

Serial Dependence in Perceptual Decisions Is Reflected in Activity Patterns in Primary Visual Cortex

Elexa St. John-Saaltink,¹ Peter Kok,¹ Hakwan C. Lau,^{1,2} and Floris P. de Lange¹

¹Radboud University, Donders Institute for Brain, Cognition and Behaviour, 6525 EN Nijmegen, The Netherlands, and ²Department of Psychology, University of California Los Angeles, Los Angeles, California 90095-1563

Sensory signals are highly structured in both space and time. These regularities allow expectations about future stimulation to be formed, thereby facilitating decisions about upcoming visual features and objects. One such regularity is that the world is generally stable over short time scales. This feature of the world is exploited by the brain, leading to a bias in perception called serial dependence: previously seen stimuli bias the perception of subsequent stimuli, making them appear more similar to previous input than they really are. What are the neural processes that may underlie this bias in perceptual choice? Does serial dependence arise only in higher-level areas involved in perceptual decision-making, or does such a bias occur at the earliest levels of sensory processing? In this study, human subjects made decisions about the orientation of grating stimuli presented in the left or right visual field while activity patterns in their visual cortex were recorded using fMRI. In line with previous behavioral reports, reported orientation on the current trial was consistently biased toward the previously reported orientation. We found that the orientation signal in V1 was similarly biased toward the orientation presented on the previous trial. Both the perceptual decision and neural effects were spatially specific, such that the perceptual decision and neural representations on the current trial were only influenced by previous stimuli at the same location. These results suggest that biases in perceptual decisions induced by previous stimuli may result from neural biases in sensory cortex induced by recent perceptual history.

Key words: fMRI; MVPA; perceptual bias; primary visual cortex; priming; trial history

Significance Statement

We perceive a stable visual scene, although our visual input is constantly changing. This experience may in part be driven by a bias in visual perception that causes images to be perceived as similar to those previously seen. Here, we provide evidence for a sensory bias that may underlie this perceptual effect. We find that neural representations in early visual cortex are biased toward previous perceptual decisions. Our results suggest a direct neural correlate of serial dependencies in visual perception. These findings elucidate how our perceptual decisions are shaped by our perceptual history.

Introduction

The visual input we receive about the world is constantly interrupted by eye movements, blinks, and the occlusion of objects within our visual field. However, we perceive objects as continuous and the visual scene as stable. How is this stability obtained? One candidate mechanism for deriving stable representations from fluctuating noisy signals is temporal smoothing (i.e., the

brain may generate a weighted average of current input with previously obtained input). This may be a beneficial strategy, given that the world is stable over short time scales (Dong and Atick, 1995). Indeed, perceptual judgments are known to be influenced by previous trial history (Gao et al., 2009; de Lange et al., 2013).

Historically, priming is a classic example of how a previously seen stimulus can alter the response to a subsequent stimulus. When stimuli are physically (or conceptually) repeated, the behavioral response is facilitated. A set of recent studies also demonstrated strong serial dependence of perception between temporally adjacent stimuli, even for reliable (suprathreshold) visual stimuli that varied randomly over time (Cicchini et al., 2014; Fischer and Whitney, 2014; Liberman et al., 2014; Rahnev et al., 2015). In particular, remarkably, these recent studies show that previous stimuli can change—in other words, *distort*—perception.

Received Dec. 8, 2015; revised April 18, 2016; accepted April 21, 2016.

Author contributions: E.S.J.S., P.K., H.C.L., and F.P.d.L. designed research; E.S.J.S. performed research; E.S.J.S., P.K., and F.P.d.L. analyzed data; E.S.J.S., P.K., H.C.L., and F.P.d.L. wrote the paper.

This work was supported by The Netherlands Organization for Scientific Research NWO 404-10-037 to F.P.d.L. and H.C.L. We thank Valentin Wyart and Aurélie Pajani for providing code to calculate orientation signal.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Floris P. de Lange, Donders Institute for Brain, Cognition and Behaviour, Radboud University, PO Box 9101, 6500 HB Nijmegen, The Netherlands. E-mail: floris.delange@donders.ru.nl.

DOI:10.1523/JNEUROSCI.4390-15.2016

Copyright © 2016 the authors 0270-6474/16/366186-07\$15.00/0




Figure 1. Experimental paradigm. **A**, Every 5–7 s, two noisy gratings independently oriented at either 45° or 135° were presented. Participants reported the orientation of the grating indicated by a response cue. **B**, Analysis focuses on features of the current and previous trial. Green dashed circle indicates the stimulus cued for report.

