



Neural Coding of Visual Objects Rapidly Reconfigures to Reflect Subtrial Shifts in Attentional Focus

Lydia Barnes¹, Erin Goddard², and Alexandra Woolgar^{1,3}

Abstract

■ Every day, we respond to the dynamic world around us by choosing actions to meet our goals. Flexible neural populations are thought to support this process by adapting to prioritize task-relevant information, driving coding in specialized brain regions toward stimuli and actions that are currently most important. Accordingly, human fMRI shows that activity patterns in frontoparietal cortex contain more information about visual features when they are task-relevant. However, if this preferential coding drives momentary focus, for example, to solve each part of a task in turn, it must reconfigure more quickly than we can observe with fMRI. Here, we used multivariate pattern analysis of magnetoencephalography data to test for rapid reconfiguration of stimulus information when a new

feature becomes relevant within a trial. Participants saw two displays on each trial. They attended to the shape of a first target then the color of a second, or vice versa, and reported the attended features at a choice display. We found evidence of preferential coding for the relevant features in both trial phases, even as participants shifted attention mid-trial, commensurate with fast subtrial reconfiguration. However, we only found this pattern of results when the stimulus displays contained multiple objects and not in a simpler task with the same structure. The data suggest that adaptive coding in humans can operate on a fast, subtrial timescale, suitable for supporting periods of momentary focus when complex tasks are broken down into simpler ones, but may not always do so. ■

INTRODUCTION

Human cognition is remarkably flexible. We can fluidly direct our focus toward what we need for our current goal, seamlessly adapt to changes in our environment, and generalize from what we know to solve new problems. Several lines of research suggest that this flexibility emerges from activity in frontoparietal cortex. Cognitively challenging tasks elicit robust activity in the “multiple demand” (MD) system—a distributed network of frontal and parietal cortex recruited by a wide range of tasks (Assem, Glasser, Van Essen, & Duncan, 2020; Fedorenko, Duncan, & Kanwisher, 2013; Duncan, 2010). Damage to this system linearly predicts fluid intelligence scores (Woolgar, Duncan, Manes, & Fedorenko, 2018; Woolgar et al., 2010), which in turn powerfully predict how well we are able to acquire new skills.

The characteristic adaptability of frontoparietal regions means that they are ideally suited to supporting flexible cognition. For example, patterns of activity in the MD system, measured with fMRI, adapt to code information that is relevant for the current task. MD patterns can encode many different aspects of a task (e.g., visual: Jackson, Rich, Williams, & Woolgar, 2016; vibrotactile: Woolgar & Zopf, 2017; for a review see Woolgar, Jackson, & Duncan,

2016), commensurate with a high degree of mixed selectivity in these regions (Fusi, Miller, & Rigotti, 2016; Rigotti et al., 2013). Moreover, MD coding for task-relevant stimuli is enhanced when stimuli are more difficult to discriminate (Woolgar, Williams, & Rich, 2015; Woolgar, Hampshire, Thompson, & Duncan, 2011) and changes to prioritize information that is at the focus of attention (Jackson & Woolgar, 2018; Woolgar, Williams, et al., 2015). Activity in at least one MD region appears to be causal for facilitating task-relevant information processing elsewhere in the MD system (Jackson, Feredoes, Rich, Lindner, & Woolgar, 2021). This may provide a source of bias to more specialized brain regions, for example, through task-dependent connectivity (Cole et al., 2013; see, e.g., the work of Baldauf & Desimone, 2014). Consequently, adaptive coding has been proposed as a central component of goal-directed attention, biasing sensory and motor brain regions to perceive and respond to information that is relevant to our current task.

A key outstanding question concerns the temporal scale of this process. Here, we explore the “attentional episodes” account of flexible behavior (Duncan, 2013), which predicts a fast temporal scale. This account draws on studies of human and artificial intelligence to propose that flexible behavior rests on our ability to break a complex task down into a series of simpler parts, and to focus, moment-to-moment, on the information needed for each part (Duncan, Chylinski, Mitchell, & Bhandari, 2017; Duncan,

¹University of Cambridge, United Kingdom, ²University of New South Wales, Sydney, Australia, ³Macquarie University, Sydney, Australia

2013; Duncan, Schramm, Thompson, & Dumontheil, 2012). Indeed, there is some evidence that this ability may underpin performance on novel problem-solving tasks. For example, explicitly breaking a complex task into simple parts removes the performance gap between people with high and low fluid intelligence scores (Duncan et al., 2017; see also the work of O'Brien, Mitchell, Duncan, & Holmes, 2020). In this matrix reasoning study, participants viewed a 2×2 grid with three of the four squares filled with an image. They had to abstract relationships between the images to fill in the remaining square. Images consisted of multiple features. In the second half of the experiment, each feature was presented separately. These segmented problems were trivial for participants to solve, regardless of whether they struggled or performed well on the difficult, unsegmented problems. This led the authors to propose that participants who were able to solve the unsegmented problems were better able to mentally break them down into their relevant parts. Adaptive coding could be a key component of this segmentation by driving momentary focus toward subsets of the available information in turn.

From these studies, it seems intuitive that flexible cognition involves identifying simple problems that we can solve and addressing them in an ordered sequence. However, we do not have clear insight into whether codes reconfigure quickly enough to prioritize relevant information throughout a task. The bulk of research on adaptive coding in humans uses fMRI. Although these studies show trial-to-trial shifts in what information can be discriminated from activity patterns (e.g., Woolgar, Williams, et al., 2015; Woolgar, Hampshire, et al., 2011), the coarse temporal resolution of fMRI does not support precise, subsecond measurement of changes in task information.

Time-resolved methods, such as electrophysiology, EEG, and magnetoencephalography (MEG), offer promising evidence for rapid changes in task representation. Nonhuman primate studies show that the same frontal neurons can encode object identity and then location within a single trial, as monkeys attended to what and then where an object was (Rao, Rainer, & Miller, 1997). These data demonstrate that the neural population can systematically change its activity pattern in synchrony with the task. However, they are taken from highly trained monkeys and could rely on a learned response rather than instantaneous shifts in a flexible brain system. More recent work by Spaak, Watanabe, Funahashi, and Stokes (2017) demonstrates that, even when the same information is encoded across phases of a task, neurons in primate lateral pFC dynamically update what they encode. This dynamic reallocation of selectivity within a trial makes plausible rapid shifts in the information that these adaptive brain regions represent. In humans, stronger coding for visual features when they are task-relevant compared to task-irrelevant emerges in MEG data as early as 100 msec from stimulus onset (Goddard, Carlson, & Woolgar, 2022; Moerel, Rich, & Woolgar, 2021; Battistoni, Kaiser, Hickey,

& Peelen, 2020; Wen, Duncan, & Mitchell, 2019), with sustained coding of the relevant feature emerging around 200–400 msec in the MEG/EEG signal (Goddard et al., 2022; Grootswagers, Robinson, Shatek, & Carlson, 2021; Moerel et al., 2021; Yip, Cheung, Ngan, Wong, & Wong, 2021). This provides preliminary evidence that population codes for task-relevant features develop rapidly, but this previous time-resolved human neuroimaging work did not require participants to shift their attention within trials, so we do not know how rapidly information codes update to redirect attention in each part of a task.

Rapid reorganization of information coding within a task has been proposed as key component of how we solve complex tasks, but the neural correlates of this have not yet been studied in the human brain. Here, we test the dynamic adaptation of task representations when what is relevant changes within single trials. We used MEG to track shifts in adaptive coding with subsecond precision across fragments of two rapidly changing tasks. Considering the strong association between task difficulty and the brain regions implicated in adaptive coding (Crittenden & Duncan, 2014; Fedorenko et al., 2013), we tested this at two levels of attentional demand. In Experiment 1, we used simple stimuli to track preferential coding of relevant information under low attentional demands. In Experiment 2, we used a complex stimulus space, abstracted decisions, and the presence of distractors to track preferential coding of relevant information under high attentional demands. Across both experiments, we asked whether neural codes for relevant stimulus information rapidly reconfigure when what is relevant changes mid-trial.

METHODS

Participants

Participants were selected to (a) have normal or corrected-to-normal visual acuity and normal color vision, (b) be right-handed, (c) have no exposure to fMRI in the previous week, (d) have no nonremovable metal objects, and (e) have no history of neurological damage or current psychoactive medication. Prospective participants were informed of the study's selection criteria, aims, and procedure, through a research participation site.

For Experiment 1, 20 participants (17 women, 3 men, mean age 25 ± 6 years) were recruited from the paid participant pool at Macquarie University (Sydney). They gave written informed consent before participating and were paid AUD\$30 for their time. Ethics approval was obtained from the Human Research Ethics Committee at Macquarie University (5201300602).

For Experiment 2, 20 participants (16 women, 4 men, mean age 31 ± 12 years) were recruited from the volunteer panel at the MRC Cognition and Brain Sciences Unit (Cambridge). They gave written informed consent before each testing session and were paid GBP£40 for their time.

Participants were additionally asked to only volunteer if they had existing structural MRI scans on the panel database. Two participants took part before completing a structural scan; one obtained a scan through another study conducted at the MRC Cognition and Brain Sciences Unit, and the other returned for a separate MRI session as part of this study. This participant gave written informed consent before completing the structural scan and was paid an additional GBP£20 for this component of their time. Ethics approval was obtained from the Psychology Research Ethics Committee at the University of Cambridge (PRE.2018.101).

Stimuli

Stimuli were created in MATLAB (The MathWorks, v2012b) and presented with Psychtoolbox (Kleiner et al., 2007; Brainard, 1997). In Experiment 1, they were displayed with an InFocus IN5108 LCD back projection monitor at a viewing distance of 113 cm. In Experiment 2, they were displayed with a Panasonic PT-D7700 projector at a viewing distance of 150 cm.

Experiment 1 stimuli consisted of four novel objects (Op de Beeck, Baker, DiCarlo, & Kanwisher, 2006; see Figure 1) that were either “cubie” or “smoothie” shaped, and green or red (red, green, blue 0–194–155 and 224–0–

98). Colors were chosen for high chromatic variation and strong contrast against the dark gray background (red, green, blue 30–30–30).

Experiment 2 stimuli consisted of 16 novel “spiky” objects, adapted from the Op de Beeck et al. (2006) “spiky” stimuli, selected at four points on a spectrum of red to green, and upright to flat (Goddard et al., 2022). Color values were numerically equally spaced in $u'v'$ color space between [0.35, 0.53] and [0.16, 0.56]. Shapes were also equally spaced to create four steps in orientation from upright to flat. Each step included 100 shape exemplars, with different spikes indicating the orientation, to discourage participants from judging orientation based on a single spike.

