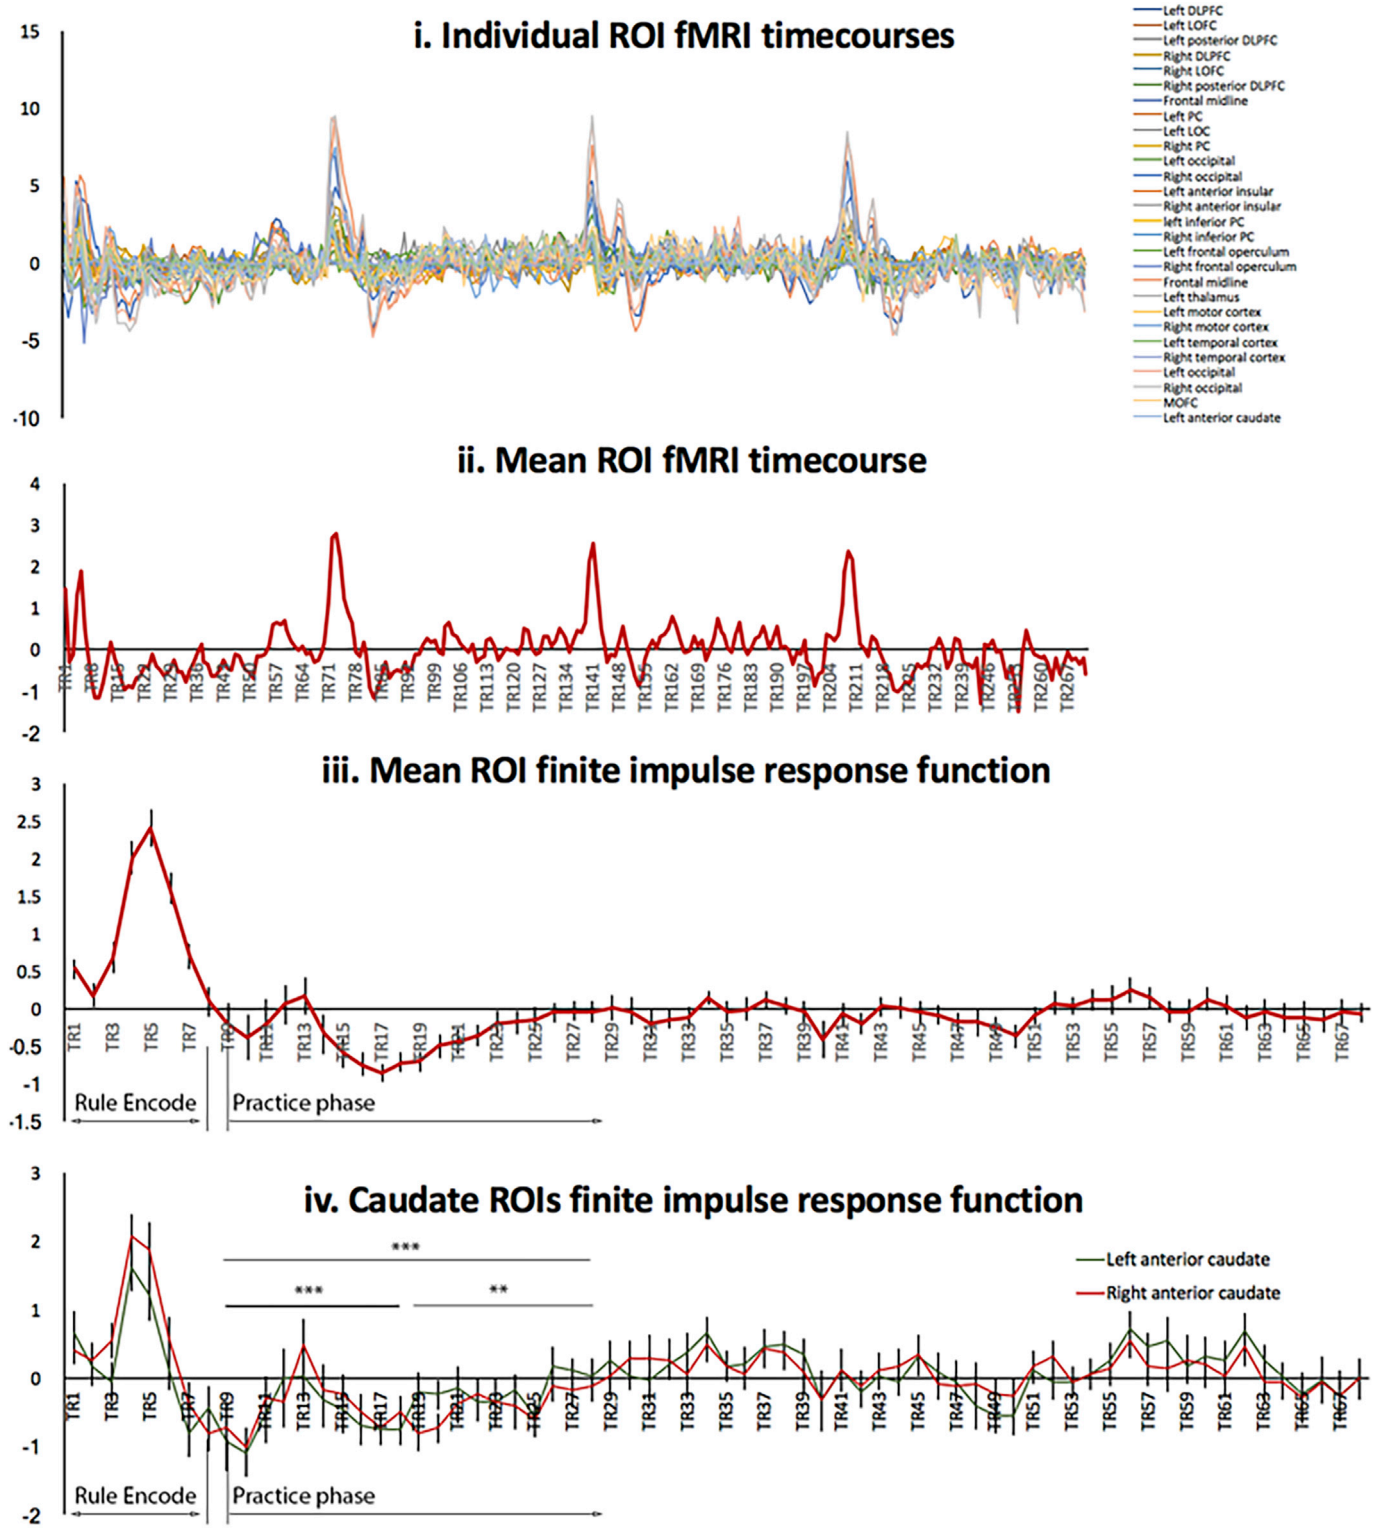

Fig. 7. Timecourse analysis.
ROI activation timecourses. (i) When ROI activations were plotted across time four strong peaks were evident at the points in time when the discrimination rules were defined. (ii) Activation timecourse averaged across all ROIs. (iii). The finite impulse response function generated by averaging across ROIs and rules. (iv) The finite impulse response function for the anterior caudate ROIs only. A primary peak was evident at the 5th time bin, 10 s after the rule slide was presented. A secondary peak was evident at the 13th time bin, 10 s after the onset of the first discrimination. This was followed by an upwards trend in activation throughout the first two practice stages. Time on the X axis is in scan repetitions (TR, each TR = 2s). Activation on the Y axis is in arbitrary units. Error bars report standard error of the mean. $^{**}p < 0.005$ $^{***}p < 0.001$ for repeated measures ANOVA across time bins and averaged for the caudate ROIs.(v) FIR within seven functionally mapped sub-divisions of the striatum (Tziortzi et al., 2014).

