

Expectation modulates repetition priming under high stimulus variability

Department of Psychology, University of Pennsylvania,
Philadelphia, PA, USA

Present address: Department of Psychology,
Durham University, Durham, England;
and Department of Psychology and Logopedics,
Faculty of Medicine, University of Helsinki,
Helsinki, Finland

Maria Olkkonen



Geoffrey K. Aguirre

Department of Neurology, University of Pennsylvania,
Philadelphia, PA, USA



Russell A. Epstein

Department of Psychology, University of Pennsylvania,
Philadelphia, PA, USA



Neural responses to stimuli are often attenuated by repeated presentation. When observed in blood oxygen level-dependent signals, this attenuation is known as fMRI adaptation (fMRIa) or fMRI repetition suppression. According to a prominent account, fMRIa reflects the fulfillment of perceptual expectations during recognition of repeated items (Summerfield, Trittschuh, Monti, Mesulam, & Egner, 2008). Supporting this idea, expectation has been shown to modulate fMRIa under some circumstances; however, it is not currently known whether expectation similarly modulates recognition performance. To address this lacuna, we measured behavioral and fMRI responses to faces while varying the extent to which each stimulus was informative about its successor. Behavioral priming was greater when repetitions were more likely, suggesting that recognition was facilitated by the expectation than an item would repeat. Notably, this effect was only observed when stimuli were drawn from a broad set of faces including many ethnicities and both genders, but not when stimuli were drawn from a narrower face set, thus making repetitions less informative. Moreover, expectation did *not* modulate fMRIa in face-selective cortex, contrary to previous studies, although an exploratory analysis indicated that it did so in a medial frontal region. These results support the idea that expectation modulates recognition efficiency, but insofar as behavioral effects of expectation were not accompanied by fMRI effects in visual cortex, they suggest that fMRIa cannot be entirely explained in terms of fulfilled expectations.

Introduction

Neural responses to repeated stimuli tend to be smaller than responses to nonrepeated stimuli. This phenomenon, which has been observed both in electrophysiological recordings (Carandini & Ferster, 1997; Henson et al., 2003) and functional magnetic resonance imaging (fMRI) (Dobbins, Schnyer, Verfaellie, & Schacter, 2004; Grill-Spector et al., 1999; Henson, 2003), has been labeled fMRI adaptation or repetition suppression (Barron, Garvert, & Behrens, 2016; Grill-Spector, Henson, & Martin, 2006). Beyond its intrinsic interest as a neurobiological phenomenon, adaptation has become an important tool in cognitive fMRI research because it allows one to examine neural representation by assessing the extent to which one stimulus is treated as a “repetition” of another. For example, fMRI adaptation has been used to investigate the specificity of responses across stimulus variation (e.g. Fang, Murray, Kersten, & He, 2005; Kourtzi & Kanwisher, 2001; MacEvoy & Epstein, 2007; Persichetti, Thompson-Schill, Butt, Brainard, & Aguirre, 2015; Weiner, Sayres, Vinberg, & Grill-Spector, 2010; Winston, Henson, Fine-Goulden, & Dolan, 2004), the dissociability of perceptual processes (Ashida, Lingnau, Wall, & Smith, 2007), and the geometry of neural similarity spaces (Aguirre, 2007; Loffler, Yourganov, Wilkinson, & Wilson, 2005).

Despite its widespread use, the neural mechanisms underlying fMRI adaptation remain controversial

Citation: Olkkonen, M., Aguirre, G. K., & Epstein, R. A. (2017). Expectation modulates repetition priming under high stimulus variability. *Journal of Vision*, 17(6):10, 1–16, doi:10.1167/17.6.10.

doi: 10.1167/17.6.10

Received October 5, 2016; published June 15, 2017

ISSN 1534-7362 Copyright 2017 The Authors



(Epstein & Morgan, 2012). Although earlier theories proposed stimulus-driven mechanisms—such as neural fatigue (Miller & Desimone, 1994; Grill-Spector & Malach, 2001), sharpening of neural representations (Desimone, 1996; Wiggs & Martin, 1998), or facilitation of neural processing (Henson & Rugg, 2003; James & Gauthier, 2006)—more recent research suggests that fMRI adaptation may be driven in part by top-down signals that reflect predictions about the stimuli that one expects to see in the world. In particular, several studies have found that the magnitude of fMRI adaptation is greater in blocks of trials when repetitions are more frequent compared with blocks when repetitions are relatively infrequent (Grotheer, Hermann, Vidnyánszky, & Kovács, 2014; Kovács, Kaiser, Kaliukhovich, Vidnyánszky, & Vogels, 2013; Larsson & Smith, 2012; Summerfield et al., 2008, Summerfield, Wyart, Johnen, & de Gardelle, 2011). Thus, for example, when trials consist of sequentially presented paired faces that can either be the same (repeat) or different (nonrepeat), and these trials are presented within blocks for which most of the trials are repeat trials or most of the trials are nonrepeat trials, the difference between repeat trials and nonrepeat trials is greater during the frequent-repeat-trial blocks. This phenomenon, known as the P(rep) effect (Kovács et al., 2013), suggests that repeated stimuli are processed more efficiently when one expects stimuli to be repeated, thus requiring less neural response. Indeed, Summerfield and colleagues argue that the P(rep) effect indicates that fMRI adaptation can be explained—perhaps entirely—in terms of stimulus prediction (Summerfield et al., 2008; although see Larsson & Smith, 2012).

Although appealing, evidence for this idea remains sparse, and several aspects of the P(rep) effect remain unexplored. Notably, no study has reported a behavioral equivalent of the P(rep) effect, perhaps because previous investigations of the effect have used orthogonal tasks with few target trials. However, when repeated stimuli are more likely, then the increased efficiency in processing repeated stimuli should be reflected not only in reduced fMRI responses, but also in faster or more accurate behavioral responses. The absence of reports of a behavioral P(rep) effect is especially notable given that there already exists a behavioral equivalent for basic fMRI adaptation: repetition priming (e.g. Desimone, 1996; Henson & Rugg, 2003; Schacter & Buckner, 1998). This refers to the faster or more accurate processing of a repeated stimulus compared to a nonrepeated stimulus (e.g. Ellis, Young, & Flude, 1990; Wiggs & Martin, 1998). Both repetition priming and fMRI adaptation have been suggested to reflect the attenuation of neuronal responses (Desimone, 1996; Henson & Rugg, 2003; Schacter & Buckner, 1998), although there is some

evidence that there might be multiple mechanisms that underlie this decrease depending on the length of the repetition interval (Epstein, Parker, & Feiler, 2008; although see Weiner et al., 2010). When behavioral priming and fMRI adaptation are assessed using similar paradigms, there are many similarities in the functional properties of these two effects. For example, both are tolerant to moderate stimulus alterations such as changes of size or retinal position, and both decay at similar rates over time (Desimone, 1996; Wiggs & Martin, 1998). Thus, if the P(rep) effect reflects increased efficiency of processing, as hypothesized by Summerfield and colleagues, then we would expect to see this manifested in both behavioral priming and fMRI adaptation. One aim of the current study was to test this prediction.

Another aspect of the P(rep) effect that remains poorly understood is the nature of the expectations that drive it. In the standard paradigm (Kovács et al., 2013; Summerfield et al., 2008), the first stimulus (S1) on each trial is followed by a second stimulus (S2), which is either the same as the first (rep trial), or different (nonrep trial). The effect of repetition is greater in blocks where repetitions are more common; this could be explained potentially by the fact that S1 is more informative about S2 in those blocks (as S2 is more likely to be the same as S1). However, there is some evidence that other factors, beyond the likelihood of repetition, play a role in mediating the expectations that S1 creates about S2. In particular, the P(rep) effect appears to depend on prior experience with the stimulus class: it is found for faces and upright roman letters, but *not* for nonsense letters in an unfamiliar script (Grotheer & Kovacs, 2014); nor is it found for chairs (Kovács et al., 2013). In the latter two cases, it might be argued that lack of expertise with the stimulus set makes it difficult for the system to develop expectations about S2 from S1. In the extreme case, for example, if the stimuli were drawn from a class for which the participant could not distinguish between stimulus exemplars, then it would be impossible to make any predictions about the identity of S2 based on the identity of S1.