Task

Experiment 1 used simple displays and stimuli, optimized for strong visual signals. Each block began with a written cue instructing participants to attend to the color of the first object and the shape of the second object, or vice versa. On each trial, participants viewed two brief displays (100 msec), each followed by a delay (800 msec; see Figure 1). Finally, they were prompted to select an object from a choice display that comprised the combination of the remembered features. All four objects appeared on the

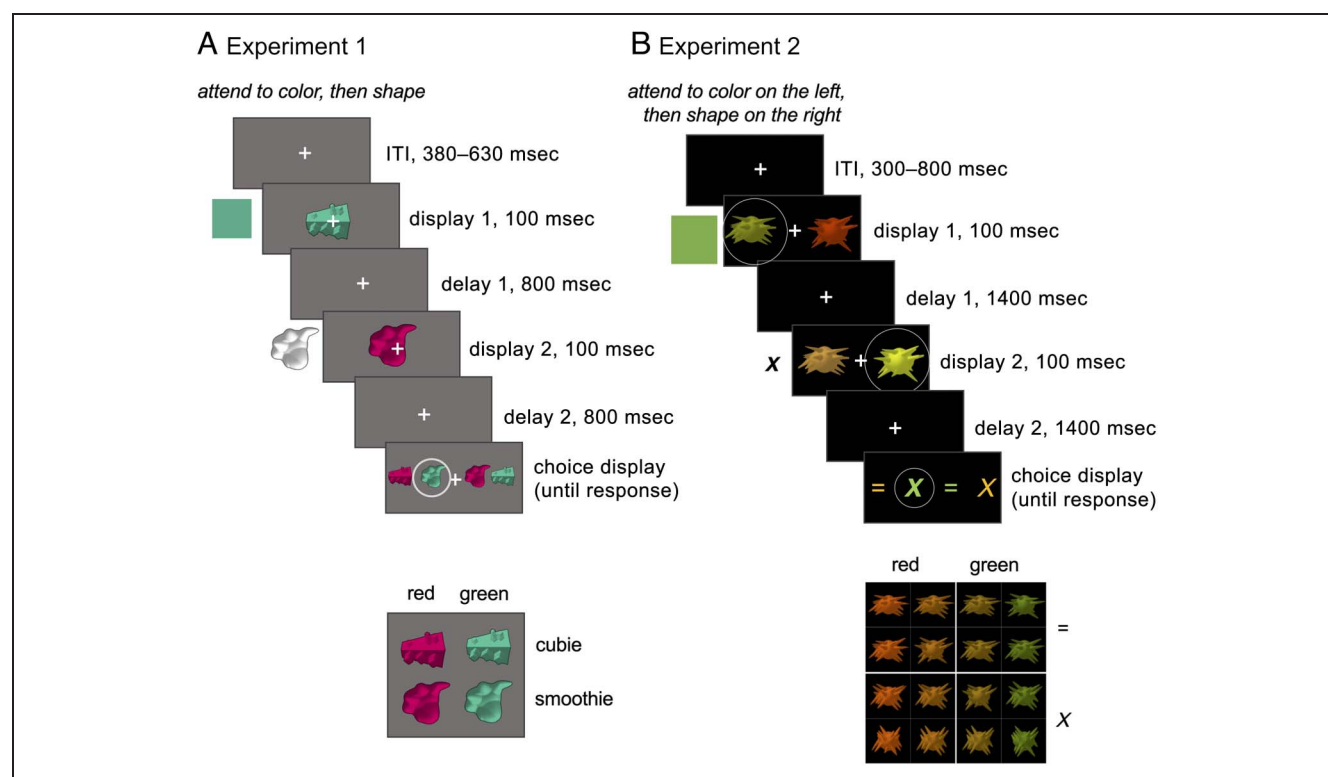


Figure 1. Stimuli and example trials for Experiments 1 and 2. Relevant information for each epoch is shown beside the display. A shows an example trial for Experiment 1, with a single object on each display. In this trial, the relevant features are “green” (Target 1) and “smoothie” (Target 2), resulting in a “green smoothie” response on the choice display. Stimuli could be red, green, “cubie,” or “smoothie.” B shows an example trial for Experiment 2, in which the participant was cued to attend to color on the left and then shape on the right. The relevant features were thus green and “X,” leading to a response of “green X” on the choice display. Stimuli varied in four steps from red to green, and from X to =, but were assigned to binary red / green, X / = categories. Circles represent the focus of attention and correct choice and were not shown to participants.

choice display, and participants selected the object that matched the color and shape they had extracted from the preceding displays. For example, under the rule “attend shape, then color,” if the first object was a “cubie” and the second object was “red,” the target on the choice display was a red cubie. Participants indicated their choice by pressing one of four buttons on a bimanual fiber optic response pad operated with the four fingers of the right hand. The mapping from object location to response button was intuitive (far left button for far left object, etc.) and consistent across trials; however, the arrangement of the four objects on the choice display varied to prevent participants preparing a motor response until the display screen was shown. Stimulus arrangements were presented in pseudorandom order and balanced within each rule such that all stimuli on the second display were equally preceded by each stimulus on the first display, and the correct choice pertained equally to all motor responses. Objects were sampled with replacement, meaning that the same object could appear in both stimulus displays, but participants could not use the trial sequence to anticipate when this would occur. If a participant made three consecutive incorrect or slow responses (> 3 sec), the task was paused and the cue was presented again until the participant verbally confirmed that they understood the rule for that block. Average accuracy and response times were displayed at the end of each block.

Experiment 2 followed the structure of Experiment 1, but used simultaneously presented objects and subtler stimulus discriminations, optimized for high attentional load. For this experiment, each display contained two objects. Participants were cued to both a location and feature, for example, “attend to shape on the right, then color on the left.” Relevant location and feature always changed from Display 1 to Display 2, creating four possible rules. Delay periods were increased to 1500 msec to allow accurate responses, following piloting of the task. Participants judged the color and shape category of the target objects’ features. The choice display contained the symbols X and =, presented in the average of the two “red” colors and the average of the two “green” colors, to represent the four possible answers. These symbols were chosen to encourage participants to make category-level decisions about the objects. As in Experiment 1, the spatial arrangement of the items on the choice display was updated on each trial.

Procedure

Experiment 1

Each participant first completed four blocks of 10 practice trials outside the shielded room. These were identical to test trials except that (a) participants received feedback of “correct,” “incorrect,” or a red screen signifying a slow response (> 3 sec), on every trial, (b) display durations in the first two practice blocks were slowed from 100 to

500 msec to ease participants into the task, and (c) response key codes were marked on the choice display to train participants in the location-response mapping. Once in the MEG scanner, participants completed eight blocks of 96 trials each, with feedback at the end of each block. Each block lasted approximately 7 min. Blocks alternated between the two rules, “attend shape, then color” and “attend color, then shape,” with the order counterbalanced across participants.

Experiment 2

Participants learned the stimulus categories (red vs. green, upright vs. flat) and the task in a separate training session. Training could be on the day of or the day before the scanner session. Training consisted of two blocks of 50 category learning trials, in which they saw a single object for 100 msec and pressed a button to indicate its shape or color category. They then began training on the core task. Within-trial delay periods began at 4 sec and reduced to 1.5 sec in three steps (3 sec, 2 sec, and 1.5 sec). Participants completed a minimum of 10 trials at each of the four speeds for each of the four rules (i.e., at least 40 trials per rule). After 10 trials were completed, the speed increased when the participant got eight trials correct in any 10 consecutive trials. Feedback was given on each trial by a brighter fixation cross for correct responses and a blue fixation cross for incorrect responses, shown for the first 100 msec of the posttrial interval. This procedure trained each participant to the same criterion without penalizing them for errors early in the block.

Once in the MEG, participants completed four blocks, each corresponding to a single rule and comprising 258 trials, lasting approximately 20 min. Rule order was balanced across participants.

MEG Data Acquisition

Experiment 1

We acquired MEG data in the Macquarie University KIT-MEG laboratory using a whole-head horizontal dewar with 160 coaxial-type first-order gradiometers with a 50-mm baseline (Model PQ1160R-N2; KIT; Uehara et al., 2003; Kado et al., 1999) in a magnetically shielded room (Fujihara Co. Ltd.). First, the tester fit the participant with a cap containing five head position indicator coils. The location of the nasion, left and right pre-auricular, and each of the head position indicators were digitized with a Polhemus Fastrak digitiser. This information was copied to the data acquisition computer to track head position during data collection. Participants lay supine during the scan and were positioned with the top of the head just touching the top of the MEG helmet. Any change in head position relative to the start of the session was checked and recorded after four blocks. MEG data were recorded at 1000 Hz.

Experiment 2

We acquired MEG data with the MRC Cognition and Brain Sciences' Elekta-Neuromag 306-sensor Vectorview system with active shielding. Ground and reference EEG electrodes were placed on the cheek and nose. Bipolar electrodes for eye movements were placed at the outer canthi, above and below the left eye. Heartbeat electrodes were on the left abdomen and right shoulder. Scalp EEG were also applied for a separate project. Head position indicators were placed on top of the EEG cap. Both head shape and the location of the head position indicators were digitized with a Polhemus Fastrak digitiser. Head position was recorded continuously through the scan and viewed after each block to ensure that the top of the participant's head stayed within 6 cm of the top of the helmet in the dewar (mean movement across task 3.94 mm, range 0.5–15 mm). Because targets in this experiment could appear to either side of fixation, we also recorded eye movements with an EyeLink 1000 eye tracker, which we calibrated before each block. If we observed more information about the stimulus at the relevant location, eye-tracking data would allow us to measure the contribution of gaze. However, our primary analysis compared features at the same location, so we did not include the eye-tracking data here.

Analyses

MEG Processing

Because of active shielding and artifacts from continuous head position indicators, data from Experiment 2 were first processed with Neuromag's proprietary filtering software (*Maxfilter*, 2010). We applied temporal signal space separation to remove environmental artifacts, used continuous head position information to correct for head movement within each block, and reoriented each block to the subjects' initial head position.

All other processing was the same across experiments. We used a minimal preprocessing pipeline to minimize the chance of removing meaningful data. This was especially appropriate in our case, as our planned multivariate analyses are typically robust to noise (Grootswagers, Wardle, & Carlson, 2016). MEG data were imported into MATLAB v2018b using Fieldtrip (Oostenveld, Fries, Maris, & Schoffelen, 2011) and bandpass filtered (0.01–200 Hz). Trials were epoched from a 100-msec prestimulus baseline to the maximum possible trial duration (Exp 1: 4800 msec, Exp 2: 5000 msec). Principle component analysis was applied to the data, retaining the first components that together captured 99% of the variance. All sensors were included in the analysis.

At the request of a reviewer, we also repeated the analyses for Experiment 2 with additional independent component analysis to remove heart- and eye-related artifacts. We then used systematic averaging before decoding (e.g., averaging across red and green trials when decoding shape) to ensure a balanced test and training set. These

additional analyses (data not shown) produced comparable results to what we report here with minimal preprocessing.

MEG Decoding

We used multivariate pattern analysis to trace the information about rule, color, and shape in each task phase. We then compared the information about color when it was relevant and irrelevant, repeating the comparison for shape. Following previous studies, we expected that rule information, which was known before each trial, would be present throughout the trial and increase briefly after visual displays (Goddard et al., 2022; Hebart, Bankson, Harel, Baker, & Cichy, 2018). We predicted that preferential coding would be reflected in improved decoding of visual features when they were relevant, compared to irrelevant (Goddard et al., 2022; Grootswagers et al., 2021; Moerel et al., 2021; Yip et al., 2021; Battistoni et al., 2020; Wen et al., 2019; Hebart et al., 2018). Increased color information when color was relevant would indicate that information was flexibly coded according to task demands. Our critical comparison, then, was how this happened for the two task phases. If information about the relevant feature was prioritized in both task epochs, this would indicate that preferential coding can reconfigure in line with subsecond shifts in what is relevant to the task.