Here, we sought to clarify the neural mechanisms underlying this perceptual effect of recent stimulus history. On the one hand, it is conceivable that recent sensory input may change the sensitivity of sensory neurons, for example, by increasing the sensitivity of neurons tuned to the previous input for a period of time following stimulus presentation (Fischer and Whitney, 2014). On the other hand, biases may only occur in downstream areas, at the stage of evidence accumulation and integration into a perceptual decision (Gold and Shadlen, 2007; Law and Gold, 2008; Hanks et al., 2011), leaving sensory processing unaffected. Here, we examined whether serial dependence in visual perception is already manifest at the level of early sensory representations using an fMRI dataset in which subjects performed a perceptual decision task on the orientation of briefly presented grating stimuli. This allowed us to determine the influence of the previous stimulus on sensory representations in early visual cortex and perceptual report on the current trial. To preview, we found that sensory representations in early visual cortex were biased by the perceptual choice on the previous trial, in a spatially specific fashion. This suggests a potential sensory mechanism for serial dependence in visual perception.

Materials and Methods

Participants. Twenty-seven healthy right-handed individuals (17 females, age 22 ± 2 years, mean \pm SD) with normal or corrected-to-normal vision gave written informed consent to participate in this study. Three participants did not complete the full fMRI session due to poor task performance or poor fixation ability; therefore, data from 24 participants were used for analyses. Experimental procedures were approved by the local ethics committee (Commissie Mensgebonden Onderzoek region Arnhem-Nijmegen, The Netherlands).

Stimuli. Stimuli consisted of two circular sinusoidal gratings (7° visual angle) presented for 200 ms at 5° along the horizontal meridian to the left and right of a central fixation point. Gratings were oriented at 45° and 135° independently of each other, and had the same orientation on 50% of trials. Grating orientations were pseudo-randomized such that all four orientation combinations (clockwise (CW)/CW; counterclockwise (CCW)/CCW; CW/CCW; CCW/CW) occurred equally often. To increase task difficulty, grating stimuli were embedded in random noise at

80% contrast. The contrast of the grating within the stimulus was presented at two levels; the same contrast was used for both gratings within a trial. The stimuli (grating + noise) were normalized such that overall contrast and luminance were constant for high and low grating contrast values (all stimuli were 80% contrast). This meant that more of the overall stimulus contrast was driven by the grating for high-contrast than for low-contrast stimuli (mean grating contrast was $\sim 5\%$). The central fixation point was displayed on a gray background throughout the experiment. Stimuli were generated using MATLAB (The MathWorks) in conjunction with Psychophysics Toolbox (Brainard, 1997). In the fMRI session, stimuli were displayed on a rear projection screen using a luminance-calibrated EIKI projector (60 Hz refresh rate, 1024×768 resolution), which participants viewed through a mirror. Stimuli were displayed on a LCD monitor (60 Hz refresh rate, 1024×768 resolution) during the behavioral session.

Experimental design. Stimuli were presented in an event-related design, with 5–7 s between trials. On each trial, a stimulus display of two gratings was briefly presented between two 500 ms periods of fixation (see Fig. 1A). At 700 ms after the onset of the stimulus display, two small chevrons pointing to the left or the right were presented for 1000 ms on either side of the fixation, which cued participants to respond to the grating that had been presented on that side of the screen. Participants performed a two-alternative forced choice (2AFC) task on the orientation of the grating specified by the response cue, using two buttons ($45^\circ/135^\circ$) on an MR-compatible button box. Response contingency was counterbalanced across participants, and participants switched response hands halfway through the experiment.

Each trial also contained a precue that consisted of two additional, smaller chevrons on either side of fixation that pointed in the same direction as the postcue on 75% of trials. The precue remained on screen during the initial fixation display and the stimulus display (i.e., 700 ms). On trials in which an orientation response was given, two small Cs appeared on either side of the fixation point 500 ms after the offset of the response cue, to prompt participants to rate their perceptual confidence on a scale from 1 to 4 (using four buttons in their other hand). These aspects of the task were included to answer a different research question than the one addressed in the current paper. Therefore, to maximize the reliability of the estimation of the neural response, we collapsed over the congruency between the precues and postcues in all analyses presented here; and for simplicity, we display only the stimuli relevant to the current analysis in Figure 1A.

Subjects participated in a behavioral session outside the scanner in the week before the fMRI session to familiarize subjects with the task and to titrate performance in the different conditions to 75% using a Quest staircase procedure (Watson and Pelli, 1983). During scanning, half of both congruent and incongruent trials were presented at low grating contrast, and the other half at high grating contrast, and contrast values were updated from the staircases after each block instead of after every trial. As mentioned above, we collapsed across cue-congruency and the different grating contrast levels in all analyses presented here.