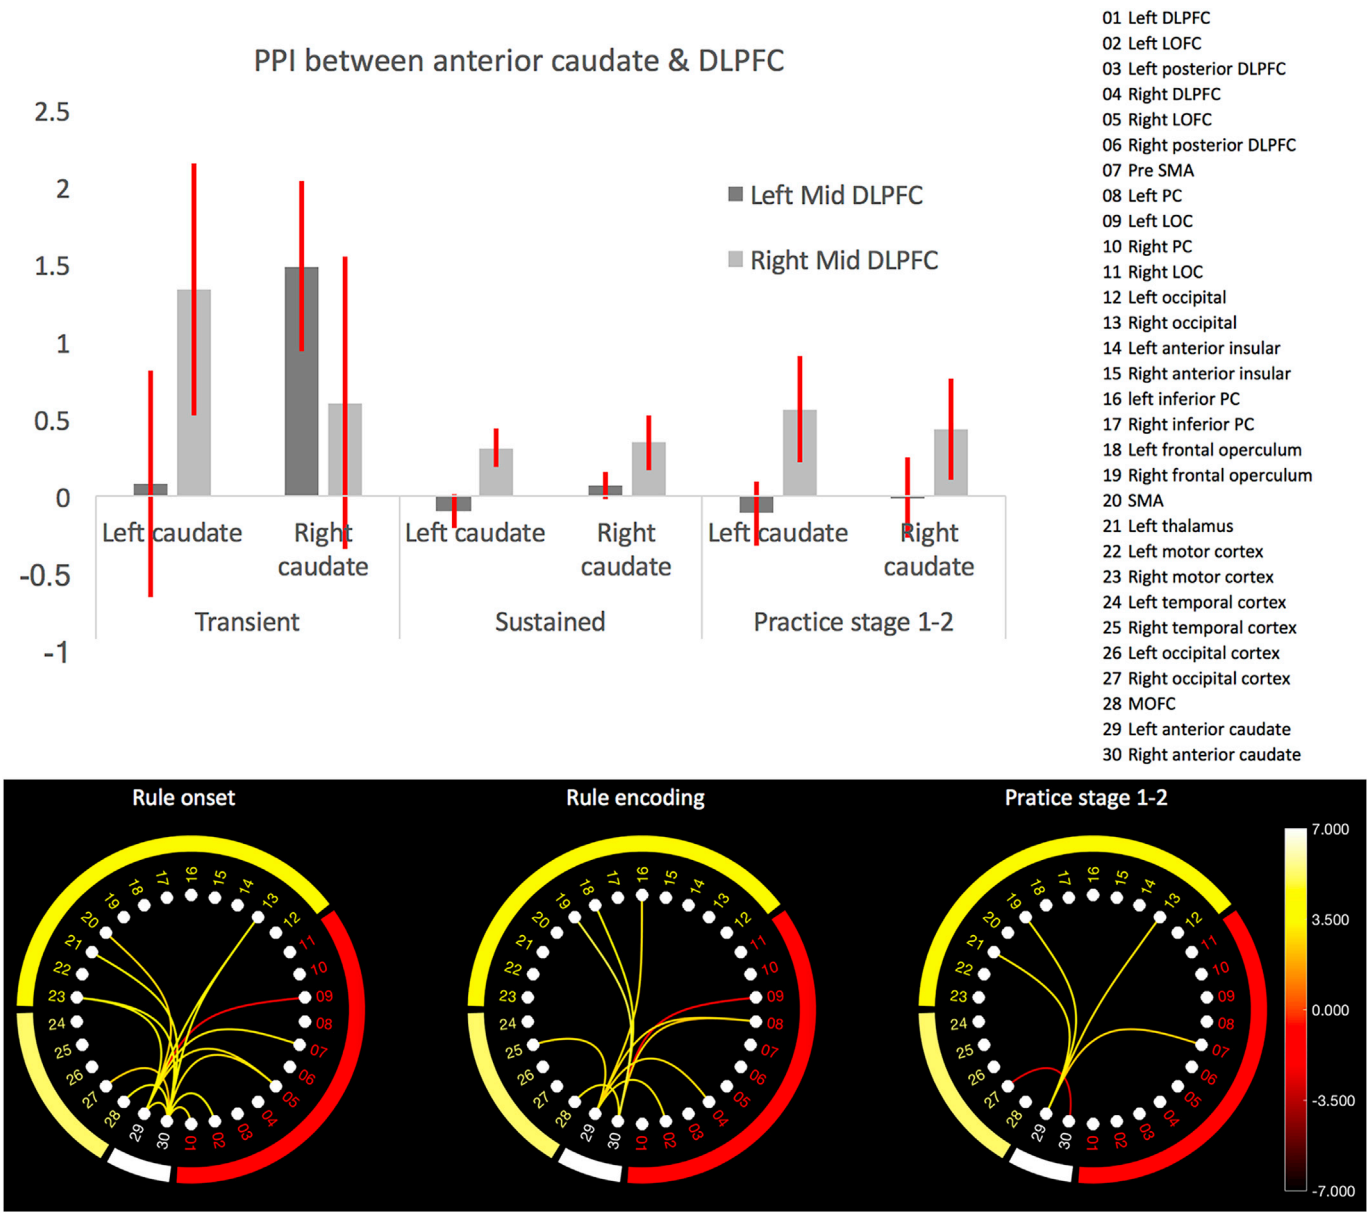

Fig. 8. Effects of learning on the functional connectivity of the anterior caudate.
Top: changes in anterior caudate connectivity with the dorsolateral prefrontal cortex during rule encoding and across practice stages 1 and 2. Error bars report the standard error of the mean. Bottom: broader connectivity changes of the anterior caudate with the discrimination-related network. Yellow = positive red = negative connectivity effects. Schemaballs are thresholded at the exploratory level of $p < 0.01$ two-tailed and uncorrected.