We first trained a linear classifier (linear discriminant analysis; see the work of Grootswagers et al., 2016) on labeled data from two feature rules—"attend color, then shape" and "attend shape, then color"—using all but one trial from each category. We then tested whether the weights that the classifier had learned to discriminate the training data generalized to the remaining unobserved trials. We repeated the process, leaving out a different pair of trials each time, until all trials had acted as the test data. We then averaged the classification accuracy across all test sets.

For color and shape classification, we trained a linear classifier on labeled data from two categories—for example, "red" and "green"—using all but one trial from each category, for each feature rule separately. For Experiment 2, we decoded pairs of shape or color, at a fixed location, for each feature and location rule. For example, we took trials under the rule "attend color on the left, then shape on the right." For items on the left on the first display, we decoded strong red versus yellow red, yellow red versus yellow green, and so on for all six pairs of color. We then averaged classifier accuracy across the six pairs into a single measure of color information coding in the left hemifield under this rule. We repeated this for each rule to obtain four traces of left hemifield color information coding, representing color information when that location and feature were relevant or irrelevant. We conducted the same pairwise decoding and averaging for color in the right hemifield. Conducting the analyses for each hemifield separately minimized the requirement for the

classifier to generalize patterns over space. Finally, we averaged the four traces of left hemifield color information coding with the corresponding right hemifield traces to produce a single trace for each attention condition: “attended location, attended feature” (the task-relevant trace), “attended location, unattended feature,” “unattended location, attended feature,” and “unattended location, unattended feature.” The two traces for color (or shape) information at the attended location parallel the two traces for each target in Experiment 1 and form the central part of our analysis.

Statistical Tests

We tested whether decoding accuracy scores were above chance using a null distribution generated from the data. To generate this, we permuted the predicted class labels so that they were randomly assigned over trials (Bae & Luck, 2019). We calculated decoding accuracy as above and repeated the process 10,000 times to produce a decoding distribution for each participant and each comparison. We then sampled 10,000 times across participants’ null distributions to form a group-level null distribution. At each time point, we calculated t-scores for classification accuracy relative to the null distribution (Stelzer, Chen, & Turner, 2013). We used a threshold-free cluster statistic (threshold step 0.1; Smith & Nichols, 2009) to flexibly set a cluster-forming threshold to identify peaks in the t-score time course that were more strong and/or sustained than expected from the null distribution ($p < .05$). This maximizes sensitivity to peaks that are most likely to reflect meaningful change while down-weighting peaks that are small or transient (Noble, Scheinost, & Constable, 2020; Vastano, Ambrosini, Ulloa, & Brass, 2020; Pernet, Latinus, Nichols, & Rousselet, 2015; Mensen & Khatami, 2013; Smith & Nichols, 2009). We then used this threshold to correct for multiple comparisons at the cluster level across the whole trial. Decoding onset was the onset of the first cluster for which decoding accuracy was reliably above chance.

For between-conditions comparisons, we contrasted the decoding trace for the target when it was the relevant or irrelevant feature using a two-sided t test, implemented in CoSMoMPPA (Oosterhof, Connolly, & Haxby, 2016) with threshold-free cluster enhancement and a threshold step of 0.1 ($p < .05$; Smith & Nichols, 2009; Figures 4 and 5).

For Experiment 2, we also conducted secondary analyses to assess the combined effects of spatial- and feature-selective, as reported in the work of Goddard et al. (2022). We conducted 2×2 ANOVAs to test, for each time bin, whether stimulus color and shape information coding was boosted (1) at the relevant compared to irrelevant location, (2) when that stimulus feature was relevant for the task compared to when it was irrelevant, and (3) when both feature and location were relevant compared to all

other attention conditions. We quantified these as main effects of Spatial and Feature-Selective Attention, and as a planned comparison between the coding of the reported feature at the attended location and the coding of that feature at that location in the other three attention conditions (following our prediction from Goddard et al., 2022). For example, we contrasted decoding for color on the left when people were attending to color on the left, with decoding for color on the left when attending to shape on the left, color on the right, and shape on the right. We present the results of these secondary analyses in Figure 6.

Lastly, in Experiment 2, we asked whether attentional effects had similar temporal profiles in Epoch 1 and Epoch 2 of the trial. We epoched the stimulus decoding traces for the target, separately around the first and second stimulus displays (0–1500 msec), using the same pretrial baseline (–100 to 0 msec) for all traces. This created four overlaid traces, a relevant and an irrelevant feature trace for Epoch 1 and Epoch 2. We conducted a 2×2 ANOVA with main effects of Relevance and Epoch. An interaction term tested our hypothesis that preferential coding of relevant information emerges earlier, or is more substantial, in one epoch compared to the other.

RESULTS

Behavioral Performance

In Experiment 1, median accuracy was 93.3% (SD 7.5%), with median RT of 829.2 msec (SD 210.7 msec). In Experiment 2, median accuracy was 75.9% (SD 10.9%), with median RT of 665.2 msec (SD 92.1 msec). In both tasks, chance accuracy was 25%.

Rule Information Coding

We trained a classifier to discriminate between feature attention rules (“attend shape, then color” from “attend color, then shape”) from MEG data to extract a time course of rule information coding (Figure 2). Because the rule was cued at the start of the block, we expected that participants might prepare their task set in advance of the stimulus display. We anticipated that rule information would be more decodable after each display, when the rule could be applied to extract relevant information (as in the work of Goddard et al., 2022). Indeed, rule information coding emerged early in both experiments, increasing after each stimulus onset, and remaining above chance throughout the trial. Rule information coding gradually ramped up after each display in Experiment 1, whereas in Experiment 2, rule information coding was elevated throughout the trial and peaked steeply after each display. For Experiment 2, we collapsed the feature rule analysis over locations to mirror Experiment 1 (Figure 2). We also decoded the location rule (i.e., “attend left, then

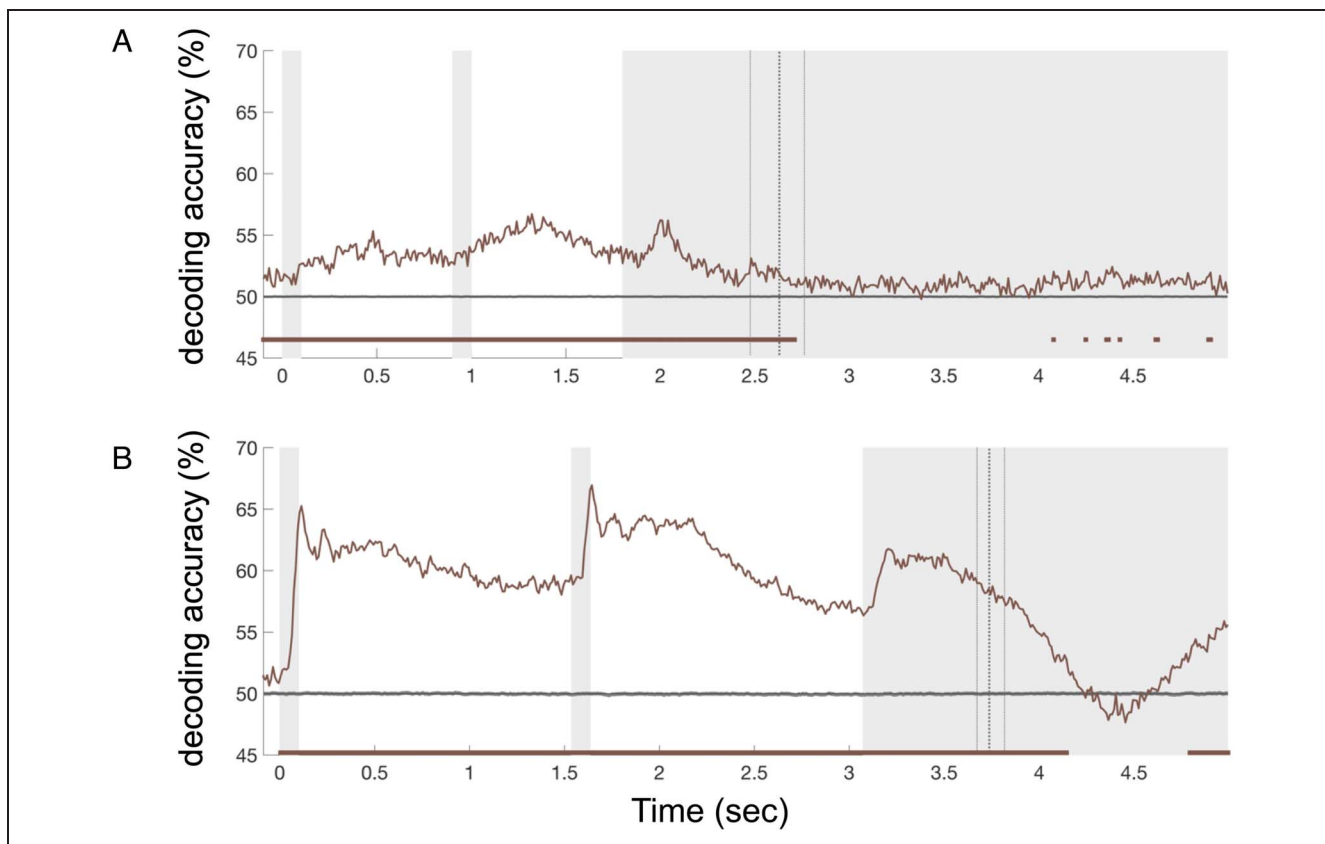


Figure 2. Feature rule decoding (“attend color then shape” vs. “attend shape then color”) for Experiment 1 (A) and Experiment 2 (B). Vertical gray patches mark the stimulus displays and the maximum possible duration of the choice display. Vertical dotted lines mark the median response time with one quartile on either side. Horizontal gray lines show chance (50%) bounded by the 95% confidence interval for the null mean, which we estimated from permutation-based null data. Time points at which decoding was reliably different to the null based on threshold-free cluster correction are marked below the trace in brown.

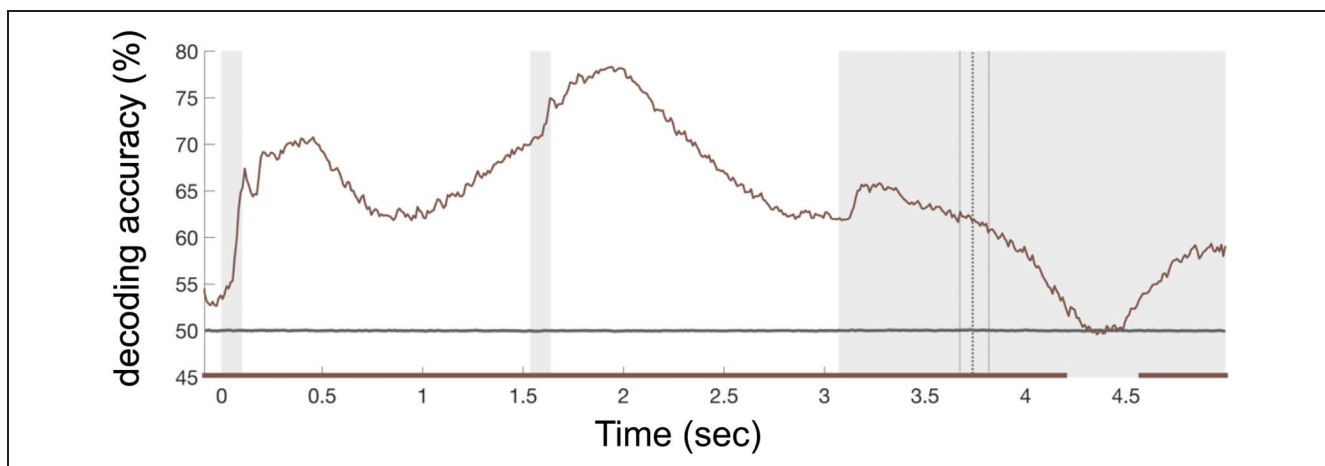



Figure 3. Location rule decoding (“attend left then right” vs. “attend right then left”) for Experiment 2. Vertical gray patches mark the stimulus displays and the maximum possible duration of the choice display. Vertical dotted lines mark the median response time with one quartile on either side. Horizontal gray lines show chance (50%) bounded by the 95% confidence interval for the null mean, which we estimated from permutation-based null data. Time points at which decoding was reliably different to the null based on threshold-free cluster correction are marked below the trace in brown.