The task was split into two runs of four blocks each, for a total of 512 trials. Between the two task runs, participants practiced remapping their response hand (the hand for grating and the hand for confidence response were switched) during the anatomical scan. Two additional scans were performed after the main experiment. A functional localizer was collected to enable identification of voxels that were maximally responsive to the grating stimuli in the left and right hemifield, and a retinotopy scan to allow delineation of early visual cortices. The localizer consisted of full-contrast gratings that were identical in size and position to those in the main experiment. Gratings were flickered at one of the stimulus locations per trial, alternating between left and right hemifield, at 2 Hz for 23.4 s. Each orientation (45° or 135°) was presented four times per location in a pseudo-random order. To ensure fixation, participants' task was to detect two letters (X, Z) in a stream of letters within the fixation bull's-eye. During the retinotopy scan, a flashing black-and-white checkerboard pattern (3 Hz) in a 90° wedge rotated on a black background in 30° steps (1 position per TR). Participants' task was to detect unpredictable changes in the color of the central fixation point (white to black). Nine cycles of CW and CCW rotation were presented. During both additional scans, participants responded to target events with a button press.

fMRI acquisition parameters. Functional images were acquired using a 3T Trio MRI system (Siemens) using a 32-channel head coil, with a $T2^*$ -weighted gradient-echo EPI sequence (TR 1.95 s, 31 transversal slices, $3 \times 3 \times 3$ mm in-plane resolution, TE 30 ms, FOV 192 mm \times 192 mm, flip angle of 80°). A high-resolution anatomical image was collected using a $T1$ -weighted MP-RAGE sequence (TR 2.3 s, TE 3.03 ms, $1 \times 1 \times 1$ mm in-plane resolution, GRAPPA acceleration factor of 2).

fMRI data preprocessing. Data were preprocessed using SPM8 (<http://www.fil.ion.ucl.ac.uk/spm>; Wellcome Trust Centre for Neuroimaging, London). The first three volumes of each task run were discarded to allow for time to achieve initial equilibrium. Functional images from the two task runs, the localizer, and retinotopy runs were spatially realigned to the mean image and temporally realigned to the first slice of each volume. The motion parameters resulting from spatial realignment were included as nuisance regressors in the GLMs. The structural image was coregistered with the functional volumes.

fMRI data analysis. An initial analysis of the functional localizer data was performed using SPM8, with regressors for left hemifield stimulation, right hemifield stimulation, and the motion parameters resulting from spatial realignment. A 128 s high-pass filter removed low-frequency signal components. Subtraction of the response to left and right hemifield stimulation was used to select stimulus-responsive voxels in each hemisphere for further analysis. In a separate analysis, *Freesurfer* (www.surfer.nmr.mgh.harvard.edu/) was used to inflate the cortical surface of each participant's $T1$ -weighted structural image and to analyze the functional data from the retinotopy session. Polar-angle maps were generated using Fourier-based methods and projected onto the surface of the inflated cortex according to established methods (Sereni et al., 1995), allowing retinotopic areas within early visual cortex to be visually identified and delineated. *Freesurfer* and SPM functions were used to convert the retinotopic labels from surface to volume space and to transform them into ROIs.

Within retinotopic ROIs V1, V2, and V3 of each hemisphere, we restricted our analyses to the 50 voxels that were the most responsive to the localizer. To remove slow drifts, preprocessed data from the localizer and task were linearly detrended. To estimate the response amplitude of each of these voxels to each single trial during the task, we applied the Least-Squares-Separate method outlined by Mumford et al. (2012) to the preprocessed data (Kok et al., 2013; Schoenmakers et al., 2014; Schlichting et

al., 2015). This method consists of running a separate GLM for every trial, such that each trial is modeled once as a regressor of interest, with all other trials combined into a single nuisance regressor. This method has been shown to improve the estimation of single-trial BOLD response, compared with a GLM with one regressor for each trial (Mumford et al., 2012). In addition to these regressors, we included separate regressors for break and end of run screens, as well as the motion parameters resulting from spatial realignment, their derivatives, and the square of these derivatives (i.e., 18 motion parameters in total). The data from the functional localizer were analyzed similarly using the Least-Squares-Separate method, with one GLM performed per trial. The resulting task and localizer β weights were normalized by z -scoring the values for each voxel.

For the main analyses, we first computed an orientation-specific signal for each trial by training a support vector machine (SVM) on the localizer data per hemisphere, and applying these SVMs to the task data to produce an SVM decision value for each task trial. We used the SVM decision value as a proxy for orientation-signal strength. To maximize the strength of this orientation signal, we determined the optimal number of voxels for each participant and each ROI. To do so, we calculated the mean orientation signal over all task trials for different numbers of voxels (5–50, in steps of 5) and selected the number of voxels at which the mean orientation signal peaked. We applied Platt Scaling (Platt, 2000) to transform SVM outputs to probabilities by passing them through a sigmoid (Niculescu-Mizil and Caruana, 2005; Charles et al., 2014).