warrant further discussion. In their commentary, the authors made predictions based on their hypothesis that the anterior caudate mediates rule application prior to consolidation. Most notably, they propose that discrimination-related activity in the caudate should increase through the initial stages of rule execution. This should be followed by decreasing activity at later stages of practice as the task program becomes automated/consolidated, during which time other striatal regions such as the posterior caudate (Kim and Hikosaka, 2015) and posterior putamen (Mattfeld and Stark, 2011) were predicted to have a more sustained involvement before a final stage of automatisation, where activity might be entirely cortex-based (Ashby et al., 2010). It was suggested that this activity during the implementation phase might be very rapid for our simple task design, occurring entirely within the first practice block, and to be concomitant with an increase in functional connectivity between the anterior caudate and other regions of the discrimination-related network. The results from our event-related analysis provided some evidence for a non-linear relationship between caudate activity and practice stage. Specifically, a significant main effect of practice stage was evident in the ROI analysis with a peak in activity during the early-mid stage. This accords with the hypothesis that the anterior caudate is involved during the early-mid application phase; however, despite the simplicity of our task design that involvement did not occur within the first application block only. This main effect was also evident in the voxelwise ANOVA across learning stages. Notably, the identified coordinates were proximal to the expected locus within the left anterior caudate (Ruge and Wolfensteller, 2016) and when analysed at a finer temporal grain, showed a smooth increase through the first half of the practice phase, i.e., when rules are transitioning to an automatic process after stabilisation of reaction times

and accuracy. Although present, the activity was much smaller than the observed for the onset of the rule slide.

The results from the finite impulse response function (FIR) were somewhat more supportive of the notion that there are practice-related effects on activation within the anterior striatum. FIR models differ to convolved models insofar as they capture changes in regional activation across time relative to a single specific onset (here the instruction slide); therefore, they can capture changes in baseline activity that are not time-locked to the onsets of the subsequent events. By including a long duration FIR across the whole epoch of instruction and practice, we were able to examine how activation changed across the learning stages without making assumptions regarding the temporal coupling to the onsets of the individual discriminations. When analysed in this way, there was an increase in activity in the caudate across time as the discriminations were practiced. This increase continued for about 1 min, i.e. the timecourse of the first two of our practice stages, each comprising 9 trials. This again could be evidence for a mediating role for the anterior caudate at an intermediate practice stage (Ruge and Wolfensteller, 2016). Notably though, the scale of this effect was again dwarfed by the primary instruction-related spike, and was largely decoupled from the timing of the discrimination trials. It also was notable that this upward trend in activity was not unique to the anterior caudate, instead being evident throughout much of the discrimination-related network, and for all other functional sub-divisions of the striatum that were examined. Moreover, the PPI analysis did not provide evidence of increased anterior caudate coupling with frontoparietal areas that are involved in the discrimination task as a function of practice; instead a trend towards decreased coupling was evident with practice and heightened coupling was evident selectively during rule encoding.

Another notable feature of the FIR model was the unexpected secondary peak in the BOLD timecourse consistent with a smaller transient activation spike at the point in time when the first discrimination was presented. Again, this spike was evident throughout the discrimination-related network and, as per the primary spike, should be viewed holistically as a global network phenomenon. Taken together, these major and minor activation-spikes indicate that the anterior caudate along with wider network are sensitive to the points in time where the participant transitions between the distinct stages of the task; e.g., the transition to rule encoding and the subsequent transition to rule implementation. As discussed above, this interpretation of a role in transitions is on concordance with our previous observation of anterior caudate activation in other task contexts. For example, we previously observed activation in the caudate during contingency reversal learning, that is, when a participant decides based on probabilistic negative feedback to switch from applying a previously rewarded discrimination rule to trying a different one due to a detected change in feedback contingencies (Hampshire et al., 2012). However, we also observed in that same study a similar activation profile when stimuli were all replaced and there was neither a contingency reversal nor negative feedback; in that context, participants simply began a new exploratory phase of behaviour. These results accord with a role in transitioning between task-stages or forming new task programs as opposed to contingency processing or rule application per se.