Figure 4. Color (A) and shape (B) decoding for Experiment 1. (A) and (B) show decoding traces for the first and second targets in the upper and lower panels. Decoding accuracies are shown for each feature when it was relevant (blue) or irrelevant (orange) for the task. Gray bars mark the stimulus and response display durations. Vertical lines show the median response time, \pm one quartile. Times at which decoding was greater than chance, $p < .05$ using a cluster-based correction for multiple comparisons, are marked below each trace in the corresponding color. Relevant information coding did not reliably exceed coding for the irrelevant feature at any time point.



right” and “attend right, then left”), which we show in Figure 3 for completeness.

Preferential Coding of Visual Features

Next, we examined the time course with which we could decode stimulus color and shape from the pattern of MEG activity. We quantified this separately when a feature was relevant or irrelevant for the participant’s task so that we could examine the effect of attention on coding of this information. We predicted that both relevant and irrelevant stimulus features would be decodable from the sensor data, but that each feature would be more readily decoded when it was relevant compared to when it was irrelevant, particularly at later time points (Goddard et al., 2022; Moerel et al., 2021; Hebart et al., 2018). In Experiment 1, robust decoding of stimulus information emerged rapidly after

the onset of each display, remaining through the initial part of the delay phase for each epoch (Figure 4). Contrary to our prediction, however, in Experiment 1, there was no reliable evidence of preferential coding of the currently relevant information, in either task epoch, for color or shape information (Figure 4). We subsequently applied a Bayesian analysis of preferential coding, comparing evidence for preferential coding to a point nil, and using a one-sided, medium width ($r = .707$) Cauchy prior over the interval $[0 \text{ Inf}]$, following Teichmann, Moerel, Baker, and Grootswagers (2021). This interval favors detection of small effects, as the bulk of the prior distribution is close to the null value of 0. This analysis showed strong evidence for the null at most time points (Bayes factor $< .1$), for all features. Few or no time points showed strong evidence (Bayes factor > 10) in favor of the hypothesis that decoding accuracy was higher when the feature was task-relevant.

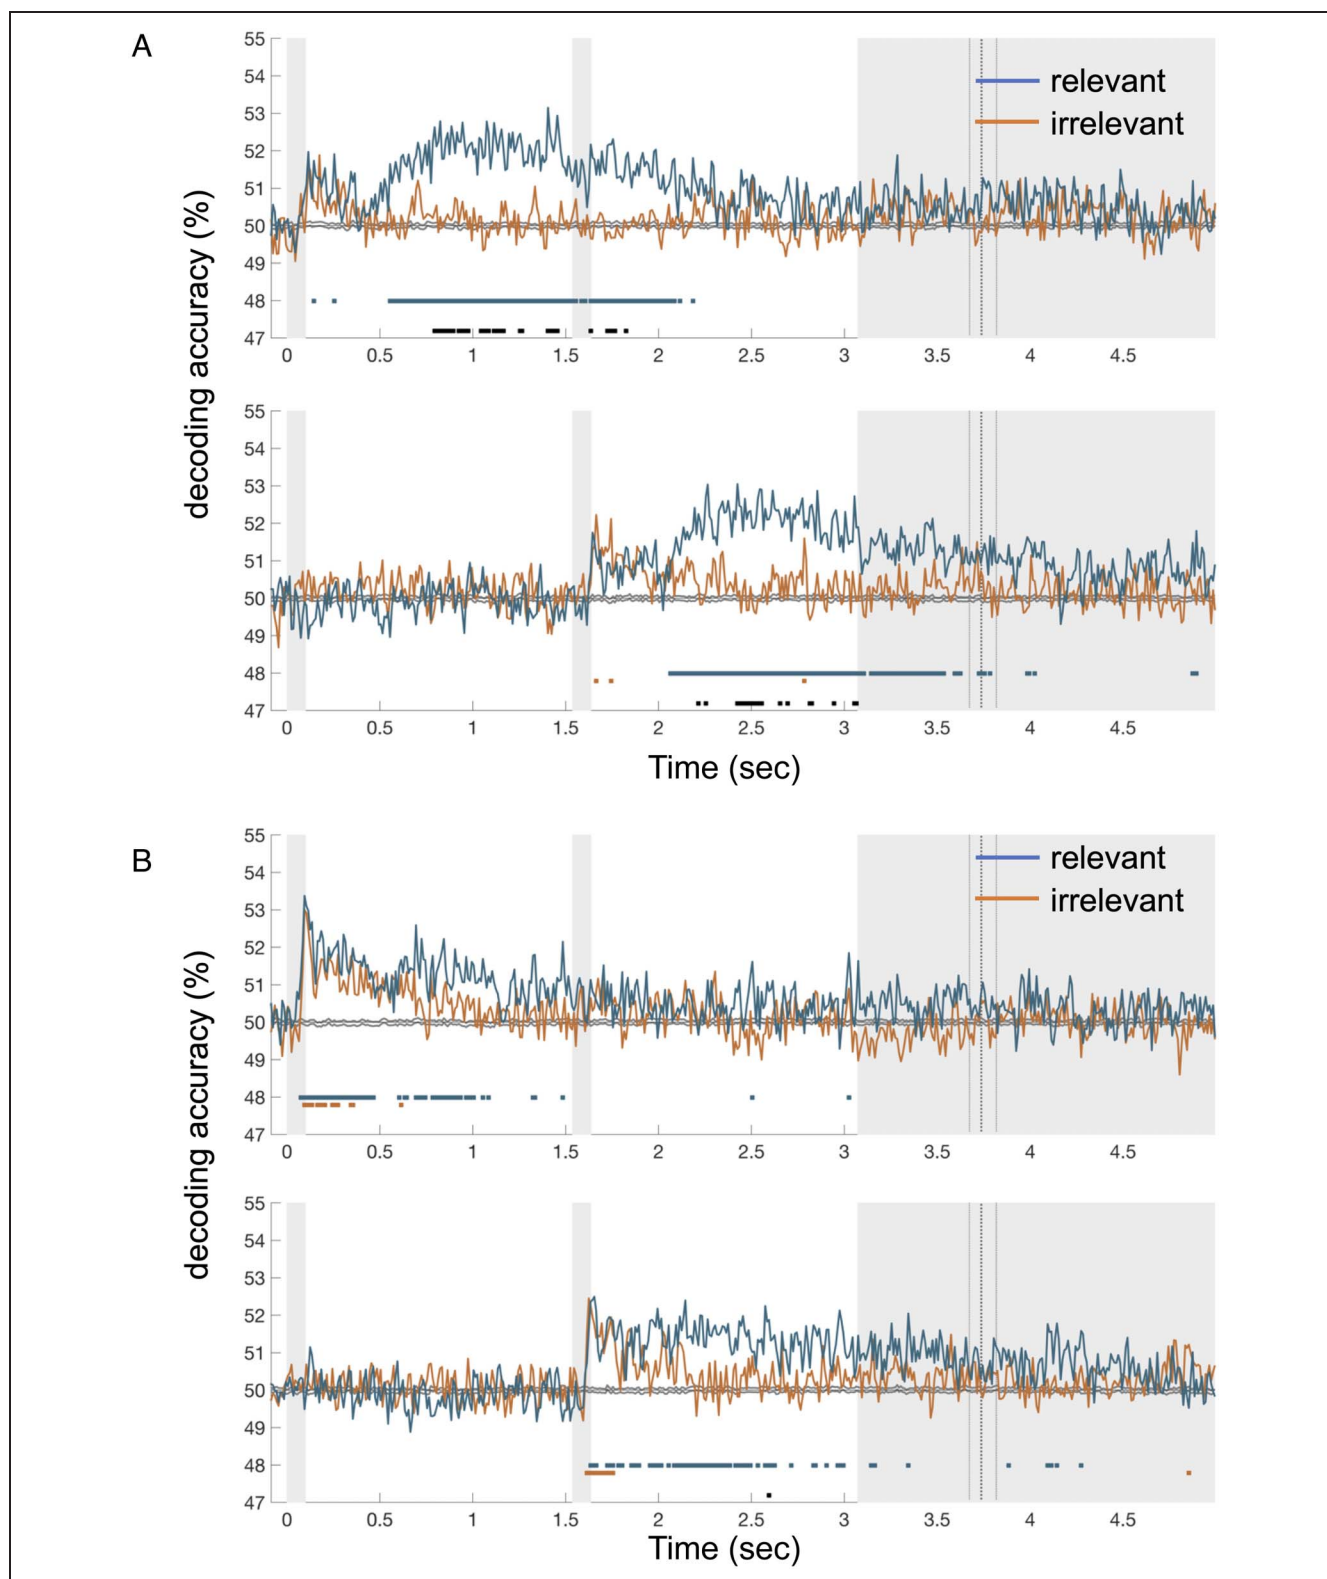


Figure 5. Color (A) and shape (B) decoding for Experiment 2. (A) and (B) show decoding traces for the first and second targets in the upper and lower panels. Decoding accuracies are shown for each feature when it was relevant (blue) or irrelevant (orange) for the task. Gray bars mark the stimulus and response display durations. Vertical lines show the median response time, \pm one quartile. Times at which decoding was greater than chance, $p < .05$, using a cluster-based correction for multiple comparisons, are marked below each trace in the corresponding color. Times at which relevant information coding was reliably above coding for the irrelevant target feature (threshold-free cluster correction, $p < .05$) are marked in black.