To test whether the amount of information provided by S1 about S2 is in fact central to the P(rep) effect, we use synthetic faces to manipulate the perceptual range of stimuli shown to each participant while keeping the stimulus class (faces) constant. In previous experiments examining the P(rep) effect with faces, stimuli were drawn from a large range of possible faces, including many ethnicities and both genders, but on any given trial the ethnicity and gender of the faces were the same. Thus, S1 was highly informative about S2, which may have encouraged participants to make predictions about S2 based on S1 (which would in turn be modulated by the likelihood of repetitions within a

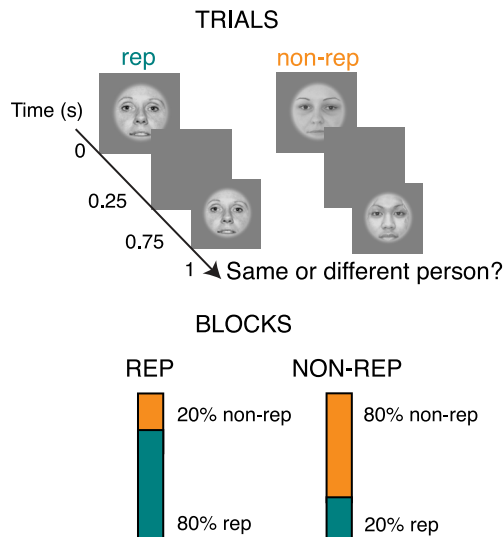


Figure 1. Stimuli and task for Experiment 1. We employed a 2×2 design to study the effect of expectation on repetition priming. Two trial types, repetitions (rep) and nonrepetitions (nonrep) were embedded in either repetition (REP) or non-repetition (NONREP) blocks. 80% of the trials in REP blocks were repetitions, whereas 20% of the trials in NONREP blocks were repetitions. Participants were instructed to respond on each trial whether the two photographs depicted the same or two different individuals as quickly and accurately as possible. The photographs were kindly provided by Mareike Grotheer and Guyla Kovács.

block). In contrast, if faces are drawn from a smaller range such that all stimuli in the experiment are the same ethnicity and gender, then S1 provides less new information, both about itself (since the gender and ethnicity of S1 are already known) and also about S2. Consequently, one might expect the P(rep) effect to be reduced in this case. To ensure that any differences between the range conditions were attributable to differences in expectation rather than within-trial differences in stimulus similarity, we defined the Euclidean distances between S1 and S2 to be the same in our parametric stimulus space for both ranges. In other words, the stimulus difference in a nonrep trial between two Caucasian males in the small range condition was the same as the stimulus difference between two East-Asian faces in the large range condition.

We performed three experiments to examine these ideas. In Experiment 1 we studied a behavioral version of the standard P(rep) paradigm to test whether expectation modulates recognition speed in the same manner as it has been previously shown to modulate fMRI responses. In Experiment 2, we tested the hypothesis that this behavioral P(rep) effect reflects the informativeness of the first stimulus in each trial, by examining the sensitivity of the effect to stimulus range.

In Experiment 3, we examined fMRI responses to the same manipulations used in Experiment 2. To anticipate, our results support the idea that expectation modulates recognition efficiency and reflects the predictive information of one stimulus regarding another. However, we also find that behavioral effects of expectation are not always accompanied by concomitant fMRI effects in visual cortex, even when fMRI adaptation effects are present. Thus, our results are most consistent with the emerging view that expectation and adaptation effects are driven by different underlying mechanisms.

Experiment 1: Behavioral P(rep) effect

The purpose of this experiment was to test whether there is an effect of expectation on recognition behavior. Specifically, we tested whether recognition priming is modulated by the probability of repetition. To facilitate comparison between the current results and earlier studies, we replicated as closely as possible the experiment of Summerfield and colleagues (2008) in which they reported a P(rep) effect on BOLD responses to faces in the fusiform face area (FFA). However, in contrast to this earlier study in which participants monitored for an occasional target (inverted face, or size-deviant face) that only occurred on some trials, here we used a task in which participants made a behavioral response on every trial.

Methods

Participants

Twenty participants (five male, ages 18–51) with normal or corrected-to-normal vision were recruited from the University of Pennsylvania participant pool. Participants gave informed consent and were paid \$10/hour. The experimental protocol (as well as protocols for subsequent experiments reported here) was approved by the university's Institutional Review Board and adhered to the Declaration of Helsinki.

Procedure

Stimuli were cropped full-front grayscale photographs of male and female faces of different ethnicities as used in previous work (Grotheer & Kovács, 2014; Kovács et al., 2013), which were kindly supplied to us by Mareike Grotheer and Guyla Kovács. Stimulus timing within a trial, as well as the block design, was identical to the original fMRI experiment (Figure 1). Specifically, on each trial, two face images differing by




Figure 2. Expectation of repetition probability affects repetition priming. Reaction times in correct trials to a same/different judgment are shown for the two trial types (white = rep, gray = nonrep) and the two block types (left = REP, right = NONREP). Error bars show ± 1 SEM across participants. The P(rep) effect is revealed by the significant trial \times block interaction.

15% in width and height were shown sequentially for 250 ms separated by a 500 ms interstimulus interval. After the offset of the second image, participants responded with a key press to indicate whether the two images depicted the same face or two different faces. Participants were instructed to make their responses as quickly and as accurately as possible. The response was followed by a 1-s inter-trial interval. If the participant failed to respond within a 1.5-s window, the program proceeded to the next trial and appended the missed trial to the end of the run.

The probability of face repetition was manipulated in blocks of 20 trials. In repetition (REP) blocks 80% of trials were repetitions in which the two face images were identical except for their size (rep trials) whereas 20% of trials were nonrepetitions (nonrep trials); in nonrepetition (NONREP) blocks, 20% of trials were repetitions while 80% were nonrepetitions. The two faces in nonrep trials were always of the same ethnicity and gender. Participants completed two runs, each consisting of four blocks of 20 trials (160 trials total).

Results and discussion

Participants were accurate in judging if the two photographs depicted the same individual or not (overall accuracy: 93%). To test the hypothesis that recognition would be facilitated by repetition and modulated by expectation, response times (RT) from the offset of the second stimulus were extracted for the correct trials. Outliers (>3 SD from the mean; 1.5% of correct trials) were excluded from analysis (Magnussen,

Idás, & Myhre, 1998). As can be seen in Figure 2, the predicted effects of repetition priming and modulation by expectation were observed: RTs were faster on rep trials than on nonrep trials, and this difference was greater during REP blocks than during NONREP blocks. These effects were confirmed by a four-way mixed-effects analysis of variance (ANOVA) with fixed factors for trial type (rep vs. nonrep), block type (REP vs. NONREP), and run number (run 1 vs. run 2), and participant as a random factor. The analysis revealed a significant main effect of trial type, $F(1, 77) = 31, p < 0.0001$, and a significant interaction between trial type and block type, $F(1, 77) = 7.9, p = 0.006$. Reaction times were overall faster in the second run, $F(1, 77) = 7.4, p = 0.01$. Furthermore, priming was stronger in the second run, shown by a significant interaction between run and trial type, $F(1, 77) = 4.7, p = 0.03$. There were no other significant effects (all other F s < 0.38 , n.s.).

For completeness, we also examine accuracy as a function of condition, though effects here were expected to be weak as performance was close to ceiling. In this case, the four-way ANOVA revealed no main effects of trial or block type on accuracy, but there was a significant trial \times block type interaction, $F(1, 77) = 7.1, p = 0.01$, reflecting that accuracy was significantly higher in the rep/REP trials compared to the other conditions. That is, participants had a higher accuracy for reporting “same” individual when the face repeated in the context of a block of trials with frequent repetitions, consistent with the P(rep) effect.

Repetition priming is a well-established phenomenon in visual perception (e.g. Ellis et al., 1990; Goshen-Gottstein & Ganel, 2000; Kristjansson & Campana, 2010; Vuilleumier, Henson, Driver, & Dolan, 2002; Wig, Grafton, Demos, & Kelley, 2005), but to our knowledge this is the first time repetition priming is shown to be affected by top-down expectations about repetition probability. Indeed, this is the first demonstration of a behavioral correlate for the fMRI P(rep) effect first shown by Summerfield et al. (2008).