Serial dependence analyses (behavior and fMRI). For all following analyses, orientation was recoded such that 45° was positive and 135° was negative. First, we constructed four regressors that captured, for each trial, the stimulus orientation at the responded to and nonresponded to locations, on both trial N and trial $N-1$ (Fig. 2A). We applied logistic regression to participants' binary perceptual choice to characterize the impact of current and previous stimuli on current perceptual choice. Parameter estimates indicate the extent to which the perceptual report on trial N is influenced by the stimulus orientations presented at each location on the current and previous trial. To correspondingly characterize the impact of current and previous stimuli on the orientation signal in early visual areas, we applied linear regression to the orientation-specific BOLD signal (i.e., SVM output) in primary visual cortex (V1).

To investigate whether the serial effect is dependent upon the previous stimulus or the previous percept of that stimulus, we separately modeled the previous stimulus (Fig. 2B, C, light blue bar) depending on response: "correct" when the previous percept was congruent with the previous stimulus; "incorrect" when the previous percept was incongruent with the previous stimulus; and "nonresponded to" when the response was made to the stimulus at the other location. Using these three regressors in combination with the two regressors that captured the stimuli on the current trial, we again applied logistic regression to participants' binary perceptual choice, and linear regression to the orientation signal in V1.

For all analyses, we used simple t tests at the group level to determine the robustness of each regressor's influence. To assess the location specificity of the effects, we used paired-sample t tests at the group level to compare regressors for reported and nonreported locations, separately for current and previous trial. Likewise, paired-sample t tests at the group level were used to compare the strength of the bias following correct and nonresponded to stimuli, and correct and incorrect stimuli. Finally, to assess whether higher-order extrastriate cortex displayed comparable serial dependence, we ran the same linear regression analyses performed in V1, separately for V2 and V3 ROIs.

Control analyses (fMRI). Because we used an event-related design, with an average ITI of 6 s, a potential concern is that the serial dependence we find in the fMRI data could be attributed to BOLD from the previous trial that is yet to return to baseline. Although this concern is partially mitigated by the fact that all trials were modeled in the context of the GLM, which attributes only the unique variance to each regressor, we performed a conservative control analysis in which we explicitly modeled variations in hemodynamic effects of the previous trial during single trial β estimation (i.e., the Least-Squares-Separate single trial GLMs, see above). Specifically, we captured any variation in the onset and duration of the BOLD response to the previous trial by modeling it with three complementary regressors: a canonical HRF and its first- and second-




Figure 2. Serial dependence effects on perceptual choice and orientation signal in V1. **A**, For each trial, four regressors captured the orientation of the stimuli at the responded to (blue) and nonresponded to (red) locations, on both trial N and trial $N-1$. **B**, **C**, Parameter estimates for the four conditions indicated in **A**. Parameter estimates indicate how strongly perceptual report on trial N is influenced by the stimulus orientations presented at each location on the current and previous trial. **B**, Perceptual report on trial N is influenced by the orientation of the stimulus presented at that location, but also by the previous stimulus presented at that location. **C**, The same as **B**, but with the orientation signal in V1 (classifier output) as the dependent variable. Error bars indicate SEM (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$).

order derivatives. If the serial dependence effects in visual cortex reflect residual BOLD activity evoked by the previous stimulus, then the trial history effects should no longer be present when this signal is removed from the single trial estimates. We therefore repeated the two linear regression analyses presented above on these single trial estimates.

Persistence of serial dependence. To evaluate the temporal limit of serial dependence on the perceptual report, we investigated the influence of the preceding four trials on the current behavioral response. We constructed regressors capturing the stimulus orientation at the responded to and nonresponded to locations on trial $N-2$, $N-3$, and $N-4$, and added them to the regressors for trials N and $N-1$ (refer to Fig. 2A). We then repeated the logistic regression on participants' binary perceptual choice.

Results

We investigated an fMRI dataset collected while participants reported the orientation one of two visual grating stimuli briefly presented to the left and right of a central fixation point. On average, participants were 78% correct ($\pm 5\%$) and responded after 595 ms (± 119 ms), indicating that participants followed task instructions.

To characterize the impact of current as well as previous stimuli on perceived orientation, we applied logistic regression to participants' binary perceptual choice. Similarly, to investigate the impact of current and previous stimuli on the orientation signal in

early visual cortical areas, we extracted orientation specific BOLD signals from visual cortex on every trial, and applied linear regression to the orientation-specific BOLD signal.

Perceived orientation was consistently biased toward the orientation of the preceding stimulus ($t_{(23)} = 5.41$, $p = 1.7e-05$; Fig. 2B). This serial dependence effect was spatially specific: perceptual decisions were more strongly influenced by previous stimuli at the same location than by stimuli at a different location ($t_{(23)} = 5.34$, $p = 2.0e-05$). In fact, perceived orientation on the current trial was slightly repelled away from the orientation of the previous stimulus at the other location ($t_{(23)} = -2.18$, $p = 0.040$).