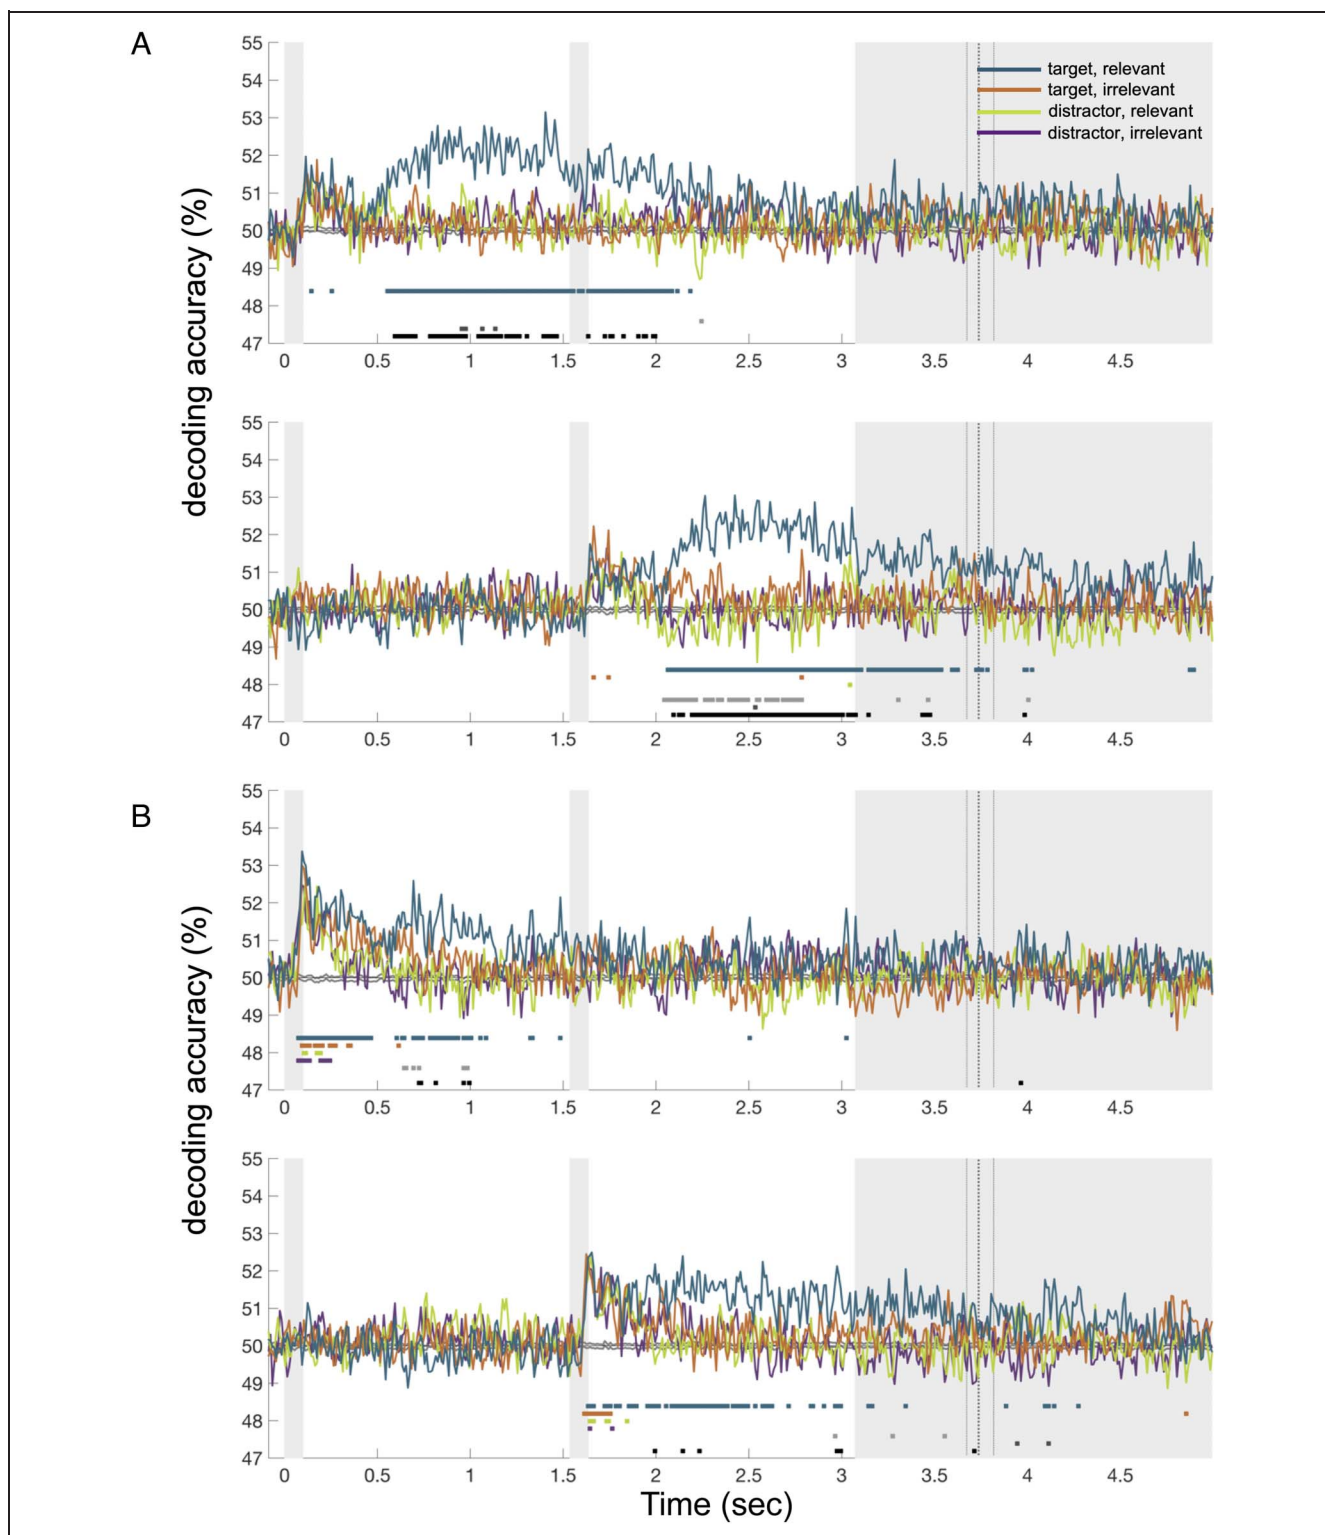


Figure 6. Experiment 2 color (A) and shape (B) decoding for the target and distractor objects on each display. Traces represent decoding accuracy for colors or shapes at the attended location (blue = relevant feature, orange = irrelevant feature), data repeated from Figure 5, as well as at the unattended location (green = attended feature, purple = unattended feature). Times at which each trace was reliably different to chance, at $p < .05$ with a threshold-free cluster correction for multiple comparisons, are marked in the corresponding color. grayscale markers indicate times with a statistically reliable effect of spatial attention (target vs. distractor, light gray), feature attention (attended vs. unattended feature, dark gray), or interaction between spatial and feature attention (relevant feature of target vs. all other features, black).

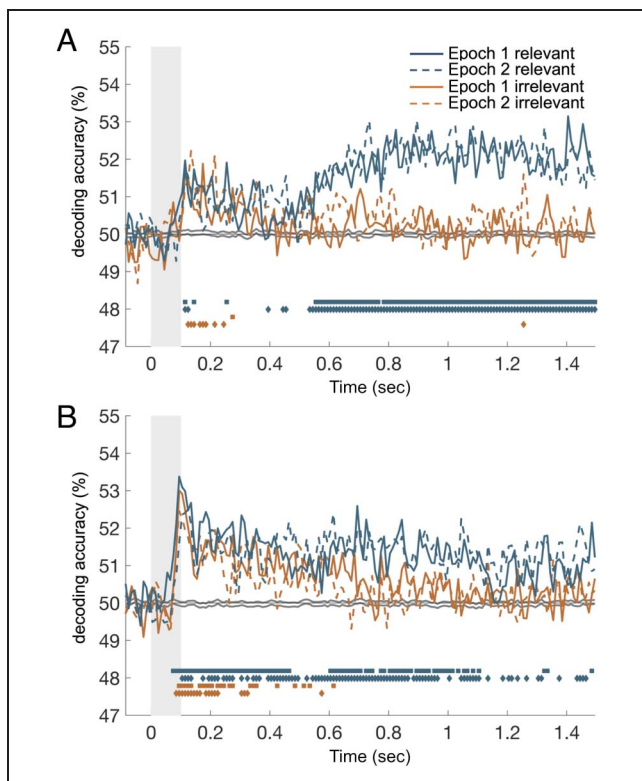


Figure 7. Color (A) and shape (B) decoding for both epochs superimposed. Blue and orange color indicate relevant and irrelevant features, and solid and dotted lines indicate Epochs 1 and 2, respectively. For each trace, time points that reliably differ from chance are marked with colored squares (Epoch 1) or diamonds (Epoch 2). There was no reliable difference between epochs, or interaction between epoch and relevance.

Experiment 2 stimulus decoding was similarly rapid (Figure 5). Although less pronounced (potentially because of the busier displays and more subtle color and shape differences), initial stimulus decoding followed a similar time course to Experiment 1. For coding of color, there was an initial stimulus-driven response peaking at 100 msec, which was similar when that information was relevant or irrelevant, and which occurred for both epochs, although these peaks did not reach statistical significance. For shape, the pattern was broadly similar and statistically significant, with an initial stimulus-driven response at 100 msec from each display onset. Critically, in contrast to Experiment 1, in Experiment 2, we now saw evidence of additional, sustained, preferential coding of relevant information. Whereas decoding for the target's color remained close to chance when that feature was irrelevant, coding for the same information when it was relevant was higher and sustained (Figure 5). Coding of relevant color information was reliably different to chance and to the irrelevant feature trace from approximately 500 msec after stimulus presentation and was sustained into the subsequent trial epoch. We observed the same pattern for shape decoding, with a sustained response only for the relevant

information in both epochs, although this was statistically reliable only in the second epoch. A follow-up analysis revealed no reliable difference between the preferential coding of color and shape.

As a secondary analysis, we additionally considered coding of the features of the distractor object. All four traces (relevant and irrelevant feature of target and distractor) are shown in Figure 6. Color and shape information was briefly decodable in all four attention conditions, after which there was a sustained preferential coding of the relevant target feature compared to the average of all other features (Figure 6, black lines). Where there were main effects of spatial or feature-selective attention, they tended to be accompanied or quickly followed by an interaction of the two attention types. Moreover, when, in an exploratory analysis, we directly compared coding of the irrelevant feature of the target with those of the distractor, or the relevant with irrelevant feature of the distractor, there are no time points where the difference was significant. This implies no advantage for the irrelevant information at the relevant location, or for the relevant information at the relevant location. This replicates similar findings in the work of Goddard et al. (2022), in which main effects of spatial and feature attention emerged briefly before an interaction showed preferential coding specifically for the information that participants needed to retain.

Rapid Coding of Features across Epochs

To compare the dynamics of attentional prioritization across the two epochs, we took the decoding traces for the target in each epoch of Experiment 2 and aligned them in time. We anticipated that the effect of attention (enhancement of relevant information) might develop later in Epoch 2, which reflected a subtrial shift of attention when participants had less time to prepare what they would attend to. However, preferential coding for relevant information in Epoch 2 was comparable to Epoch 1 (Figure 7). We did not observe a main effect of epoch, or an interaction between epoch and relevance. This does not rule out the possibility that shifting attention mid-trial incurs some delay in preferential coding in other circumstances, for example, with more difficult tasks or a shorter within-trial interstimulus interval. However, it demonstrates that humans can rapidly reconfigure their neural codes to prioritize coding of a new stimulus dimension mid-trial, even while holding the previously attended stimulus information in mind. Commensurate with non-human primate work, this highlights our capacity to dynamically code task-relevant information.

DISCUSSION

Understanding how task-sensitive neural codes reconfigure is a key step in tracing how the brain supports adaptive behavior. Here, we conducted two experiments to ask whether the brain can rapidly reconfigure neural codes

for relevant stimulus features when what is relevant changes. In both experiments, participants judged the shape, then color, or vice versa of two targets presented in sequence. When shape and color judgments were easy (Experiment 1), we observed strong coding of all object information. We found no reliable evidence for preferential coding of task-relevant features. By contrast, when the shape and color judgments were difficult and additional distractors were present (Experiment 2), we did see preferential coding for the relevant feature. Crucially, stronger coding for the relevant feature occurred in both phases of the trial, although participants were shifting attention between features mid-trial.