An important question regards why evidence for effects of rule application within the striatum were weaker here that in some previous studies of IBL. One possibility is that this relates to the simple nature of the discrimination rules that we used in our current task. Our simpler design had a total of four possible stimulus-response mappings for each discrimination rule, whereas Ruge et al. applied more complex rule set (Ruge and Wolfensteller, 2010, 2013); it is not inconceivable that anterior caudate activity scales as the repertoire of competing mappings increases – this would accord with a role in decision making as proposed by other groups (Hiebert et al., 2017). Relatedly, it is likely that the application of a more complex rule set meant that the discrimination rules remained somewhat ambiguous to the participant during the early practice phase. If this were the case, rule formation/encoding in working memory would still be underway across the first set of trials. In support of this view, it is notable that in their study, accuracy was not at ceiling in the early stage of practice and negative feedback was presented to enable the participant to update their internal model of the task program. In this respect, their task design was a mixed instruction-based/reinforcement-based design. Although such designs are necessary when learning more complex rule sets (Sliwinska et al., 2017), as per classic reinforcement learning tasks (Ruge and Wolfensteller, 2016) they can render interpretation of observed brain activation more ambiguous than simple instruction-based learning designs such as the one used here.

A further point raised during review was that the lack of replication of heightened DLPFC-caudate coupling during the practice phase could relate to our approach to PPI analysis, where generalised PPI coefficients were generated in each direction for any pair of ROIs and then averaged. This contrasts with the unidirectional approach taken in some studies. To address this issue, and the possibility of a non-linear relationship between practice stage and frontostriatal connectivity, the PPI coefficients were extracted uni-directionally and separately for each DLPFC-caudate ROI pair. Eight one way repeated measures ANOVAs were applied, one for each connection, with the factor practice stage (4). All eight main effects of stage were non-significant (all $p > 0.2$). Statistical power should be considered. We included fewer rules in our design, and consequently fewer repetitions of each practice stage, which enables the learning process to be examined across a longer temporal window. Here, we observed robust PPI effects for the anterior caudate ROIs during encoding; however, it is the case that designs that involve learning a greater number of rule sets might well detect other, more subtle, functional connectivity effects during the practice phase.

In sum, we report that there is a widespread spike in activity throughout a network of cortical and striatum brain regions when new rule slides are presented. This is followed by more subtle dissociations between these brain regions when rules are being encoded and practiced. Frontoparietal networks are active early in the practice phase. Visual, motor and cingulo-opercular areas are active throughout, whereas the default mode network is most active towards the end of the learning curve. There is tentative evidence of increased anterior caudate activity at an intermediate practice stage, after the frontoparietal activity and behavioural response time peaks. However, this effect was much smaller than the primary instruction-related spike and in the FIR analysis was evident for many other brain regions. We know of no current model of the brain mechanisms that support IBL that predicts the pronounced spike in activation that is evident when instructions are initially presented or to a lesser extent at the onset of the first discrimination trial. Given the scale of these effects any model of IBL mechanisms must account for them to be plausible. We believe that this may be achieved through inclusion of a global destabilisation-stabilisation mechanism when new rules are formed or updated in working memory. A sensible future direction is to determine which brain regions are involved in driving the widespread network dynamic effects that are observed during instruction-based learning, and testing whether these are the same drivers under other types of intentional learning. This may be achieved by application of directed connectivity analysis with tasks that are designed to produce a greater number of rule definition events, i.e., with shorter practice periods, and preferably across a greater number and broader range of rule-switching conditions.

Acknowledgements

This work was funded by an ITMAT Experimental Medicine grant from the NIHR Imperial Biomedical Research Centre and MRC grant MR/R005370/1. IRV was funded by the Wellcome Trust (103045/13/Z).

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.neuroimage.2019.03.002.

References

Ashby, F.G., Turner, B.O., Horvitz, J.C., 2010. Cortical and basal ganglia contributions to habit learning and automaticity. Trends Cognit. Sci. 14, 208–215.