A notable difference between the current study and previous investigations of the P(rep) effect is that we used a same/different task instead of an orthogonal task, and we required participants to make a behavioral response on every trial, rather than just on-target trials. Because participants were explicitly reporting the difference between rep and nonrep trials, they may have been especially aware of this manipulation. However, we think this fact is unlikely to affect the generalizability of our results, as it seems unlikely that participants in the previous study were unaware of the difference between rep and non-rep trials, even though they were not required to report this difference. Moreover, we do not believe that this task would have made participants more aware of the block manipulation, and indeed no participant reported noticing this

manipulation. The same/different task may have led to an overestimate of the magnitude of the priming effect, because priming is confounded by response, and it is known that on comparison tasks reaction times are faster for “same” responses than for “different” responses (“fast-same” effect; Proctor, 1981). However, such an effect could only explain the main effect of trial type, not the P(rep) effect revealed through the trial-block interaction.

Experiment 2: Modulating predictive information by varying stimulus range

Having established that expectation can modulate repetition priming by showing a behavioral equivalent to the fMRI P(rep) effect, we next examined the nature of the expectations that drive this effect. In particular, we tested the hypothesis that the P(rep) effect reflects the information conveyed by the first stimulus about the second, and thus should be modulated by the range of stimuli presented to each participant. To do this, we tested a group of participants with faces drawn from a small range (all Caucasian males) and another group of participants with faces drawn from a large range (different ethnicities and both genders) similar to the range used in previous fMRI P(rep) experiments (e.g. Larsson & Smith, 2012; Summerfield et al., 2008). We reasoned that if the P(rep) effect relates to the predictive information of the stimulus, then it should be greater when stimuli are drawn from a larger range than when they are drawn from a small range. Because this design required us to systematically vary the similarity between faces, we used artificial face images created by FaceGen software rather than face photographs. Importantly, to keep the within-trial difference between S1 and S2 constant for both range conditions, we always used the same Euclidean distance in FaceGen units between S1 and S2 in non-repetition trials. Consequently, only the between-trial differences (i.e., the range) varied between conditions.

Methods

Participants

Experiment 2 consisted of three versions of the basic 2×2 factorial design with different sets of participants in each version. Twenty participants (five male, aged 18–51) from Experiment 1 participated in version 3 during the same session (see next section for description of versions). Another group of 20 participants (eight male, aged 18–54) from the University of Pennsylvania participant pool ran version 2. A subset of both groups

(9 and 11 participants, respectively) ran version 1 (seven male, aged 18–54). We ensured that the two groups did not differ in terms of the effects of interest in version 1 (group did not interact with any of the effects of interest in a three-way ANOVA with trial type, block type, and group as variables (all F s < 2.2 and p s > 0.14).

Procedure

We used the same 2×2 factorial design as in Experiment 1 in three versions of the experiment (Figure 3). In all three versions, trials of two different types (rep/nonrep, see what follows) were shown within REP or NONREP blocks. The trial timing, block structure, and run structure were unchanged.

The three versions of the experiment differed in: (a) the range from which the face stimuli were drawn, and (b) whether the two stimuli on repeat trials were exactly the same or approximately the same. In version 1 (large range, exact repetitions), the stimuli were sampled from a large range of faces that included many different ethnicities and both genders. On rep trials, the exact same face was shown twice (but at different sizes), and on nonrep trials two faces were shown that were separated by a large jump in the parametric stimulus space (see what follows) and thus had the appearance of two different individuals. Thus, this version was similar to Experiment 1. In version 2 (small range, exact repetitions), the stimuli were sampled from a small range of faces giving them all the appearance of Caucasian males. Once again, the exact same face was shown on rep trials, and the faces shown on non-rep trials had the appearance of different individuals. Version 3 (small range, approximate repetitions) was similar to version 2, though in this case the rep trials were not exact repetitions but small transitions between faces [$5 SD$, approximately 1.5 just-noticeable-differences (JND)]. Importantly, the difference in face space between the first and second stimulus in non-rep trials was the same in all three versions. Thus, variation of the P(rep) effect across the versions could only be attributed to differences in between-trial range, not within-trial stimulus differences.

In all versions, participants were instructed to indicate whether the two consecutively presented faces depicted the same individual. Because in version 3, faces were never exact repetitions even on rep trials, these participants were told prior to the experiment that faces might not repeat exactly, but that they should try to judge identity across these small variations in geometry. For instance, if the two faces appeared to have slightly different expressions but otherwise appeared the same, participants should respond “same individual.”




Figure 3. Stimuli and task for Experiment 2. (a) We manipulated the informativeness of S1 about S2 by sampling the faces across trials either from a small face range (left panel) or from a large face range (right panel). In the case of a small range, S1 does not distinguish strongly among the possibilities for the second stimulus even with high repetition probability, because all stimuli fall in the same general stimulus class (in this case, Caucasian males). The possible transitions in this space are all similar. But if stimulus range is large across trials, S1 distinguishes more strongly among alternatives for S2, because intertrial stimulus differences are larger. For instance, if S1 is an Asian male, this constrains the stimulus space for S2 to the same category, which may cause the participant to be more sensitive to the repetition probability manipulation between blocks. (b) We ran three versions of the experiment to investigate the dependence of the expectation effect on the overall stimulus range (large vs. small) and the similarity of the faces in repetition trials (approximate vs. exact repetitions). The trial \times block factorial design within each version was identical to the one in Experiment 1.

Stimuli

We used FaceGen (full version, Singular Inversions, Inc., Toronto, ON, Canada) to generate face stimuli with various degrees of similarity. FaceGen provides a parametric description of faces in a space of 130 principal components that control face geometry, color, and contrast. FaceGen allows for the generation of arbitrary faces by setting the 130 components to arbitrary values within a given range. We used this feature to generate faces with specified distances (i.e., differences) from the origin and from other faces.

FaceGen uses units of standard deviation (SD), which indicate how much of the variance in the PCA solution each dimension explains. We will use these inherent FaceGen units to describe the magnitude of our stimulus modulations.

The faces created by FaceGen vary continuously in appearance. Thus, before selecting stimuli for the experiment, it was first necessary to understand which stimulus differences would be either interpreted as being different individuals and which would be interpreted as being the same individual. Normative

perceptual data are not available for parametric variation of faces on the arbitrary FaceGen axes, so we conducted a behavioral norming study to determine this. Our goal was to establish the stimulus differences between faces that equated to (a) a just noticeable difference (JND) and (b) a sufficiently large difference that participants were near a performance ceiling in a same/different identity task.

First, we determined the average JND by measuring discrimination thresholds in an odd-one-out, three-interval forced choice task (Giesel, Hansen, Gegenfurtner, 2009; Roseboom, Linares, & Nishida, 2015). Four reference faces were selected by eye within a radius of 5 *SD* from the origin. For each reference, we generated a large number of test faces on a randomly oriented axis centered on the reference. The endpoints of the continuum were ± 6 *SD* units away from the reference. Test faces were constrained to be within 10 *SD* from the origin of the space to avoid visually bizarre shape distortions of the synthetic faces that occur for extreme values on the stimulus axes. Each trial had three intervals, in which the same reference face was displayed in two intervals and a test face in the third interval. The stimuli were presented at fixation in a randomized order. Each face was displayed for 1 s, followed by a 100 ms interstimulus interval. The participant indicated on the keyboard which face was different from the other two. Trials for the four references were interleaved. The test face on each trial was selected according to one of two adaptive staircases, which converged on the 50% and 79% percentiles of the psychometric function, respectively. In other words, when the participants started to perform close to chance, the differences between the reference and test were made larger, and when the task became easy, the differences were made smaller. The JNDs for each reference were defined as the average over the reversal points of the two staircases, and the 65% percentile of a cumulative Gaussian fitted to the proportion-different responses. One JND approximately corresponded to a 70% discrimination threshold, or 3 *SD* in face space units (Supplementary Figure 1a).

Next, to determine a distance in stimulus space that would be interpreted as different individuals, we determined with three new participants the average stimulus difference at which reaction times to a same/different individual question plateaued. We reasoned that this would indicate that participants were able to discriminate individuals without difficulty. We selected two reference faces and 15 comparison stimuli for each reference. The comparison stimuli were chosen with even spacing from a line between the reference and an extreme comparison face 12 *SD* away, selected from a randomized direction away from the reference. Participants saw each reference with each comparison once

(105 pairings per reference). We measured reaction times in a go/no-go task where participants were asked to press a key when the two images depicted different individuals and withdraw a response if they depicted the same individual. Supplementary Figure 1b shows the average reaction times for all stimulus differences for the three participants. The reaction time data were remarkably flat, although extreme responses fell off around 10 *SD*. We selected 11 *SD* as our largest step size as a compromise between the reaction time data and the limits of the synthetic face space.