Strikingly, the orientation signal in V1 was similarly biased toward the orientation presented on the previous trial at the same location (V2: $t_{(23)} = 2.08$, $p = 0.049$; V3: $t_{(23)} = 2.88$, $p = 0.0085$). This effect was not present at the opposite location (V2: $t_{(23)} = 0.65$, $p = 0.52$; V3: $t_{(23)} = 0.26$, $p = 0.79$), although the difference between the two locations was not statistically significant in V2 ($t_{(23)} = 0.63$, $p = 0.54$) and only approached significance in V3 ($t_{(23)} = 2.02$, $p = 0.056$).

Is the serial effect on perceptual choice and neural representation dependent upon the previous stimulus, or instead upon the previous percept? Most often, perception follows the stimulus, precluding such an analysis. However, trials in which perception diverges from stimulus input offer an opportunity to tease these factors apart. In the following, we will only consider effects of stimuli presented on previous trials at the same spatial location as the currently responded to stimulus, given the spatial specificity of the serial dependence effect demonstrated above. To investigate this question, we separately modeled the previous stimulus on the basis of the on perceptual report, resulting in regressors for "correct," "incorrect," and "nonresponded to" trials (see Materials and Methods). As can be seen in Figure 3B, the perceptual decision on the current trial was consistently biased toward the previous stimulus when it was correctly perceived ($t_{(23)} = 6.52$, $p = 1.2e-06$), but biased away from the physically presented stimulus toward the perceived stimulus when stimulus and perceptual choice diverged on the previous trial ($t_{(23)} = -4.98$, $p = 4.9e-05$), resulting in a significant difference between these conditions ($t_{(23)} = 6.92$, $p = 4.7e-07$). When no explicit perceptual decision was made on the previous stimulus (e.g., on the current trial the left grating was responded to,

but on the previous trial the right grating was responded to), there was still a strong bias toward the previous stimulus ($t_{(23)} = 4.91, p = 5.8e-05$), ruling out an explanation of response bias. The bias toward the previous stimulus was stronger when it was correctly reported than when it was not responded to ($t_{(23)} = 2.48, p = 0.021$).

The orientation signal decoded from BOLD activity in V1 displayed a similar profile, with a significant bias of orientation signals toward the previous stimulus when the previous stimulus was correctly perceived ($t_{(23)} = 2.38, p = 0.026$) or passively viewed ($t_{(23)} = 2.30, p = 0.031$) but no reliable effect when stimulus and choice diverged ($t_{(23)} = 0.40, p = 0.69$). This, however, did not culminate in a significant difference between correctly versus incorrectly responded trials ($t_{(23)} = 0.51, p = 0.61$), potentially due to the higher variability of the neural orientation signal compared with the behavioral report, in combination with the relatively low number of error trials (23% on average). There was no difference in the strength of the bias following correctly reported compared with nonreported stimuli ($t_{(23)} = 0.71, p = 0.48$). Again, a similar pattern of results was found in V2 and V3, albeit nonsignificantly in V2. Orientation signals were biased toward previous stimuli when they were correctly perceived (V2: $t_{(23)} = 1.05, p = 0.31$; V3: $t_{(23)} = 2.37, p = 0.027$) or passively viewed (V2: $t_{(23)} = 1.51, p = 0.14$; V3: $t_{(23)} = 2.45, p = 0.023$), with no reliable effect when stimulus and choice could be dissociated (V2: $t_{(23)} = 0.51, p = 0.62$; V3: $t_{(23)} = 1.54, p = 0.14$).

To ensure that our neural results were not dependent on residual BOLD signal from the evoked response to the previous stimulus, we modeled the neural response to the previous trial during the single trial β estimation. If the serial dependence effects in visual cortex reflect left over BOLD activity evoked by the previous stimulus, then the trial history effects should no longer be present when this signal is removed from the single trial estimates. However, there was still a reliable effect of the previous stimulus on the current neural response in V1 ($t_{(23)} = 2.38, p = 0.026$).

To investigate how long the influence of previous trials persists, we looked at the influence of stimuli from the four preceding trials on the current perceptual report. There is a consistent positive bias toward the orientation of preceding stimuli from the preceding three trials ($N-2: t_{(23)} = 5.10, p = 3.7e-05$; $N-3: t_{(23)} = 5.04, p = 4.2e-05$), specific to stimuli at the location that is responded to on trial N ($N-2: t_{(23)} = 5.19, p = 2.9e-05$; $N-3: t_{(23)} = 3.90, p = 0.00072$; Fig. 4). This bias drops off after three trials ($N-4: t_{(23)} = 1.41, p = 0.17$). Serial dependence of the orientation signal in visual cortex does not persist beyond the directly preceding trial ($N-2: t_{(23)} = 0.77, p = 0.45$), possibly due to the fact that the fMRI orientation signals were generally noisier than behavior.