Badre, D., Kayser, A.S., D'Esposito, M., 2010. Frontal cortex and the discovery of abstract action rules. Neuron 66, 315–326.

Boettiger, C.A., D'Esposito, M., 2005. Frontal networks for learning and executing arbitrary stimulus-response associations. J. Neurosci. 25, 2723–2732.

Brainard, D.H., 1997. The Psychophysics toolbox. Spatial Vis. 10, 433–436.

Brass, M., Wenke, D., Spengler, S., Waszak, F., 2009. Neural correlates of overcoming interference from instructed and implemented stimulus-response associations. J. Neurosci. 29, 1766–1772.

Brett, M., Anton, J., Valabregue, R., Poline, J., 2002. Region of Interest Analysis Using an SPM Toolbox. 8th International Conference on Functional Mapping of the Human Brain, Sendai, Japan [abstract].

Brovelli, A., Nazarian, B., Meunier, M., Boussaoud, D., 2011. Differential roles of caudate nucleus and putamen during instrumental learning. Neuroimage 57, 1580–1590.

Chein, J.M., Schneider, W., 2005. Neuroimaging studies of practice-related change: fMRI and meta-analytic evidence of a domain-general control network for learning. Brain Res Cogn Brain Res 25, 607–623.

Cole, M.W., Laurent, P., Stocco, A., 2013. Rapid instructed task learning: a new window into the human brain's unique capacity for flexible cognitive control. Cognit. Affect Behav. Neurosci. 13, 1–22.

Collins, A.G.E., Frank, M.J., 2018. Within- and across-trial dynamics of human EEG reveal cooperative interplay between reinforcement learning and working memory. Proc. Natl. Acad. Sci. U. S. A. 115, 2502–2507.

Constantinidis, C., Klingberg, T., 2016. The neuroscience of working memory capacity and training. Nat. Rev. Neurosci. 17, 438–449.

Eliassen, J.C., Souza, T., Sanes, J.N., 2003. Experience-dependent activation patterns in human brain during visual-motor associative learning. J. Neurosci. 23, 10540–10547.

Friston, K.J., Buechel, C., Fink, G.R., Morris, J., Rolls, E., Dolan, R.J., 1997. Psychophysiological and modulatory interactions in neuroimaging. Neuroimage 6, 218–229.

Gao, W., Lin, W., 2012. Frontal parietal control network regulates the anti-correlated default and dorsal attention networks. Hum. Brain Mapp. 33, 192–202.

Ghahremani, D.G., Monterosso, J., Jentsch, J.D., Bilder, R.M., Poldrack, R.A., 2010. Neural components underlying behavioral flexibility in human reversal learning. Cerebr. Cortex 20, 1843–1852.

Hampshire, A., Chaudhry, A.M., Owen, A.M., Roberts, A.C., 2012. Dissociable roles for lateral orbitofrontal cortex and lateral prefrontal cortex during preference driven reversal learning. Neuroimage 59, 4102–4112.

Hampshire, A., Hellyer, P.J., Parkin, B., Hiebert, N., MacDonald, P., Owen, A.M., Leech, R., Rowe, J., 2016. Network mechanisms of intentional learning. Neuroimage 127, 123–134.

Hampshire, A., Owen, A.M., 2006. Fractionating attentional control using event-related fMRI. Cerebr. Cortex 16, 1679–1689.

Hampshire, A., Thompson, R., Duncan, J., Owen, A.M., 2008. The target selective neural response-similarity, ambiguity, and learning effects. PLoS One 3, e2520.

Hartstra, E., Kuhn, S., Verguts, T., Brass, M., 2011. The implementation of verbal instructions: an fMRI study. Hum. Brain Mapp. 32, 1811–1824.

Hiebert, N.M., Owen, A.M., Seergobin, K.N., MacDonald, P.A., 2017. Dorsal striatum mediates deliberate decision making, not late-stage, stimulus-response learning. Hum. Brain Mapp. 38, 6133–6156.

Kim, H.F., Hikosaka, O., 2015. Parallel basal ganglia circuits for voluntary and automatic behaviour to reach rewards. Brain 138, 1776–1800.

Kitzbichler, M.G., Henson, R.N., Smith, M.L., Nathan, P.J., Bullmore, E.T., 2011. Cognitive effort drives workspace configuration of human brain functional networks. J. Neurosci. 31, 8259–8270.