Once we had determined stimulus differences corresponding to the perception of the same or different individual, we then selected stimuli for use in the main experiment. We employed some constraints in stimulus generation to avoid geometric distortions in the synthetic face space. To keep skin color, contrast, and texture constant within ethnicity groups, we only varied the dimensions that control face geometry (2–42). We also kept the first two dimensions (roughly, face width and length) constant throughout, as these dimensions masked changes on the other dimensions by virtue of being conspicuous.

Stimuli in version 1 (large range) varied in ethnicity and gender across trials. We first found the coordinates for prototypical faces in each ethnicity group (Afro-Caribbean, East-Asian, South-Asian, European) and gender. Next we generated the base stimuli for each trial as follows. The first face in a trial was selected randomly from the 40 usable dimensions (specifically, we randomly set the 40 coordinates to values within our predefined limits to avoid distortions). The next face was chosen randomly a given number of *SD* units away from the first face. The number of units was zero in case of a perfect repetition; 5 in case of a small change, and 11 in case of a large change. We further made sure that each face was at least three units (one JND) away from any other face in the same run, excepting repetition trials. To then generate a face with a particular ethnicity and gender, we recentered the base face according to the prototypical face in that ethnicity/gender category via vector summation. This allowed us to separately control the distance between any two faces in a trial, and the ethnicity and gender across trials. The stimuli for versions 2 (small range, exact repetitions) and 3 (small range, approximate repetitions) were generated similarly but without the shift to a new ethnicity/gender on every trial, resulting in Caucasian-appearing male faces.

Results and discussion

To test for priming and expectation effects, response times from the offset of the second stimulus were extracted for correct trials. In versions 1 and 2, trials

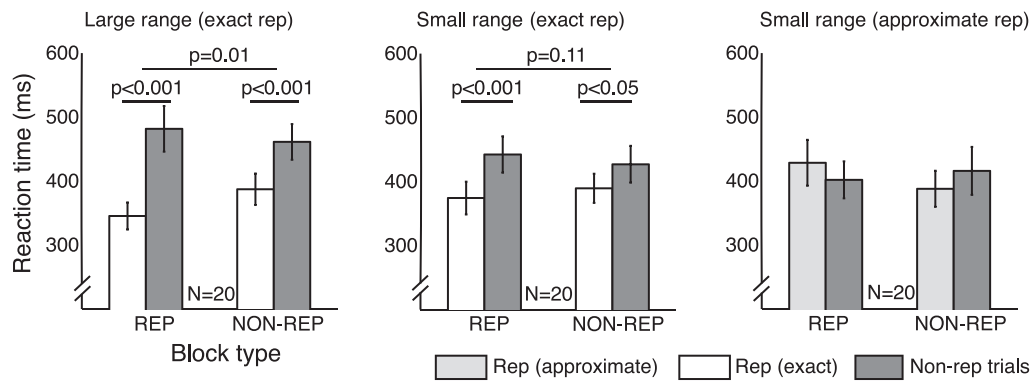


Figure 4. Effects of priming and expectation on reaction times for the three experimental versions. Trial type is indicated with bar color, and the two sets of bars show the data for the two block types. Left: large range (exact rep); middle: small range (exact rep); right: small range (approximate rep). The P(rep) effect is significant when the range is large, but not when the range is small, as indicated by the significant trial \times block interaction in the large range (exact rep) condition. Error bars show $\pm 1SEM$ across participants.

were scored correct if participants responded “same” to exact repetitions and “different” to nonrepetitions. In version 3, trials were scored correct if participants responded “same” for the approximate repetitions and “different” for the nonrepetitions. Response time outliers ($>3 SD$ from the mean, 2% of correct trials) were excluded from analysis. Accuracy was 82% for version 1 (large range, exact rep), 90% for version 2 (small range, exact rep) and 75% for version 3 (small range, approximate rep). This difference between versions was significant in a four-way ANOVA with trial type, block type, and version as fixed effects and participant as a random effect, $F(2, 375) = 23.4$, $p < 0.0001$.

As predicted, the RT effects differed notably across the three versions of the experiment (Figure 4). In the large range (exact rep) version, which was most similar to Experiment 1, we replicated the effects of repetition priming and expectation: RTs on correct trials were faster when stimuli were repeated within a trial (repetition priming) and this RT advantage was greater on REP blocks compared with NONREP blocks [P(rep) effect]. These effects were confirmed by a four-way ANOVA with factors for trial type (rep vs. nonrep), block type (REP vs. NONREP) and run number (1 vs. 2) and participant as random factor. This analysis revealed a significant main effect of trial type, $F(1, 76) = 36$, $p < 0.001$, and a significant interaction between trial type and block type, $F(1, 76) = 7.0$, $p = 0.01$. Reaction times in the second run were, as expected, marginally faster than the first, $F(1, 76) = 3.6$, $p = 0.07$. No other effects were significant ($F_s < 3.6$, n.s.). In the small range (exact rep) version, on the other hand, the repetition priming effects was significant [main effect of trial type $F(1, 77) = 8.2$, $p = 0.005$], but the P(rep) effect was not, though the trend was in the right direction (interaction of trial and block type, $F(1, 77) = 2.7$, $p = 0.1$). Finally, in the small range

(approximate rep) version, neither response priming nor P(rep) effects were observed [main effect of trial type $F = 0$, $p = 0.95$; trial \times block type interaction $F = 5.9$, $p = 0.02$, significant but in the wrong direction].

Thus, the results confirm the hypothesis that P(rep) effects would be reduced when stimuli were drawn from a narrower range, thus making repetitions less informative. To confirm the difference between the three versions of the experiment, we tested the effect of version in a mixed-effects ANOVA with trial type and block type as fixed effects, and participant and version as random effects with effects up to third level included. The effect of trial was significant, $F(1, 277) = 22.7$, $p = 0.001$, as was the interaction between trial and version, $F(2, 277) = 5.8$, $p = 0.007$, and the three-way interaction between trial type, block type, and version, $F(2, 277) = 5.1$, $p = 0.007$. In other words, the strength of the P(rep) effect depended on the experimental version.

These results demonstrate that the modulation of priming by expectation is affected not only by the probability of repetition, but also by the range of stimuli that one expects to see. This supports the idea that the P(rep) effect is related to the predictive information provided by the first stimulus about the second. When stimuli are varied, then the first stimulus in a trial sharply reduces the possible appearance of the second stimulus, and thus is highly informative. In contrast, when stimuli are less varied, then knowing the first stimulus provides less information, because one already knows that all stimuli are drawn from a small set of similar items. The P(rep) effect is greater in the first case than the second.

Insofar as behavioral priming is a proxy for fMRI adaptation, these results suggest that the modulatory effect of expectation on fMRI adaptation might only emerge when there is enough stimulus range in the stimulus set. In the following fMRI experiment, we test this hypothesis.

Experiment 3: Expectation and fMRI adaptation

The previous two experiments established a behavioral equivalent for the P(rep) effect and showed that it is modulated by the predictive information conveyed by the stimuli. Here, we turn to fMRI, to see if we can observe similar effects in the neural response to faces. Our primary goal was to determine whether expectation effects are modulated by the stimulus range, as we established for behavior. To this end, we used the same parametric face space as in Experiment 2, with a similar design.

Methods

Participants

Thirty-six participants with normal or corrected-to-normal vision (ages 19–30, 17 males) were recruited from the University of Pennsylvania participant pool. These were divided into three groups of 12, and each group participated in one version of the experiment. The participants signed informed consent and were paid \$30. The experimental protocol was approved by the university IRB and adhered to the declaration of Helsinki.

Stimuli and procedure

Functional MRI data were obtained while participants performed a modified variant of Experiment 2, with some changes in timing and task. These modifications were designed to match our experiment as closely as possible with the fMRI paradigms used previously by Kovács et al. (2013) and Summerfield et al. (2008). Briefly, participants viewed artificially generated faces from the FaceGen set, which were shown in repetition (rep) or nonrepetition (nonrep) trials, which were in turn shown in repetition (REP) or nonrepetition (NONREP) blocks. As in Experiment 2, there were three versions of the experiment: (a) large range with exact repeats, (b) small range with exact repeats, (c) small range with approximate repeats. Each participant performed one of these versions.