Figure 3. Serial dependence is governed by previous percept rather than stimulus or response. **A**, The regressor for the stimulus on trial $N-1$ at the responded location on trial N (Fig. 2B, C, light blue) was subdivided on the basis of the perceptual decision on trial $N-1$: correct response (horizontal stripes), incorrect response (grid), responded to other location (no fill). **B**, **C**, Parameter estimates for the three conditions outlined in **A**. **B**, Previous stimuli that were correctly reported or not responded to exert a positive bias on current perception. For incorrect trials, perception is biased toward previous percept rather than stimulus. **C**, The same as in **B**, but with the orientation signal in V1 (classifier output) as the dependent variable. Here only the positive biases are present. Error bars indicate SEM (* $p < 0.05$, *** $p < 0.001$).




Figure 4. The influence of previous trials on perceptual choice persists for three trials. For this analysis, the orientation of the stimuli at the responded to (blue-green) and nonresponded to (orange-yellow) locations on trial $N-2$, $N-3$, and $N-4$ was added to the regressors in Figure 2A. The positive bias of previous stimuli on the current perceptual report persisted for three trials. Error bars indicate SEM (* $p < 0.05$, *** $p < 0.001$).

Discussion

Our perception at any given moment is influenced by both current and previous sensory signals. In this study, we investigated the neural mechanism underlying this serial dependence in perceptual decisions. Through extracting orientation-specific signals from visual cortex, we determined that the attractive bias exerted by the previous percept is present at the level of early sensory representations. This result sheds light on the mechanism behind the serial dependencies that have been reported in the literature, such as that choice on a current trial is influenced by the directly preceding stimulus (Gao et al., 2009; de Lange et al., 2013; Fischer and Whitney, 2014).

What specifically is carried over from one trial to the next? We distinguish three possibilities. First, serial dependence could be driven by the previous stimulus, via the bottom-up signal. Another option is that the perceptual choice made on the basis of the previous stimulus is carried over to the next trial. The crucial distinction according to this explanation is that serial dependence is more attributable to the percept of the previous stimulus than the previous stimulus per se. The third possibility is that serial dependence relies on the behavioral response that is coupled to the perceptual choice. In this explanation, the serial effect is motoric (as opposed to sensory or perceptual) in nature.

If serial dependence is driven by the bottom-up input, then there should be a consistent influence of previous stimuli based on their orientation. This is indeed what we found: there was a reliable bias toward the orientation of previous stimulus presented at the same location. However, perceptual choice is highly correlated to sensory input (i.e., generally, our perceptual experience of the world is a good reflection of the stimulation received by our sensory cortices). This makes it difficult to determine whether it is the bottom-up signal, or the perceptual choice, that carries over across trials. Because the orientation task used here was deliberately difficult (the contrast of the orientation signal within the noisy stimulus patches was titrated such that participants were 75% correct; the gratings were on average 5% contrast in 80% contrast noise), it resulted in a proportion of trials in which stimulus and choice diverged. These trials allow us to dissociate perceptual choice from the sensory input. On these trials, subsequent perceptual decision was biased toward the previous (incorrect) perceptual choice, instead of the stimulus that was presented (Fig. 3B). Furthermore, the attractive bias elicited by the previous trial was stronger following a correct perceptual choice than when the stimulus was not responded to. This suggests that perception, as opposed to sensory input per se, is what is carried over across time, exerting a positive bias on subsequent perceptual decisions.

If, on the other hand, the behavioral response (i.e., the button that was pressed) is carried over between trials, one would expect serial dependence to be location unspecific (i.e., if pressing button 1 would prime a subsequent button 1 response, this would not be localized to one visual hemifield, but would instead transfer across stimulus locations, as response location was randomized over trials). However, all effects of previous stimuli on both perceptual report and orientation signal in this study were location specific (see Fig. 2B), which rules out that simple motor response biases were responsible for our results. For the same reason, this rules out nonspecific decision biases, such as a predisposition to report the same orientation repeatedly, and instead suggests a spatially specific perceptual carryover.