Tracing this process with MEG allowed us to see the temporal evolution of preferential coding in the human brain, showing with millisecond resolution how attention emerges and redirects. Even with this precise temporal detail, Experiment 2 demonstrated a remarkably similar time course for selection of relevant information for the first and second stimuli. We might expect that preferential encoding of the relevant feature in the second epoch would be slower and/or less selective than in the first. For example, a lag or reduction in selectivity could reflect residual attention to the feature that was relevant for the first epoch, or time taken to transition to selective encoding of the second feature. Instead, we did not find any evidence of slower or reduced selectivity in the second epoch, suggesting that, in this paradigm, reconfiguration was fast enough for the relevant feature of the second stimulus to be selected as efficiently as for the first. These findings indicate that, when adaptive coding is engaged, task-relevant information is preferentially coded with remarkable speed even as task demands change within single trials. This provides possible infrastructure for the fast, subtrial switching of attentional sets necessary for a goal-directed behavior (Duncan, 2013).

Although participants successfully performed both tasks, Experiment 1 did not elicit reliable preferential coding of relevant over irrelevant stimulus features. Curiously, both tasks showed strong and sustained representation of the rule (“attend color, then shape”), although only one task showed an effect of rule on stimulus coding. Current explanations of top-down control emphasize both maintaining task information and enhancing relevant stimulus information. For example, both rule and relevant stimulus information can typically be decoded from MD regions in human fMRI (Woolgar & Zopf, 2017; Jackson et al., 2016; Woolgar, Afshar, Williams, & Rich, 2015; Woolgar, Thompson, Bor, & Duncan, 2011) and from frontal cortex in nonhuman primate single-unit recordings (Stokes et al., 2013; Everling, Tinsley, Gaffan, & Duncan, 2006). Disrupting prefrontal function causes reduction in task-relevant information coding (Jackson et al., 2021), and incorrect rule or stimulus information coding predicts incorrect behavioral responses (Woolgar, Dermody, Afshar, Williams, & Rich, 2019). Moreover, the structure of frontal stimulus information predicts subsequent occipital

stimulus information as attentional selection of relevant features emerges (Goddard et al., 2022). In view of these findings, it is plausible that selection occurs through rule information that is maintained by domain-general regions, which in turn selectively enhance relevant stimulus information in both domain-general and task-specific regions. In contrast, in Experiment 1, we observed a dissociation: clear rule coding, but no evidence of enhanced coding of the relevant stimulus features, although the rule defined which stimulus features participants should attend to. Rule decoding increased after the stimulus displays in both tasks, particularly in Experiment 2. These increases could reflect neural responses diverging as participants applied the feature rule to the stimuli, in a way that did not enhance coding of the relevant stimulus features to an extent that our methods could reliably detect. Conversely, increases in rule decoding could be related to a more general shift, such as the widespread reduction in cortical response variance at the onset of a stimulus (Churchland et al., 2010). This highlights the utility of tracing both attentional rule information and rule-related changes in stimulus information, to characterize the impact of the rule on attentional selection. As Experiment 1 shows, the presence of decodable attentional rules does not necessarily translate to preferential coding of relevant stimulus information.

There were several differences between the two experiments that may have contributed to the different results. Experiment 2 was more difficult: Participants responded well above chance level in both tasks, but overall performance was lower in Experiment 2 even after intensive training on the task. In Experiment 1, stimuli were drawn from a set of four objects, with strongly differentiated colors and shapes, and a single object was shown on each display. Because of this small stimulus set, on 25% of trials, the objects on Display 1 and Display 2 were identical, making the task trivial. On the remaining trials, participants had to select differential information from each display to respond accurately. However, there was significantly less information on each display, and less confusability among colors and shapes, than in Experiment 2. Thus, responding to the relevant information could well engage different attentional mechanisms across the two tasks.

Increased selection with increased stimulus complexity is a common theme in many theories of attention. For example, behavioral data demonstrate that although participants can find and respond to targets more quickly in simple displays compared to complex displays, they are also more easily influenced by salient distractors (Lavie, 1995; Lavie & Tsai, 1994). Neuroimaging evidence also suggests that distractors are not processed as deeply when a task becomes more difficult: BOLD activity associated with a distractor stimulus category no longer differentiates repeating and unrepeating distractors when target visibility drops (Yi, Woodman, Widders, Marois, & Chun, 2004). Load theory (Lavie, Beck, & Konstantinou, 2014; Lavie, 1995), takes these findings to argue that selection is

qualitatively different for simple and complex stimuli. In simple environments, perceptual capacity not spent on relevant information spills over to other stimuli. As complexity increases, through the number, similarity, or visibility of the stimuli, we voluntarily direct our fixed capacity toward relevant features and ignore salient distractors.

Load theory does not strictly specify that all features that fall within perceptual capacity limits are equally represented. Based on behavioral responses to distractors under low load, we might predict that relevant and irrelevant features in simple displays are equally encoded, so that preferential coding only occurs when we exceed our perceptual capacity. Our differential findings in Experiments 1 and 2 could be consistent with this view, if Experiment 1 displays fell within most participants' perceptual capacity while Experiment 2 displays exceeded it. However, neuroimaging data so far do not support the idea that we require complex displays to engage preferential coding. Indeed, multivariate analyses of fMRI data show that relevant feature coding in visual cortex (primary visual area and the lateral occipital complex) can be enhanced in simple displays, with this enhancement extending to frontoparietal cortex when stimulus discrimination is difficult (Jackson et al., 2016; Woolgar, Williams, et al., 2015). Recent sensor-space MEG data also show enhanced coding of the relevant stimulus category (objects or letters) although the displays contained only two easily distinguishable objects (Grootswagers et al., 2021). Based on these previous results, we would predict that feature-selective attention produces a relative enhancement of relevant perceptual information in simple displays, although both relevant and irrelevant information can be perceived and recalled. This raises an interesting question: If both simple and complex displays can elicit preferential coding (that we can detect with both fMRI and MEG), why is stimulus coding in our Experiment 1 unaffected by relevance?

Theories focusing on the object-based nature of attention (Baldauf & Desimone, 2014; Chen, 2012) may offer a better explanation for why coding two features of a single object, as in our Experiment 1, and coding two objects, as in the work of Grootswagers et al. (2021), would follow different rules. Behavioral studies demonstrate that we can often report irrelevant features of a target object without any apparent performance cost, suggesting that all features of the object are processed in parallel before we chose specific elements to respond to (Chen, 2012; Duncan, 1984). Under this object-based account of attention, it is unsurprising that we did not observe different responses to the same visual feature when it was the relevant or irrelevant dimension of a target object. Rather, we should expect to see preferential coding of the target object over the distractor. We can see this in the work of Goddard et al. (2022), in which a spatial attention effect emerges before coding of the relevant target feature outstrips all other traces. This same pattern is suggested by our secondary analyses, where brief main effects of spatial attention emerge before preferential coding of the relevant

target feature (Figure 6, Epoch 2 color and Epoch 1 shape). However, object-based accounts struggle to account for the preferential coding of single dimensions of stimuli (e.g., Jackson & Woolgar, 2018; Jackson et al., 2016), which we observed at later time points in Experiment 2.

Biased competition (Reynolds, Chelazzi, & Desimone, 1999; Kastner, Weerd, Desimone, & Ungerleider, 1998; Desimone & Duncan, 1995) provides a possible unifying framework for the load-driven and object-based characteristics of attention. Similar to load theory, this account proposes that complex stimuli trigger attentional selection. Rather than appealing to a threshold for perceptual capacity, biased competition suggests that, as distinct representations of stimulus features in early visual cortex feed forward to shared neural populations in higher visual cortex, competition emerges for what feature will be represented at the higher level, forcing selection to occur (Scalf, Torralbo, Tapia, & Beck, 2013; Reynolds, O'Reilly, Cohen, & Braver, 2012; Desimone & Duncan, 1995). Because integration co-occurs with broadening receptive fields, even spatially segregated shapes can project to the same neurons and compete for in-depth processing. In our study, the two-object displays of difficult-to-discriminate stimuli in Experiment 2 might elicit more competition than the single-object displays in Experiment 1, creating the opportunity for selection, even within the target objects.

Importantly, Duncan (2006) integrates space-, object-, and feature-based attention under the biased competition framework, highlighting that competition drives selection across disparate forms of attention, which can operate independently or in concert. This broader perspective of attention as a family of processes implemented through biased competition has since been embraced by Kravitz and Behrmann (2011), who demonstrate that space-, object-, and feature-based attention can combine to enhance object processing. Combined effects of spatial and feature-based attention have also been observed in nonhuman primates' lateral intraparietal area (Ibos & Freedman, 2016). Goddard et al. (2022) similarly show multiplicative effects of spatial and feature-selective attention give rise to selective coding of only the relevant feature at the relevant location. Using the same stimuli, we replicated this finding, showing that coding of the relevant feature at the relevant location is enhanced relative to the irrelevant feature at that location (Figure 5) and the relevant and irrelevant features of the distractor (secondary analyses, Figure 6), whereas there was no advantage for the irrelevant feature at the relevant location, or relevant feature at the irrelevant location.

From a broader perspective, each of these theories incorporates the suggestion that selection processes do not always alter stimulus representations. In Experiment 1, we saw that people were able to perform a task that required selection without visible impact on the representation of stimulus features. This was consistent with the idea that there were enough resources to process both aspects of those stimuli to a sufficient level before

choosing what would impact behavior. According to the theories above, this capacity to process multiple stimulus features to a high level could depend on the number of features, on object binding, or on a lack of competition, each of which could have facilitated neural coding of stimuli in Experiment 1. Neural network simulations additionally offer some insight into the cost of selection, showing that strong coding of currently relevant task features induces slow reconfiguration to code subsequently relevant information (Musslick, Jang, Shvartsman, Shenhav, & Cohen, 2018). Therefore, there may be a computational benefit to avoiding brain-wide reconfiguration of attentional sets (e.g., within trials) where possible. An adaptive system may be characterized not only by the ability to flexibly prioritize processing of currently relevant information, but the flexibility to only do so when processing demands require it.

We should highlight that the two experiments in this study differed in aspects other than the number and complexity of the stimuli. The experiments were coded, recruited, and run at different testing sites, meaning that the participants, screens, and scanners were unique to each. We were careful to control the stimulus parameters and match the data preprocessing. However, we cannot rule out the possibility that some property of the participant group or scanning equipment impacted the results. In addition, we extended the poststimulus delay periods in Experiment 2 relative to Experiment 1, to account for a large increase in task difficulty. This makes it difficult to directly compare the two tasks. A within-subject study with matched timings will be important in the future to statistically compare preferential coding in simple and complex tasks, and narrow down the circumstances in which patterns do or do not rapidly reconfigure within single trials.