On each trial, a face image was shown for 250 ms, followed by a 500 ms interstimulus interval, and then a second face image for 250 ms. Intertrial interval varied between 2 and 4 s (mean = 3 s). Experimental trials in which the two stimuli differed by 15% in size (as in Experiment 2) were interspersed with target trials in which the two stimuli differed by 60% in size. To ensure vigilance, participants were asked to press a button for target trials and withhold response for nontarget trials. Target trials constituted 20% of all trials and were evenly dispersed among the trial and block types.

Average performance on the monitoring task was 95%, 93%, and 96% in the three experimental versions. Runs below 87.5% (four misses) were excluded from analysis; participants that had more than one run below this criterion were excluded completely from further analysis. Based on this criterion, we excluded 2/36 runs in version 1, 5/36 runs in version 2, and the complete dataset from one participant in version 2. To reach our planned number of participants in version 2, data from a 13th participant were acquired.

As before, during REP blocks 80% of the experimental (i.e., nontarget) trials were repeat trials and during NONREP blocks 20% of the experimental trials were repeat trials. The first two trials in each block were constrained to be from the higher-probability category (for instance, the first two trials in a REP block were repetition trials). Each block consisted of 20 trials (four target and 16 experimental). Twelve REP and 12 NONREP blocks were interleaved in ABAB order, and split over three scan runs. Text announcing “New block” appeared briefly between each block. Each scan run consisted of 160 trials and took 11 min.

After the three experimental runs, participants completed two functional localizer runs, each comprising 20 blocks of faces, scenes, objects, and scrambled objects interleaved with fixation blocks. Within each image block, 15 images were presented for 490 ms with a 490-ms inter-stimulus interval. Participants performed a one-back task on image repetition. Each localizer run took 5.25 min.

fMRI acquisition and analysis

Data acquisition. Scanning was performed at the Hospital of the University of Pennsylvania on a 3T Siemens Trio scanner (Siemens Medical Solutions USA, Inc., Malvern, PA) equipped with a 32-channel head coil. Stimuli were displayed on a Sanyo PLC XT35 LCD projector (SANYO Electric Co., Ltd., Osaka, Japan) via a mirror mounted on the head coil. The viewing area of the display was 50.5×38 cm or $23 \times 17^\circ$ of visual angle. Stimulus size was approximately $6 \times 6^\circ$ of visual angle. High-resolution T1-weighted images for anatomical localization were acquired using a three-dimensional magnetization-prepared rapid acquisition gradient echo pulse sequence [repetition time (TR), 1620 ms; echo time (TE), 3.09 ms; inversion time, 950 ms; voxel size, $1 \times 1 \times 1$ mm; matrix size, $192 \times 256 \times 160$]. T2*-weighted images sensitive to blood oxygenation level-dependent contrasts were acquired using a gradient echo echoplanar pulse sequence (TR, 3000 ms; TE, 30 ms; voxel size, $3 \times 3 \times 3$ mm; matrix size, $64 \times 64 \times 44$).

Data preprocessing. Functional images were corrected for differences in slice timing by resampling slices in time to match the first slice of each volume. Images were then realigned to the first volume of the first scan run.

Nuisance variables for motion were computed with the Artifact Detection Tool (http://www.nitrc.org/projects/artifact_detect/) and entered in the general linear model (see what follows). Data were smoothed with an 8 mm full-width at half-maximum Gaussian filter.

Functional regions of interest. Data from the functional localizer scans were used to identify the left and right fusiform face areas (FFA). BOLD responses to the different image types (faces, houses, objects, scrambled objects, fixation blocks) were modeled with a general linear model with boxcar regressors for the five event types and nuisance regressors for head motion in FSL (Jenkinson, Beckmann, Behrens, Woolrich, & Smith, 2012). The right and left FFAs were defined for each participant individually using the face > object contrast and a group-based anatomical constraint of face-selective activation derived from a large number (42) of localizer participants in our laboratory (Julian, Fedorenko, Webster, & Kanwisher, 2012). The t-maps from the faces>objects contrast were averaged over two independent localizer runs and thresholded at $p < 0.0001$. Regions of interest (ROI) for each participant were defined as the conjunction of the thresholded map with the group-based ROIs, registered to each participant's own space. For random-effects whole brain analyses, functional data from each participant were registered to the Montreal Neurological Institute (MNI) 2-mm standard.

fMRI adaptation. We used the methods from Summerfield et al. (2008) to process the functional MRI data after preprocessing. Analyses were run in FSL and included high-pass filters that removed low temporal frequencies (Gaussian-weighted least-squares straight line fitting, with $\sigma=50.0s$) as well as nuisance regressors for motion parameters and outlier volumes. Briefly, the functional MRI data for each participant and run were modelled by GLMs with six regressors: one regressor for each combination of trials and blocks (rep/REP, rep/NON-REP, non-rep/REP, non-rep/NON-REP), a regressor for the “new block” indicator, and a regressor for target trials. The regressors were rectangular step functions convolved with the double-gamma hemodynamic response function in FSL. Each face pair was treated as a composite event. Beta weights for each regressor were calculated for each voxel in the FFA and then averaged over all voxels within the ROI. As there were no significant differences between the left and right hemispheres in terms of adaptation, data were averaged over the hemispheres, as in previous studies on the P(rep) effect.

Results and discussion

Figure 5a shows the average beta values in the 2×2 factorial design extracted from the functionally defined

FFAs (shown in Figure 5b) for the three different versions of the experiment. There was a significant main effect of trial type, which tests for fMRI adaptation, in two of the three versions in a mixed-effects ANOVA with trial type and block type as fixed effects and participant as random effect [version 1—large range, exact rep: $F(1, 22) = 5.2, p = 0.029$; version 3—small range, approximate rep: $F(1, 22) = 8.6, p = 0.006$]; there was no significant fMRI adaptation in version 2 [small range, exact rep; $F(1, 22) = 1, p = 0.32$].

The P(rep) effect was characterized by Summerfield et al. (2008) as stronger fMRI adaptation when repetitions are more frequent (i.e., during REP blocks). We tested this in the fMRI data by seeing if there was an interaction between trial type and block type in the FFA. No trial \times block type interaction in the predicted direction was found in any of the three versions [for version 1 and 3, $F(1, 22) < 0.03, ps > 0.86$; for version 2, $F(1, 22) = 3.3, p = 0.08$ but note that the interaction went in the nonpredicted direction].

The absence of a P(rep) effect in any of the three versions was surprising to us. We considered the possibility that this failure to replicate the results of Summerfield et al. (2008) might be due to minor differences in the anatomical localization of the FFA by extracting beta values from a 6-mm sphere centered on the average FFA from Summerfield et al. (2008) (Figure 5b). Figure 5c shows the average beta values for this ROI in version 1, where the interaction was expected to be strongest based on the behavioral data, and the fact that this condition most closely resembles Summerfield et al.'s original experiment. There was a main effect of trial type, $F(1, 22) = 5.8, p = 0.02$, but no trial \times block interaction, $F(1, 22) = 0, p = 0.97$. The P(rep) effect in this ROI was also nonsignificant in the other two experiment versions (F values $< 0.44, p$ values > 0.5).

Figure 6 summarizes these results by showing the main effect of faces (faces > blank contrast), the strength of adaptation, and the trial \times block interaction in FFA. Contrary to the original results of Summerfield et al. (2008) and in distinction to our behavioral results, there was no trial type \times block type interaction in the FFA in any version, although there were robust responses to the synthetic face stimuli, and robust adaptation in two of the three versions. (As noted already, the trend toward an interaction in version 2 was in the wrong direction.)