The current results are in line with other recent reports suggesting a perceptual nature of serial dependence (Burr and Cicchini, 2014; Cicchini et al., 2014; Fischer and Whitney, 2014;

Liberman et al., 2014). It should be noted that perceptual decisions in this experiment were measured using a 2AFC task, which can be contrasted with the more continuous measures of perception used in these previous reports. Given that it is necessarily less fine-grained, a binary measure of perception may be less sensitive to subtle perceptual biases. However, because the oriented gratings in our stimuli were low contrast, embedded in high-contrast white noise (which contains signal for all possible orientations) it is feasible that small (serial dependence) biases would lead to the false perception of the orthogonal grating orientation (Pajani et al., 2015). Because in the current study participants were presented with only two (orthogonal) orientations, the biases reported here likely reflect carryover of perceptual decisions about whether the grating was oriented CW (45°) or CCW (135°), rather than subtle perceptual biases on the order of a few degrees, such as those previously described in continuous-report designs (Fischer and Whitney, 2014). A paradigm that combines a continuous measure of perception with a neuroimaging measure of stimulus representations in low-level sensory cortices, such as one in which neural correlates of subtle perceptual biases have previously been measured (Kok et al., 2013), could be a promising avenue for future investigation of serial dependence.

What may be the neural mechanisms underlying the serial dependence of perceptual choice? One hypothesis is that the sensitivity of sensory neurons tuned to the previous percept may be increased for a brief period following stimulus presentation, thereby influencing current perception (Fischer and Whitney, 2014). In line with this idea, we found that the orientation of the previous perceptual decision biases the representation of stimuli in early visual cortex. Interestingly, similar effects of the decision variable on sensory responses have been observed during the period in which a perceptual decision unfolds (Nienborg and Cumming, 2009; Wimmer et al., 2015). We speculate that this biasing of sensory responses due to the decision may persist, thereby biasing subsequent sensory processing.

Notably, we found that nonreported gratings also have an influence on the subsequent perceptual decision and orientation signal in visual cortex. While this result could be seen as evidence for the stimulus-driven account of serial dependence, it should be borne in mind that nonreported gratings were still attentively perceived by the participants. Namely, subjects were only informed about which stimulus to report after the stimuli had been removed from the screen, therefore necessitating an implicit perceptual decision about both stimuli.

The effects in primary visual cortex and in perceptual report were spatially specific: perceptual decisions and neural representations on the current trial were only influenced by previous stimuli at the same location. This was equally true of the influence that stimuli from two and three trials back had on the perceptual report. This location specificity may appear at odds with a previous report (Fischer and Whitney, 2014), in which serial dependence transferred across spatial locations (i.e., that serial dependence smooths across time and space). One explanation for this could be due to differences in participants' attentional state between the two designs: when serial dependence was found to transfer across locations, only one stimulus was attended on each trial, whereas in the current study, both stimuli required a certain level of attention because either stimulus could be cued for report. Future designs that manipulate both the number of attended stimuli and total number of stimuli may elucidate the cause of this discrepancy.

A potential limitation of the present study is the possibility that the serial dependence we measured in visual cortex is a result

of a residual BOLD response to the previous stimulus. However, there are several reasons why we believe that our results are not the result of autocorrelation in the BOLD signal. First, to increase the separability of single trials, the interval between trials was jittered such that trials were presented every 5–7 s. Second, each trial was independently modeled using a technique that maximizes extraction of the signal unique to each trial (Mumford et al., 2012). This approach to estimating the trial-specific BOLD signal, in combination with the jitter between trials, should allow the signal from individual trials to be dissociated from neighboring trials. Furthermore, the serial dependence in visual cortex persisted even after we applied a conservative approach to regress out the BOLD response to the previous trial. Therefore, it appears plausible that the bias in sensory cortex is generated at the moment of bottom-up stimulation, rather than reflecting a spillover of activity to the previous stimulus.

Our perception of the world is partly determined by our (often implicit) priors about the statistical regularities in the environment (Yuille and Kersten, 2006; Chalk et al., 2010; Kok et al., 2013). Serial dependence, such as reported here, can be understood as one such prior—that the world is stable over short time scales. Interestingly, serial dependence is not restricted to low-level stimuli, such as used here, but also extends to complex and naturalistic stimuli, such as faces (Liberman et al., 2014) and numerosity (Cicchini et al., 2014). Other sequential effects, such as repetition suppression, may similarly be cast as the result of this same prior (Summerfield et al., 2008; Todorovic et al., 2011; Henson, 2015): if the world is generally stable, objects are more likely to repeat than change. Given the stability of the sensory world, such a prior could make visual processing more robust by filtering out temporal noise.