An interesting question is whether we would see the same rapid reconfiguration of what information is preferentially encoded in a less stable context. Here, participants applied the same rule (e.g., “attend color, then shape”) throughout a block of more than 200 trials, before switching to a new rule. This has the advantage of allowing participants to prepare for each trial, enabling us to use the rapid preferential coding of relevant information in Epoch 1 as a baseline against which to compare Epoch 2. However, the repeating rule could have more extensive consequences. An interleaved rule design (e.g., cued trial-by-trial) could potentially uncover limits to rapid reconfiguration, for example, if people struggle to quickly prioritize new information without warning, or are unable to fully prepare one or both parts of the task in advance. In addition, it is well established that frontoparietal BOLD activation is sensitive to difficulty, typically with a U-shaped function, where activation peaks when tasks are difficult but not overwhelming (Van Snellenberg et al., 2015; Jaeggi et al., 2007; Callicott et al., 1999). Thus, to the extent that the current results reflect the engagement of this network, it seems likely that the additional challenge of

reconfiguring task sets on each trial would further impact the results depending where on this function the task sits. Further empirical work is needed to establish the extent to which our results generalize to other designs.

Here, we have shown that human adaptive population codes can reconfigure within a single trial. This supports the current theory, which emphasizes the potential of focusing on each step in a task to produce complex and creative behavior. Surprisingly, where attention effects were seen, the dynamics were comparable for between-trials and within-trial shifts of attentional focus. This provides a potential neural substrate for the rapid creation of attentional episodes in multipart tasks. However, significant effects of attention were only obtained in a demanding version of the task. Although many factors differed between the experiments, the difference could reflect the inherent cost of reconfiguring attention, meaning that it is not always an optimal strategy to engage. Future work will be important to identify what conditions push us toward preferentially coding the relevant information. Spatio-temporally resolved methods, such as source-reconstructed MEG or MEG-fMRI fusion (Moerel et al., 2021; Mohsenzadeh, Mullin, Lahner, Cichy, & Oliva, 2019; Cichy, Pantazis, & Oliva, 2016), paired with systematic manipulation of task difficulty, could further elucidate how domain-general and task-specific brain regions interact to select relevant information under varying task demands. Rapid stimulus streams or self-directed attention shifting could further probe how rapidly the brain can reconfigure neural codes for preferential processing. Furthermore, relating the speed of reconfiguration to measures of fluid ability could clarify the functional importance of adaptive coding timescales. Together with our findings, this will offer rich insight into the biological bases of a mind that adapts to connect our goals with the world around us.

Reprint requests should be sent to Lydia Barnes, MRC Cognition and Brain Sciences Unit, University of Cambridge, 15 Chaucer Road, Cambridge, CB2 7EF UK, or via e-mail: lydiabarnes01@gmail.com.

Funding Information

Medical Research Council (<https://dx.doi.org/10.13039/501100000265>), grant number: SUAG/052/G101400. Macquarie University (<https://dx.doi.org/10.13039/501100001230>), grant number: Research Training Pathway Scholarship. Australian Research Council (<https://dx.doi.org/10.13039/501100000923>), grant numbers: DE200100139, FT170100105.

Diversity in Citation Practices

Retrospective analysis of the citations in every article published in this journal from 2010 to 2021 reveals a persistent pattern of gender imbalance: Although the proportions of

authorship teams (categorized by estimated gender identification of first author/last author) publishing in the *Journal of Cognitive Neuroscience (JoCN)* during this period were $M(\text{an})/M = .407$, $W(\text{oman})/M = .32$, $M/W = .115$, and $W/W = .159$, the comparable proportions for the articles that these authorship teams cited were $M/M = .549$, $W/M = .257$, $M/W = .109$, and $W/W = .085$ (Postle and Fulvio, *JoCN*, 34:1, pp. 1–3). Consequently, *JoCN* encourages all authors to consider gender balance explicitly when selecting which articles to cite and gives them the opportunity to report their article's gender citation balance. The authors of this article report its proportions of citations by gender category to be as follows: $M/M = .549$, $W/M = .137$, $M/W = .059$, and $W/W = .255$.

REFERENCES

- Assem, M., Glasser, M. F., Van Essen, D. C., & Duncan, J. (2020). A domain-general cognitive core defined in multimodally parcellated human cortex. *Cerebral Cortex*, 30, 4361–4380. <https://doi.org/10.1093/cercor/bhaa023>, PubMed: 32244253
- Bae, G.-Y., & Luck, S. J. (2019). Decoding motion direction using the topography of sustained ERPs and alpha oscillations. *Neuroimage*, 184, 242–255. <https://doi.org/10.1016/j.neuroimage.2018.09.029>, PubMed: 30223063
- Baldauf, D., & Desimone, R. (2014). Neural mechanisms of object-based attention. *Science*, 344, 424–427. <https://doi.org/10.1126/science.1247003>, PubMed: 24763592
- Battistoni, E., Kaiser, D., Hickey, C., & Peelen, M. V. (2020). The time course of spatial attention during naturalistic visual search. *Cortex*, 122, 225–234. <https://doi.org/10.1016/j.cortex.2018.11.018>, PubMed: 30563703
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433–436. <https://doi.org/10.1163/156856897X00357>, PubMed: 9176952
- Callicott, J. H., Mattay, V. S., Bertolino, A., Finn, K., Coppola, R., Frank, J. A., et al. (1999). Physiological characteristics of capacity constraints in working memory as revealed by functional MRI. *Cerebral Cortex*, 9, 20–26. <https://doi.org/10.1093/cercor/9.1.20>, PubMed: 10022492
- Chen, Z. (2012). Object-based attention: A tutorial review. *Attention, Perception, & Psychophysics*, 74, 784–802. <https://doi.org/10.3758/s13414-012-0322-z>, PubMed: 22673856
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., et al. (2010). Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nature Neuroscience*, 13, 369–378. <https://doi.org/10.1038/nn.2501>, PubMed: 20173745
- Cichy, R. M., Pantazis, D., & Oliva, A. (2016). Similarity-based fusion of MEG and fMRI reveals spatio-temporal dynamics in human cortex during visual object recognition. *Cerebral Cortex*, 26, 3563–3579. <https://doi.org/10.1093/cercor/bhw135>, PubMed: 27235099
- Cole, M. W., Reynolds, J. R., Power, J. D., Repovs, G., Anticevic, A., & Braver, T. S. (2013). Multi-task connectivity reveals flexible hubs for adaptive task control. *Nature Neuroscience*, 16, 1348–1355. <https://doi.org/10.1038/nn.3470>, PubMed: 23892552
- Crittenden, B. M., & Duncan, J. (2014). Task difficulty manipulation reveals multiple demand activity but no frontal lobe hierarchy. *Cerebral Cortex*, 24, 532–540. <https://doi.org/10.1093/cercor/bhs333>, PubMed: 23131804
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193–222. <https://doi.org/10.1146/annurev.ne.18.030195.001205>, PubMed: 7605061
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, 113, 501–517. <https://doi.org/10.1037/0096-3445.113.4.501>, PubMed: 6240521
- Duncan, J. (2006). EPS mid-career award 2004: Brain mechanisms of attention. *Quarterly Journal of Experimental Psychology*, 59, 2–27. <https://doi.org/10.1080/17470210500260674>, PubMed: 16556554
- Duncan, J. (2010). The multiple-demand (MD) system of the primate brain: Mental programs for intelligent behaviour. *Trends in Cognitive Sciences*, 14, 172–179. <https://doi.org/10.1016/j.tics.2010.01.004>, PubMed: 20171926
- Duncan, J. (2013). The structure of cognition: Attentional episodes in mind and brain. *Neuron*, 80, 35–50. <https://doi.org/10.1016/j.neuron.2013.09.015>, PubMed: 24094101
- Duncan, J., Chylinski, D., Mitchell, D. J., & Bhandari, A. (2017). Complexity and compositionality in fluid intelligence. *Proceedings of the National Academy of Sciences, U.S.A.*, 114, 5295–5299. <https://doi.org/10.1073/pnas.1621147114>, PubMed: 28461462
- Duncan, J., Schramm, M., Thompson, R., & Dumontheil, I. (2012). Task rules, working memory, and fluid intelligence. *Psychonomic Bulletin & Review*, 19, 864–870. <https://doi.org/10.3758/s13423-012-0225-y>, PubMed: 22806448
- Everling, S., Tinsley, C. J., Gaffan, D., & Duncan, J. (2006). Selective representation of task-relevant objects and locations in the monkey prefrontal cortex. *European Journal of Neuroscience*, 23, 2197–2214. <https://doi.org/10.1111/j.1460-9568.2006.04736.x>, PubMed: 16630066
- Fedorenko, E., Duncan, J., & Kanwisher, N. (2013). Broad domain generality in focal regions of frontal and parietal cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, 110, 16616–16621. <https://doi.org/10.1073/pnas.1315235110>, PubMed: 24062451
- Fusi, S., Miller, E. K., & Rigotti, M. (2016). Why neurons mix: High dimensionality for higher cognition. *Current Opinion in Neurobiology*, 37, 66–74. <https://doi.org/10.1016/j.conb.2016.01.010>, PubMed: 26851755
- Goddard, E., Carlson, T. A., & Woolgar, A. (2022). Spatial and feature-selective attention have distinct, interacting effects on population-level tuning. *Journal of Cognitive Neuroscience*, 34, 290–312. https://doi.org/10.1162/jocn_a_01796, PubMed: 34813647
- Grootswagers, T., Robinson, A. K., Shatek, S. M., & Carlson, T. A. (2021). The neural dynamics underlying prioritisation of task-relevant information. *Neurons, Behavior, Data Analysis, and Theory*, 5, 1–17. <https://doi.org/10.51628/001c.21174>
- Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2016). Decoding dynamic brain patterns from evoked responses: A tutorial on multivariate pattern analysis applied to time series neuroimaging data. *Journal of Cognitive Neuroscience*, 29, 677–697. https://doi.org/10.1162/jocn_a_01068, PubMed: 27779910
- Hebart, M. N., Bankson, B. B., Harel, A., Baker, C. I., & Cichy, R. M. (2018). The representational dynamics of task and object processing in humans. *eLife*, 7, e32816. <https://doi.org/10.7554/eLife.32816>, PubMed: 29384473
- Ibos, G., & Freedman, D. J. (2016). Interaction between spatial and feature attention in posterior parietal cortex. *Neuron*, 91, 931–943. <https://doi.org/10.1016/j.neuron.2016.07.025>, PubMed: 27499082
- Jackson, J., Feredoes, E., Rich, A. N., Lindner, M., & Woolgar, A. (2021). Concurrent neuroimaging and neurostimulation reveals a causal role for dlPFC in coding of task-relevant