Thus, despite the presence of a behavioral effect of repetition expectation, we were unable to replicate Summerfield et al.'s (2008) finding of a P(rep) effect in the fMRI response of the FFA. We therefore considered the possibility that a neural correlate of the trial type \times block type interaction might be present at some other cortical site. We conducted whole-brain, exploratory analyses to identify voxels exhibiting a trial type

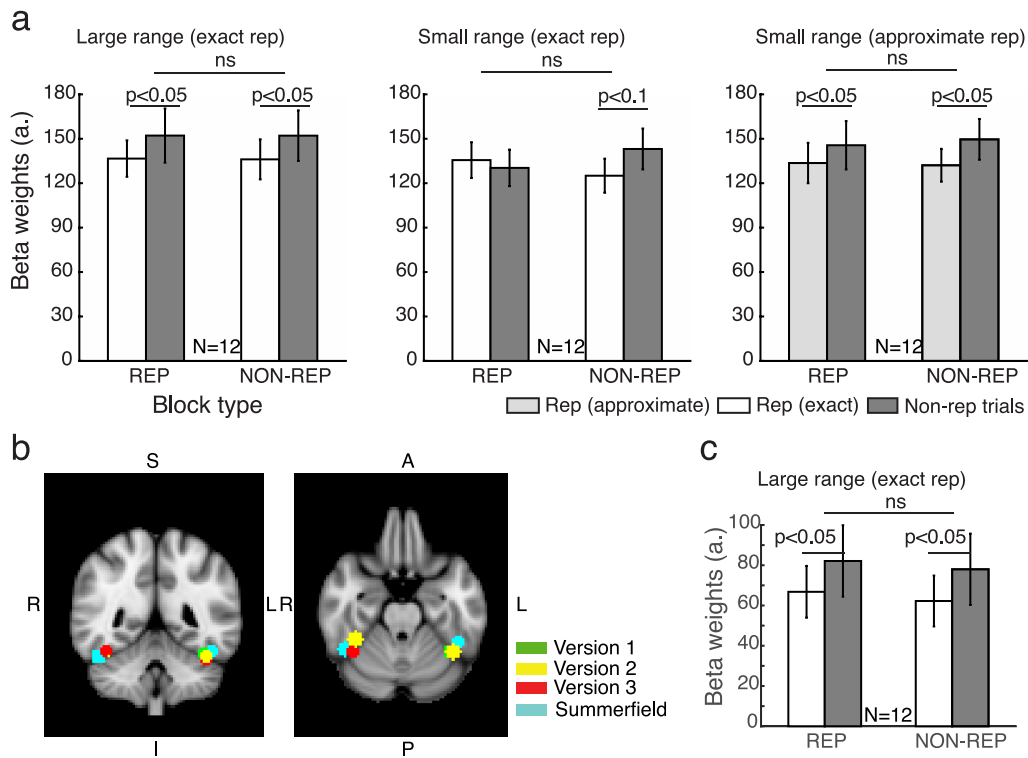


Figure 5. ROI analyses. (a) Average beta values extracted from the individually defined FFAs for the 2×2 factorial design in the three versions of Experiment 3. Trial type is indicated with bar color, and the two sets of bars show the data for the two block types. Left: large range (exact rep); middle: small range (exact rep); right: small range (approximate rep). The P(rep) effect is not significant in any of the three versions. Error bars show $\pm 1SEM$ across participants. (b) FFAs were defined in each experiment with thresholded activation in a faces > objects contrast derived from an independent functional localizer, combined with a group parcel. Spheres (6 mm) are shown around the centroids in the three experimental conditions, as well as from the Summerfield et al. (2008) study. This is to illustrate the average location of the FFAs; the ROI analyses were done on individually defined FFAs. The MNI coordinates of the average centroids were [L: $-38 -52 -22$, R: $40 -52 -20$], [L: $-40 -52 -24$, R: $38 -42 -24$], and [L: $-40 -52 -26$, R: $40 -52 -20$] for Version 1, 2, and 3. (c) Average beta values in the large range experiment extracted from a 6-mm sphere centered on the FFA ROI in Summerfield et al. (2008) (MNI coordinates [L: $-44 -44 -20$, R: $46 -50 -24$]).

\times block type interaction consistent with the P(rep) effect (i.e., greater adaptation in REP blocks than in NONREP blocks). We used the FSL algorithm *randomise* to generate statistical maps with a voxel-wise corrected threshold of $p < 0.05$ derived from the null

distribution of the maximum voxel-wise t statistic (Nichols & Holmes, 2001; Winkler, Ridgway, Webster, Smith, Nichols, 2014).

At a corrected threshold of $p < 0.05$, no ventral occipitotemporal voxels showed the trial \times block

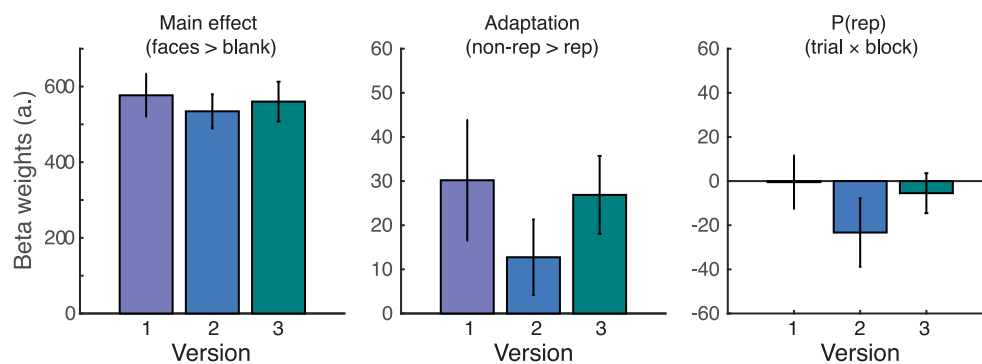


Figure 6. Main effect of seeing a face (left), main effect of trial type (fMRI adaptation, middle), and the trial \times block type interaction [P(rep) effect, right] in the FFA are shown averaged across 12 participants in each experiment. Although adaptation is significant in two out of three versions, the P(rep) effect is not significant in any version. Error bars show $\pm 1 SEM$.




Figure 7. Whole-brain analysis. A whole-brain analysis revealed a frontal region that showed a significant trial \times block type interaction in the large range experiment. Highlighted voxels show the interaction thresholded between $p < 0.001$ uncorrected and $p < 0.05$ corrected for multiple comparisons. The peak voxel was at MNI [-12 54 40].

interaction in any experiment, but a frontal region close to the left anterior paracingulate showed the interaction in the large range experiment (peak voxel at MNI [-12 54 40]; Figure 7). We further generated statistical maps for the comparison between the large-range experiment and the two small-range experiments. The difference between the large-range experiment and the small range, approximate repetitions experiment was significant at the corrected $p < 0.05$ threshold in the same frontal region (peak voxel at MNI [-10 56 40]). The comparison between the large range to the small range, exact repetition experiment did not reach significance at the corrected threshold, but was significant in the same frontal voxels at an uncorrected threshold of $p < 0.01$.

In sum, we found fMRI adaptation in functionally localized FFA, but no accompanying modulation of fMRI adaptation by repetition probability. Thus, we failed to replicate previous findings for face photographs in the FFA (Kovács et al., 2013; Larsson & Smith, 2012; Summerfield et al., 2008). However, a post-hoc whole-brain analysis identified a region in medial prefrontal cortex (specifically, anterior paracingulate) that showed an interaction between trial type and block type in the large range experiment, which had the largest behavioral P(rep) effect in Experiment 2. Medial prefrontal cortex has been previously implicated in the processing of familiar faces (Gobbini, Leibenluft, Santiago, & Haxby, 2004; Todorov, Gobbini, Evans, & Haxby, 2007), and more interestingly, in the detection of violated expectations (Wacongne et al., 2011), and in general error monitoring in cognitive tasks (Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis, 2004). Because our whole-brain analysis was exploratory, this result needs to be verified in future experiments that target a priori defined frontal ROIs, but this result, together with previous literature, points to a possible frontal component in the P(rep) effect.

General discussion

In this study, we report three main results. First, in Experiment 1, we demonstrated that expectations have consequences for recognition behavior in addition to their already-established consequences for fMRI response: When repetitions were more likely, participants exhibited greater repetition priming than when they were less likely. Second, in Experiment 2, we demonstrated that this effect relates to the information conveyed by S1 about S2: When stimuli were drawn from a large set of possible stimuli, thus making the first item on a repetition trial highly informative about the second item, the behavioral P(rep) effect was larger, whereas when stimuli were drawn from a more restricted set of stimuli, making the first item less informative, the behavioral P(rep) effect was smaller. Third, in Experiment 3, we found some evidence for a P(rep) effect in the fMRI response in prefrontal cortex, but we failed to replicate previous findings of a P(rep) effect in the fMRI response in the FFA even though fMRI adaptation was found in this region. These results have important implications for our understanding of the P(rep) effect and its relationship to simple effects of repetition (i.e., repetition priming and fMRI adaptation).

Most generally, these results provide support for the idea that top-down expectations are important modulators of both recognition behavior and neural response to perceptual stimuli. Although the P(rep) effect has been demonstrated for neural-related signals such as BOLD response and event-related potentials, to our knowledge it has not been previously demonstrated for recognition behavior. Here, we find that behavioral repetition priming can be modulated by expectation about repetition probability: Priming is greater when repetitions are more frequent, and thus more expected. While the overall priming effect could be explained in part by a “fast-same” effect (Proctor, 1981), this could

not account for the P(rep) effect revealed by the interaction of trial type and block type.