Koshino, H., Minamoto, T., Yaoi, K., Osaka, M., Osaka, N., 2014. Coactivation of the default mode network regions and working memory network regions during task preparation. Sci. Rep. 4, 5954.

Law, J.R., Flanery, M.A., Wirth, S., Yanike, M., Smith, A.C., Frank, L.M., Suzuki, W.A., Brown, E.N., Stark, C.E., 2005. Functional magnetic resonance imaging activity during the gradual acquisition and expression of paired-associate memory. J. Neurosci. 25, 5720–5729.

Levine, M., 1975. A Cognitive Theory of Learning. Lawrence Erlbaum Associates Software & Alternative Media, Incorporated.

Liu, Z., Braunlich, K., Wehe, H.S., Seger, C.A., 2015. Neural networks supporting switching, hypothesis testing, and rule application. Neuropsychologia 77, 19–34.

Mattfeld, A.T., Stark, C.E., 2011. Striatal and medial temporal lobe functional interactions during visuomotor associative learning. Cerebr. Cortex 21, 647–658.

McLaren, D.G., Ries, M.L., Xu, G., Johnson, S.C., 2012. A generalized form of context-dependent psychophysiological interactions (gPPI): a comparison to standard approaches. Neuroimage 61, 1277–1286.

Mohr, H., Wolfensteller, U., Betzel, R.F., Misic, B., Sporns, O., Richiardi, J., Ruge, H., 2016. Integration and segregation of large-scale brain networks during short-term task automatization. Nat. Commun. 7, 13217.

Passingham, R.E., Rowe, J.B., Sakai, K., 2013. Has brain imaging discovered anything new about how the brain works? Neuroimage 66, 142–150.

Pessoa, L., Gutierrez, E., Bandettini, P., Ungerleider, L., 2002. Neural correlates of visual working memory: fMRI amplitude predicts task performance. Neuron 35, 975–987.

Raichle, M.E., 2015. The brain's default mode network. Annu. Rev. Neurosci. 38, 433–447.

Ruge, H., Wolfensteller, U., 2010. Rapid formation of pragmatic rule representations in the human brain during instruction-based learning. Cerebr. Cortex 20, 1656–1667.

Ruge, H., Wolfensteller, U., 2013. Functional integration processes underlying the instruction-based learning of novel goal-directed behaviors. Neuroimage 68, 162–172.

Ruge, H., Wolfensteller, U., 2015. Distinct fronto-striatal couplings reveal the double-faced nature of response-outcome relations in instruction-based learning. Cognit. Affect Behav. Neurosci. 15, 349–364.

Ruge, H., Wolfensteller, U., 2016. Towards an understanding of the neural dynamics of intentional learning: considering the timescale. Neuroimage 142, 668–673.

Seger, C.A., Spiering, B.J., 2011. A critical review of habit learning and the Basal Ganglia. Front. Syst. Neurosci. 5, 66.

Sliwinska, M.W., Violante, I.R., Wise, R.J.S., Leech, R., Devlin, J.T., Geranmayeh, F., Hampshire, A., 2017. Stimulating multiple-demand cortex enhances vocabulary learning. J. Neurosci. 37, 7606–7618.

Spreng, R.N., Stevens, W.D., Chamberlain, J.P., Gilmore, A.W., Schacter, D.L., 2010. Default network activity, coupled with the frontoparietal control network, supports goal-directed cognition. Neuroimage 53, 303–317.

Taylor, S.E., Pham, L.B., Rivkin, I.D., Armor, D.A., 1998. Harnessing the imagination. Mental simulation, self-regulation, and coping. Am. Psychol. 53, 429–439.

Toni, I., Ramnani, N., Josephs, O., Ashburner, J., Passingham, R.E., 2001. Learning arbitrary visuomotor associations: temporal dynamic of brain activity. Neuroimage 14, 1048–1057.

Tziortzi, A.C., Haber, S.N., Searle, G.E., Tsoumpas, C., Long, C.J., Shotbolt, P., Douaud, G., Jbabdi, S., Behrens, T.E., Rabiner, E.A., Jenkinson, M., Gunn, R.N., 2014. Connectivity-based functional analysis of dopamine release in the striatum using diffusion-weighted MRI and positron emission tomography. Cerebr. Cortex 24, 1165–1177.

Wolfensteller, U., Ruge, H., 2012. Frontostriatal mechanisms in instruction-based learning as a hallmark of flexible goal-directed behavior. Front. Psychol. 3, 192.