References

- Brainard DH (1997) The Psychophysics Toolbox. *Spatial Vision* 10:433–436. [CrossRef Medline](#)
- Burr D, Cicchini GM (2014) Vision: efficient adaptive coding. *Curr Biol* 24:R1096–R1098. [CrossRef Medline](#)
- Chalk M, Seitz AR, Seriès P (2010) Rapidly learned stimulus expectations alter perception of motion. *J Vis* 10:1–18. [CrossRef Medline](#)
- Charles L, King JR, Dehaene S (2014) Decoding the dynamics of action, intention, and error detection for conscious and subliminal stimuli. *J Neurosci* 34:1158–1170. [CrossRef Medline](#)
- Cicchini GM, Anobile G, Burr DC (2014) Compressive mapping of number to space reflects dynamic encoding mechanisms, not static logarithmic transform. *Proc Natl Acad Sci U S A* 111:7867–7872. [CrossRef](#)
- de Lange FP, Rahnev DA, Donner TH, Lau H (2013) Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. *J Neurosci* 33:1400–1410. [CrossRef](#)
- Dong DW, Atick JJ (1995) Statistics of natural time-varying images. *Network* 6:345–358. [CrossRef](#)
- Fischer J, Whitney D (2014) Serial dependence in visual perception. *Nat Neurosci* 17:738–743. [CrossRef Medline](#)
- Gao J, Wong-Lin K, Holmes P, Simen P, Cohen JD (2009) Sequential effects in two-choice reaction time tasks: decomposition and synthesis of mechanisms. *Neural Comput* 21:29. [CrossRef Medline](#)
- Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574. [CrossRef Medline](#)
- Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN (2011) Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J Neurosci* 31:6339–6352. [CrossRef Medline](#)
- Henson RN (2015) Repetition suppression to faces in the fusiform face area: a personal and dynamic journey. *Cortex*. Advance online publication. Retrieved Oct. 31, 2015. doi: 10.1016/j.cortex.2015.09.012. [CrossRef Medline](#)
- Kok P, Brouwer GJ, Van Gerven MA, de Lange FP (2013) Prior expectations bias sensory representations in visual cortex. *J Neurosci* 33:16275–16284. [CrossRef](#)
- Law CT, Gold JI (2008) Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat Neurosci* 11:505–513. [CrossRef Medline](#)
- Liberman A, Fischer J, Whitney D (2014) Serial dependence in the perception of faces. *Curr Biol* 24:2569–2574. [CrossRef Medline](#)
- Mumford JA, Turner BO, Ashby FG, Poldrack RA (2012) Deconvolving BOLD activation in event-related designs for multivoxel pattern classification analyses. *Neuroimage* 59:2636–2643. [CrossRef Medline](#)
- Niculescu-Mizil A, Caruana R (2005) Predicting good probabilities with supervised learning, pp 625–632. *Proceedings of the 22nd International Conference on Machine Learning, Bonn, Germany, August 7–11, 2005*. [CrossRef](#)
- Nienborg H, Cumming BG (2009) Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* 459:89–92. [CrossRef Medline](#)
- Pajani A, Kok P, Kouider S, de Lange FP (2015) Spontaneous activity patterns in primary visual cortex predispose to visual hallucinations. *J Neurosci* 35:12947–12953. [CrossRef Medline](#)
- Platt JC (2000) Probabilities for SV machines. In: *Advances in large margin classifiers* (Smola AJ, Bartlett P, Schölkopf B, Schuurmans D, eds), pp 61–74. Cambridge, MA: Massachusetts Institute of Technology.
- Rahnev D, Koizumi A, McCurdy LY, D'Esposito M, Lau H (2015) Confidence leak in perceptual decision making. *Psychol Sci* 26:1664–1680. [CrossRef Medline](#)
- Schlichting ML, Mumford JA, Preston AR (2015) Learning-related representational changes reveal dissociable integration and separation signatures in the hippocampus and prefrontal cortex. *Nat Commun* 6:8151. [CrossRef Medline](#)
- Schoenmakers S, Guclu U, van Gerven M, Heskes T (2014) Gaussian mixture models and semantic gating improve reconstructions from human brain activity. *Front Comput Neurosci* 8:173. [CrossRef Medline](#)
- Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB (1995) Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268:889–893. [CrossRef Medline](#)
- Summerfield C, Trittschuh EH, Monti JM, Mesulam MM, Eger T (2008) Neural repetition suppression reflects fulfilled perceptual expectations. *Nat Neurosci* 11:1004–1006. [CrossRef Medline](#)
- Todorovic A, Van Ede F, Maris E, de Lange FP (2011) Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: an MEG study. *J Neurosci* 31:9118–9123. [CrossRef Medline](#)
- Watson AB, Pelli DG (1983) QUEST: a Bayesian adaptive psychometric method. *Percept Psychophys* 33:113–120. [CrossRef Medline](#)
- Wimmer K, Compte A, Roxin A, Peixoto D, Renart A, de la Rocha J (2015) Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT. *Nat Commun* 6:6177. [CrossRef Medline](#)
- Yuille A, Kersten D (2006) Vision as Bayesian inference: analysis by synthesis? *Trends Cogn Sci* 10:301–308. [CrossRef Medline](#)