Our results also provide some insight into the nature of the expectations that drive the P(rep) effect. When faces were drawn from many different ethnicities and both genders (Experiment 1, and the large range conditions of Experiment 2), then we observed a significant P(rep) effect, but when the faces were all Caucasian males (the small range conditions of Experiment 2), we did not observe such an effect, even though repetition priming was still robust. We hypothesize that this modulation of the P(rep) effect by range reflects the difference in information conveyed by the first item of a repetition in the two contexts: When stimuli are drawn from a broad set of possible stimuli, then knowing the first item of a repetition provides a good deal of information, because it allows one to cut down this universe of possible stimuli to a much smaller set. For example, in the large range condition, the ethnicity and gender of the first stimulus constrains the second stimulus to have the same characteristics. In contrast, if all the stimuli are Caucasian males, then seeing the first stimulus in a trial does not provide any new information about the category of the second stimulus, as one already knows that it will be a Caucasian male. This may make the visual system more likely to use S1 to make predictions about S2 in the first case than in the second case.

Previous work has interpreted the P(rep) effect in terms of predictive coding (Summerfield et al., 2008; Todorovic & de Lange, 2012; Wacongne et al., 2011; also see Rao & Ballard, 1999). In this view, repetition suppression occurs because the smaller prediction error for a repeated stimulus compared to a novel stimulus is reflected as a smaller neural or BOLD signal, and the P(rep) effect reflects stronger predictions in REP blocks compared with NONREP blocks. The present results are generally consistent with the predictive coding framework insofar that they show that the extent to which S1 is informative about S2 affects the strength of the P(rep) effect. However, at present, we do not have a precise understanding of how the range manipulation affects the prediction error on viewing S2.

Our findings have the potential to illuminate previous attempts to replicate the P(rep) effect, which have obtained mixed results. Many of the previous studies that have reported positive results have used face photographs drawn from a large stimulus set, where all stimuli are highly distinct, and thus the first stimulus is informative about the second stimulus in repetition blocks. In contrast, studies that have failed to replicate the effect have used nonface stimuli (Grotheer & Kovács, 2014; Kaliukhovich & Vogels, 2011; Kovács et al., 2013). In these cases, seeing one stimulus may create only an imprecise prediction about the next stimulus on the grounds that we are less

specialized to process these stimuli. Supporting this view, Grotheer and Kovács (2014) recently found expectation modulation for familiar but not for unfamiliar letters. Stimulus familiarity might conceivably be important for the accumulation of expectations because people tend to process familiar stimuli with more accuracy and resolution than unfamiliar stimuli (e.g., Ashby & Maddox, 2005; Goldstone, 1998). Thus, a familiar stimulus may be more informative about the precise identity of an upcoming stimulus than an unfamiliar stimulus.

Despite these reliable behavioral effects, we were not able to replicate previous studies that found modulation of fMRI adaptation by expectation in the FFA. What accounts for this failure? It is unlikely to be due solely to the fact that our stimuli were synthetic, as Summerfield et al. (2011) found a P(rep) effect with synthetic faces in an EEG paradigm. However, their synthetic faces had hair and were not parametrically sampled from the face space, rendering their stimuli more similar to photographs, and with larger differences between stimuli than in our experiment. In addition, it is possible that our synthetic faces attracted slightly less attention than real faces — animate images tend to elicit more attention than inanimate images (Mack, Pappas, Silverman, & Gay, 2002; New, Cosmides, & Tooby, 2007), and attention seems to be important for the P(rep) effect (Larsson & Smith, 2012). This might also explain why we found stronger evidence for expectation modulation with behavioral priming compared with the fMRI ROI analysis. In the behavioral study, participants made judgments about identity, which probably drew their attention toward face features, whereas doing the orthogonal task in the fMRI experiment may have directed their attention away from these features.

Whatever the reason for the absence of the P(rep) effect in the FFA, it is important to note that this was obtained in the setting of reliable fMRI adaptation, which was observed in 2 of the 3 versions of Experiment 3. Furthermore, in an exploratory analysis, we found a frontal region that showed the P(rep) effect in the large range experiment, which tentatively suggests a role for frontomedial cortex in tracking predictions. Notably, this region only exhibited a P(rep) effect for the large range version, for which the behavioral P(rep) effect also was strongest. Thus, we were able to identify a brain locus correlate of the P(rep) effect, but it was not anatomically coterminous with the brain locus of fMRI adaptation.

These results are most consistent with the view that the P(rep) effect on fMRI response and fMRI adaptation are two independent effects. Although they are sometimes observed in the same brain region (the FFA), this concurrence of anatomical locus is not mandatory. This conclusion is in line with recent

findings that repetition suppression and expectation suppression follow different time courses in human auditory EEG (Todorovic & de Lange, 2012) and in monkey inferotemporal cortex (Bell, Summerfield, Morin, Malecek, & Ungerleider, 2016), and with the failure to show modulation of fMRI adaptation/repetition suppression by expectation in humans (Grotheer & Kovács, 2014; Kovács et al., 2013) or monkeys (Kaliukhovich, & Vogels, 2011). They are also in line with Larsson and Smith's (2012) finding that the P(rep) effect was eliminated when attention was diverted away from the stimulus whereas the fMRI adaptation effect was maintained. Although their experiment examined long adaptation durations rather than short adaptation durations, their results are generally consistent with ours insofar as they demonstrate that adaptation and expectation have dissociable effects on the fMRI signal. We conclude that fMRI adaptation cannot be entirely explained in terms of predictive coding, but rather reflects stimulus-driven similarities between stimuli, as traditional accounts would maintain.

Keywords: face perception, fMRI adaptation, repetition suppression, response priming, predictive coding

Acknowledgments

This project was supported by NIH R21-EY-022751-02. Many thanks to Jack Ryan for help with the fMRI data collection. We thank Guyla Kovács, Marcelo G. Mattar, Andrew Bock, and Steven Marchette for helpful discussions and Anne-Marieke Schiffer and Toni P. Saarela for comments on a previous version of the manuscript.

Commercial relationships: none.
Corresponding author: Maria Olkkonen.
Email: maria.olkkonen@durham.ac.uk
Address: Department of Psychology, Durham University, Durham, UK.

References

- Aguirre, G. K. (2007). Continuous carry-over designs for fMRI. *NeuroImage*, *35*, 1480–1494.
- Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology*, *56*(1), 149–178.
- Ashida, H., Lingnau, A., Wall, M. B., & Smith, A. T. (2007). fMRI adaptation reveals separate mechanisms for first-order and second-order motion. *Journal of Neurophysiology*, *97*(2), 1319–1325.
- Barron, H. C., Garvert, M. M., & Behrens, T. E. J. (2016). Repetition suppression: A means to index neural representations using BOLD? *Philosophical Transactions of the Royal Society B*, *371*, 20150.
- Bell, A. H., Summerfield, C., Morin, E. L., Malecek, N. J., & Ungerleider, L. G. (2016). Encoding of stimulus probability in macaque inferior temporal cortex. *Current Biology*, *26*(17), 2280–2290.
- Carandini, M., & Ferster, D. (1997). A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science*, *276*, 949–952.
- Desimone, R. (1996). Neural mechanisms for visual memory and their role in attention. *Proceedings of the National Academy of Sciences, USA*, *93*(24), 13494–13499.
- Dobbins, I. G., Schnyer, D. M., Verfaellie, M., & Schacter, D. L. (2004). Cortical activity reductions during repetition priming can result from rapid response learning. *Nature*, *428*, 316–319.
- Ellis, A. W., Young, A. W., & Flude, B. M. (1990). Repetition priming and face processing: Priming occurs within the system that responds to the identity of a face. *The Quarterly Journal of Experimental Psychology Section A*, *42*(3), 495–512.
- Epstein, R. A., & Morgan, L. K. (2012). Neural responses to visual scenes reveals inconsistencies between fMRI adaptation and multivoxel pattern analysis. *Neuropsychologia*, *50*(4), 530–543.
- Epstein, R. A., Parker, W. E., & Feiler, A. M. (2008). Two kinds of fMRI repetition suppression? Evidence for dissociable neural mechanisms. *Journal of Neurophysiology*, *99*(6), 2877–2886, doi:10.1152/jn.90376.200.
- Fang, F., Murray, S. O., Kersten, D., & He, S. (2005). Orientation-tuned fMRI adaptation in human visual cortex. *Journal of Neurophysiology*, *97*, 4188–4195.
- Giesel, M., Hansen, T., & Gegenfurtner, K. R. (2009). The discrimination of chromatic textures. *Journal of Vision*, *9*(9):11, 1–28, doi:10.1167/9.9.11. [PubMed] [Article]
- Gobbini, M. I., Leibenluft, E., Santiago, N., & Haxby, J. V. (2004). Social and emotional attachment in the neural representation of faces. *NeuroImage*, *22*(4), 1628–1635.
- Goldstone, R. L. (1998). Perceptual learning. *Annual Review of Psychology*, *49*, 585–612.
- Goshen-Gottstein, Y., & Ganel, T. (2000). Repetition priming for familiar and unfamiliar faces in a sex-judgment task: Evidence for a common route for