- information. *Communications Biology*, 4, 1–16. <https://doi.org/10.1038/s42003-021-02109-x>, PubMed: 34002006
- Jackson, J., Rich, A. N., Williams, M. A., & Woolgar, A. (2016). Feature-selective attention in frontoparietal cortex: Multivoxel codes adjust to prioritize task-relevant information. *Journal of Cognitive Neuroscience*, 29, 310–321. https://doi.org/10.1162/jocn_a_01039, PubMed: 27626230
- Jackson, J., & Woolgar, A. (2018). Adaptive coding in the human brain: Distinct object features are encoded by overlapping voxels in frontoparietal cortex. *Cortex*, 108, 25–34. <https://doi.org/10.1016/j.cortex.2018.07.006>, PubMed: 30121000
- Jaeggi, S. M., Buschkuhl, M., Etienne, A., Ozdoba, C., Perrig, W. J., & Nirkko, A. C. (2007). On how high performers keep cool brains in situations of cognitive overload. *Cognitive, Affective, & Behavioral Neuroscience*, 7, 75–89. <https://doi.org/10.3758/CABN.7.2.75>, PubMed: 17672380
- Kado, H., Higuchi, M., Shimogawara, M., Haruta, Y., Adachi, Y., Kawai, J., et al. (1999). Magnetoencephalogram systems developed at KIT. *IEEE Transactions on Applied Superconductivity*, 9, 4057–4062. <https://doi.org/10.1109/77.783918>
- Kastner, S., Weerd, P. D., Desimone, R., & Ungerleider, L. G. (1998). Mechanisms of directed attention in the human extrastriate cortex as revealed by functional MRI. *Science*, 282, 108–111. <https://doi.org/10.1126/science.282.5386.108>, PubMed: 9756472
- Kleiner, M., Brainard, D. H., Pelli, D. G., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in psychtoolbox-3. *Perception*, 36, 1–16.
- Kravitz, D. J., & Behrmann, M. (2011). Space-, object-, and feature-based attention interact to organize visual scenes. *Attention, Perception, & Psychophysics*, 73, 2434–2447. <https://doi.org/10.3758/s13414-011-0201-z>, PubMed: 22006523
- Lavie, N. (1995). Perceptual load as a necessary condition for selective attention. *Journal of Experimental Psychology*, 21, 451–468. <https://doi.org/10.1037/0096-1523.21.3.451>, PubMed: 7790827
- Lavie, N., Beck, D. M., & Konstantinou, N. (2014). Blinded by the load: Attention, awareness and the role of perceptual load. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, 369, 20130205. <https://doi.org/10.1098/rstb.2013.0205>, PubMed: 24639578
- Lavie, N., & Tsal, Y. (1994). Perceptual load as a major determinant of the locus of selection in visual attention. *Perception & Psychophysics*, 56, 183–197. <https://doi.org/10.3758/BF03213897>, PubMed: 7971119
- Maxfilter (2.2). (2010). [Computer software]. Elekta Neuromag.
- Mensen, A., & Khatami, R. (2013). Advanced EEG analysis using threshold-free cluster-enhancement and non-parametric statistics. *Neuroimage*, 67, 111–118. <https://doi.org/10.1016/j.neuroimage.2012.10.027>, PubMed: 23123297
- Moerel, D., Rich, A. N., & Woolgar, A. (2021). Selective attention and decision-making have separable neural bases in space and time. *BioRxiv*. <https://doi.org/10.1101/2021.02.28.433294>
- Mohsenzadeh, Y., Mullin, C., Lahner, B., Cichy, R. M., & Oliva, A. (2019). Reliability and generalizability of similarity-based fusion of MEG and fMRI data in human ventral and dorsal visual streams. *Vision*, 3, 8. <https://doi.org/10.3390/vision3010008>, PubMed: 31735809
- Musslick, S., Jang, J. S., Shvartsman, M., Shenhav, A., & Cohen, J. D. (2018). Constraints associated with cognitive control and the stability-flexibility dilemma. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society* (pp. 806–811). Madison, WI.
- Noble, S., Scheinost, D., & Constable, R. T. (2020). Cluster failure or power failure? Evaluating sensitivity in cluster-level inference. *Neuroimage*, 209, 116468. <https://doi.org/10.1016/j.neuroimage.2019.116468>, PubMed: 31852625
- O'Brien, S., Mitchell, D. J., Duncan, J., & Holmes, J. (2020). Cognitive segmentation and fluid reasoning in childhood. *PsyArXiv*. <https://doi.org/10.31234/osf.io/dt84m>
- Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J.-M. (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*, 2011, 156869. <https://doi.org/10.1155/2011/156869>, PubMed: 21253357
- Oosterhof, N. N., Connolly, A. C., & Haxby, J. V. (2016). CoSMoMVA: Multi-modal multivariate pattern analysis of neuroimaging data in MATLAB/GNU octave. *Frontiers in Neuroinformatics*, 10, 27. <https://doi.org/10.3389/fninf.2016.00027>, PubMed: 27499741
- Op de Beeck, H. P., Baker, C. I., DiCarlo, J. J., & Kanwisher, N. G. (2006). Discrimination training alters object representations in human extrastriate cortex. *Journal of Neuroscience*, 26, 13025–13036. <https://doi.org/10.1523/JNEUROSCI.2481-06.2006>, PubMed: 17167092
- Pernet, C. R., Latinus, M., Nichols, T. E., & Rousselet, G. A. (2015). Cluster-based computational methods for mass univariate analyses of event-related brain potentials/fields: A simulation study. *Journal of Neuroscience Methods*, 250, 85–93. <https://doi.org/10.1016/j.jneumeth.2014.08.003>, PubMed: 25128255
- Rao, S. C., Rainer, G., & Miller, E. K. (1997). Integration of what and where in the primate prefrontal cortex. *Science*, 276, 821–824. <https://doi.org/10.1126/science.276.5313.821>, PubMed: 9115211
- Reynolds, J. H., Chelazzi, L., & Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience*, 19, 1736–1753. <https://doi.org/10.1523/JNEUROSCI.19-05-01736.1999>, PubMed: 10024360
- Reynolds, J. R., O'Reilly, R. C., Cohen, J. D., & Braver, T. S. (2012). The function and organization of lateral prefrontal cortex: A test of competing hypotheses. *PLoS One*, 7, e30284. <https://doi.org/10.1371/journal.pone.0030284>, PubMed: 22355309
- Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., et al. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497, 585–590. <https://doi.org/10.1038/nature12160>, PubMed: 23685452
- Scalf, P. E., Torralbo, A., Tapia, E., & Beck, D. M. (2013). Competition explains limited attention and perceptual resources: Implications for perceptual load and dilution theories. *Frontiers in Psychology*, 4, 243. <https://doi.org/10.3389/fpsyg.2013.00243>, PubMed: 23717289
- Smith, S. M., & Nichols, T. E. (2009). Threshold-free cluster enhancement: Addressing problems of smoothing, threshold dependence and localisation in cluster inference. *Neuroimage*, 44, 83–98. <https://doi.org/10.1016/j.neuroimage.2008.03.061>, PubMed: 18501637
- Spaak, E., Watanabe, K., Funahashi, S., & Stokes, M. G. (2017). Stable and dynamic coding for working memory in primate prefrontal cortex. *Journal of Neuroscience*, 37, 6503–6516. <https://doi.org/10.1523/JNEUROSCI.3364-16.2017>, PubMed: 28559375
- Stelzer, J., Chen, Y., & Turner, R. (2013). Statistical inference and multiple testing correction in classification-based multi-voxel pattern analysis (MVPA): Random permutations and cluster size control. *Neuroimage*, 65, 69–82. <https://doi.org/10.1016/j.neuroimage.2012.09.063>, PubMed: 23041526

- Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., & Duncan, J. (2013). Dynamic coding for cognitive control in prefrontal cortex. *Neuron*, *78*, 364–375. <https://doi.org/10.1016/j.neuron.2013.01.039>, PubMed: 23562541
- Teichmann, L., Moerel, D., Baker, C., & Grootswagers, T. (2021). An empirically-driven guide on using Bayes factors for M/EEG decoding. *BioRxiv*. <https://doi.org/10.1101/2021.06.23.449663>
- Uehara, G., Adachi, Y., Kawai, J., Shimogawara, M., Higuchi, M., Haruta, Y., et al. (2003). Multi-channel SQUID systems for biomagnetic measurement. *IEICE Transactions on Electronics, E86-C*, 43–54.
- Van Snellenberg, J., Slifstein, M., Read, C., Weber, J., Thompson, J., Wager, T., et al. (2015). Dynamic shifts in brain network activation during supracapacity working memory task performance. *Human Brain Mapping*, *36*, 1245–1264. <https://doi.org/10.1002/hbm.22699>, PubMed: 25422039
- Vastano, R., Ambrosini, E., Ulloa, J. L., & Brass, M. (2020). Action selection conflict and intentional binding: An ERP study. *Cortex*, *126*, 182–199. <https://doi.org/10.1016/j.cortex.2020.01.013>, PubMed: 32088407
- Wen, T., Duncan, J., & Mitchell, D. J. (2019). The time-course of component processes of selective attention. *Neuroimage*, *199*, 396–407. <https://doi.org/10.1016/j.neuroimage.2019.05.067>, PubMed: 31150787
- Woolgar, A., Afshar, S., Williams, M. A., & Rich, A. N. (2015). Flexible coding of task rules in frontoparietal cortex: An adaptive system for flexible cognitive control. *Journal of Cognitive Neuroscience*, *27*, 1895–1911. https://doi.org/10.1162/jocn_a_00827, PubMed: 26058604
- Woolgar, A., Dermody, N., Afshar, S., Williams, M. A., & Rich, A. N. (2019). Meaningful patterns of information in the brain revealed through analysis of errors. *BioRxiv*. <https://doi.org/10.1101/673681>
- Woolgar, A., Duncan, J., Manes, F., & Fedorenko, E. (2018). Fluid intelligence is supported by the multiple-demand system not the language system. *Nature Human Behaviour*, *2*, 200–204. <https://doi.org/10.1038/s41562-017-0282-3>, PubMed: 31620646
- Woolgar, A., Hampshire, A., Thompson, R., & Duncan, J. (2011). Adaptive coding of task-relevant information in human frontoparietal cortex. *Journal of Neuroscience*, *31*, 14592–14599. <https://doi.org/10.1523/JNEUROSCI.2616-11.2011>, PubMed: 21994375
- Woolgar, A., Jackson, J., & Duncan, J. (2016). Coding of visual, auditory, rule, and response information in the brain: 10 Years of multivoxel pattern analysis. *Journal of Cognitive Neuroscience*, *28*, 1433–1454. https://doi.org/10.1162/jocn_a_00981, PubMed: 27315269
- Woolgar, A., Parr, A., Cusack, R., Thompson, R., Nimmo-Smith, I., Torralva, T., et al. (2010). Fluid intelligence loss linked to restricted regions of damage within frontal and parietal cortex. *Proceedings of the National Academy of Sciences, U.S.A.*, *107*, 14899–14902. <https://doi.org/10.1073/pnas.1007928107>, PubMed: 20679241
- 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. <https://doi.org/10.1016/j.neuroimage.2010.04.035>, PubMed: 20406690
- Woolgar, A., Williams, M. A., & Rich, A. N. (2015). Attention enhances multi-voxel representation of novel objects in frontal, parietal and visual cortices. *Neuroimage*, *109*, 429–437. <https://doi.org/10.1016/j.neuroimage.2014.12.083>, PubMed: 25583612
- Woolgar, A., & Zopf, R. (2017). Multisensory coding in the multiple-demand regions: Vibrotactile task information is coded in frontoparietal cortex. *Journal of Neurophysiology*, *118*, 703–716. <https://doi.org/10.1152/jn.00559.2016>, PubMed: 28404826
- Yi, D.-J., Woodman, G. F., Widders, D., Marois, R., & Chun, M. M. (2004). Neural fate of ignored stimuli: Dissociable effects of perceptual and working memory load. *Nature Neuroscience*, *7*, 992–996. <https://doi.org/10.1038/nn1294>, PubMed: 15286791
- Yip, H. M. K., Cheung, L. Y. T., Ngan, V. S. H., Wong, Y. K., & Wong, A. C.-N. (2021). The effect of task on object processing revealed by EEG decoding. *BioRxiv*, <https://doi.org/10.1101/2020.08.18.255018>