- the processing of sex and identity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(5), 1198–1214.
- Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: Neural models of stimulus-specific effects. *Trends in Cognitive Sciences*, 10(1), 14–23.
- Grill-Spector, K., Kushnir, T., Edelman, S., Avidan, G., Itzhak, Y., & Malach, R. (1999). Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron*, 24, 187–203.
- Grill-Spector, K., & Malach, R. (2001). fMR-adaptation: A tool for studying the functional properties of human cortical neurons. *Acta Psychologica*, 107(1-3), 293–321, doi.org/10.1016/S0001-6918(01)00019-1.
- Grotheer, M., Hermann, P., Vidnyánszky, Z., & Kovács, G. (2014). Repetition probability effects for inverted faces. *NeuroImage*, 102P2, 416–423.
- Grotheer, M. & Kovács, G. (2014). Repetition probability effects depend on prior experiences. *Journal of Neuroscience*, 34(19), 6640–6646.
- Henson, R. N. (2003). Neuroimaging studies of priming. *Progress in Neurobiology*, 70, 53–81.
- Henson, R. N. A., & Rugg, M. D. (2003). Neural response suppression, haemodynamic repetition effects, and behavioural priming. *Neuropsychologia*, 41(3), 263–270.
- Henson, R. N., Goshen-Gottstein, Y., Ganel, T., Otten, L. J., Quayle, A., & Rugg, M. D. (2003). Electrophysiological and haemodynamic correlates of face perception, recognition and priming. *Cerebral Cortex*, 13, 793–805.
- James, T. W., & Gauthier, I. (2006). Repetition-induced changes in BOLD response reflect accumulation of neural activity. *Human Brain Mapping*, 27(1), 37–46, doi.org/10.1002/hbm.20165.
- Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., & Smith, S. M. (2012). FSL. *NeuroImage*, 62(2), 782–790.
- Julian, J. B., Fedorenko, E., Webster, J., & Kanwisher, N. (2012). An algorithmic method for functionally defining regions of interest in the ventral visual pathway. *NeuroImage*, 60(4), 2357–2364.
- Kaliukhovich, D. a. and Vogels, R. (2011). Stimulus repetition probability does not affect repetition suppression in macaque inferior temporal cortex. *Cerebral Cortex*, 21(7): 1547–1558.
- Kourtzi, Z., & Kanwisher, N. (2001). Representation of perceived object shape by the human lateral occipital complex. *Science (New York, N.Y.)*, 293(5534), 1506–1509.
- Kovács, G., Kaiser, D., Kaliukhovich, D. a., Vidnyánszky, Z., & Vogels, R. (2013). Repetition probability does not affect fMRI repetition suppression for objects. *The Journal of Neuroscience*, 33(23), 9805–9812.
- Kristjansson, A., & Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search. *Attention, Perception & Psychophysics*, 72(1), 5–18.
- Larsson, J., & Smith, A. T. (2012). fMRI repetition suppression: neuronal adaptation or stimulus expectation? *Cerebral Cortex*, 22(3), 567–576.
- Loffler, G., Yourganov, G., Wilkinson, F., & Wilson, H. R. (2005). fMRI evidence for the neural representation of faces. *Nature Neuroscience*, 8(10), 1386–90.
- MacEvoy, S. P., & Epstein, R. A. (2007). Position selectivity in scene- and object-responsive occipitotemporal regions. *Journal of Neurophysiology*, 98, 2089–2098.
- Mack, A., Pappas, Z., Silverman, M., & Gay, R. (2002). What we see: Inattention and the capture of attention by meaning. *Consciousness and Cognition*, 11(4), 488–506.
- Magnussen, S., Idås, E., & Myhre, S. H. (1998). Representation of orientation and spatial frequency in perception and memory: A choice reaction-time analysis. *Journal of Experimental Psychology. Human perception and performance*, 24(3), 707–718.
- Miller, E. K., & Desimone, R. (1994). Parallel neuronal mechanisms for short-term memory. *Science*, 263, 520–522.
- New, J., Cosmides, L., & Tooby, J. (2007). Category-specific attention for animals reflects ancestral priorities, not expertise. *Proceedings of the National Academy of Sciences, USA*, 104(42), 16598–16603.
- Nichols, T. E., & Holmes, A. P. (2001). Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Human Brain Mapping*, 25, 1–25.
- Persichetti, A. S., Thompson-Schill, S. L., Butt, O. H., Brainard, D. H., & Aguirre, G. K. (2015). Functional magnetic resonance imaging adaptation reveals a noncategorical representation of hue in early visual cortex. *Journal of Vision*, 15(6), 18, 1–19, doi:10.1167/15.6.18. [PubMed] [Article]
- Proctor, R. W. (1981). A unified theory for matching-task phenomena. *Psychological Review*, 88(4), 291–326.
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding

- in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87.
- Ridderinkhof, R., Ullsperger, M., Crone, E., & Nieuwenhuis, S. (2004). The role of the medial frontal cortex in cognitive control. *Science*, 306(5695), 443–447.
- Roseboom, W., Linares, D., & Nishida, S. (2015). Sensory adaptation for timing perception. *Proceedings of the Royal Society B: Biological Sciences*, 21(12), 2833–2841.
- Schacter, D. L., & Buckner, R. L. (1998). Priming and the brain. *Neuron*, 20(2), 185–195.
- Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M. M., & Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*, 11(9), 1004–1006.
- Summerfield, C., Wyart, V., Johnen, V. M., & de Gardelle, V. (2011). Human scalp electroencephalography reveals that repetition suppression varies with expectation. *Frontiers in Human Neuroscience*, 5(July), 67.
- Todorov, A., Gobbini, M. I., Evans, K. K., & Haxby, J. V. (2007). Spontaneous retrieval of affective person knowledge in face perception. *Neuropsychologia*, 45(1), 163–173.
- Todorovic, A., & de Lange, F. P. (2012). Repetition suppression and expectation suppression are dissociable in time in early auditory evoked fields. *The Journal of Neuroscience*, 32(39), 13389–13395.
- Vuilleumier, P., Henson, R. N., Driver, J., & Dolan, R. J. (2002). Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nature Neuroscience*, 5(5), 491–499.
- Wacongne, C., Labyt, E., van Wassenhove, V., Bekinschtein, T., Naccache, L., & Dehaene, S. (2011). Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proceedings of the National Academy of Sciences, USA*, 108(51), 20754–20759.
- Weiner, K. S., Sayres, R., Vinberg, J., & Grill-Spector, K. (2010). fMRI-adaptation and category selectivity in human ventral temporal cortex: regional differences across time scales. *Journal of Neurophysiology*, 103(6), 3349–3365.
- Wig, G. S., Grafton, S. T., Demos, K. E., & Kelley, W. M. (2005). Reductions in neural activity underlie behavioral components of repetition priming. *Nature Neuroscience*, 8(9), 1228–1233.
- Wiggs, C. L., & Martin, A. (1998). Properties and mechanics of perceptual priming. *Current Opinion in Neurobiology*, 8, 227–233.
- Winkler, A. M., Ridgway, G. R., Webster, M. A., Smith, S. M., & Nichols, T. E. (2014). Permutation inference for the general linear model. *NeuroImage*, 92, 381–397.
- Winston, J. S., Henson, R. N. a., Fine-Goulden, M. R., & Dolan, R. J. (2004). fMRI-adaptation reveals dissociable neural representations of identity and expression in face perception. *Journal of Neurophysiology*, 92, 1830–1839.