

Neural Dynamics of Serial Dependence in Numerosity Perception

Michele Fornaciai and Joonkoo Park

Abstract

■ Serial dependence—an attractive perceptual bias whereby a current stimulus is perceived to be similar to previously seen ones—is thought to represent the process that facilitates the stability and continuity of visual perception. Recent results demonstrate a neural signature of serial dependence in numerosity perception, emerging very early in the time course during perceptual processing. However, whether such a perceptual signature is retained after the initial processing remains unknown. Here, we address this question by investigating the neural dynamics of serial dependence using a recently developed technique that allowed a reactivation of hidden memory states. Participants performed a numerosity discrimination task during EEG recording, with task-relevant dot array stimuli preceded by

a task-irrelevant stimulus inducing serial dependence. Importantly, the neural network storing the representation of the numerosity stimulus was perturbed (or pinged) so that the hidden states of that representation can be explicitly quantified. The results first show that a neural signature of serial dependence emerges early in the brain signals, starting soon after stimulus onset. Critical to the central question, the pings at a later latency could successfully reactivate the biased representation of the initial stimulus carrying the signature of serial dependence. These results provide one of the first pieces of empirical evidence that the biased neural representation of a stimulus initially induced by serial dependence is preserved throughout a relatively long period. ■

INTRODUCTION

In perception, the incoming sensory information is noisy, and biological sensors like the eye are highly unstable (i.e., due to frequent eye and head movements). Yet, the brain provides us with a stable and seamless conscious experience of the visual world, and how the brain achieves this stability remains an open question. Recent studies suggest that one way to stabilize perception is to integrate stimulus information over time to smooth out noise from neural signals (Burr & Cicchini, 2014; Fischer & Whitney, 2014). This operation may indeed represent a successful strategy (Cicchini, Mikellidou, & Burr, 2018; Burr & Cicchini, 2014), thanks to the stability of our visual environment over short timescales. However, in the experimental context where visual stimuli can be arbitrarily modulated, such an operation predicts a systematic bias in perceptual tasks, with current stimuli appearing more similar to previous ones—a bias that is commonly referred to as “attractive serial dependence.”

A growing literature outlines several characteristics of serial dependencies in visual perception. First, such an attractive bias seems to generalize across many and very different visual features, starting from basic attributes, such as orientation (Cicchini, Mikellidou, & Burr, 2017; Fritsche, Mostert, & de Lange, 2017; Fischer & Whitney,

2014), position (Manassi, Liberman, Kosovicheva, Zhang, & Whitney, 2018), motion (Alais, Leung, & Van der Burg, 2017), and numerosity (Fornaciai & Park, 2018a, 2018b, 2019a, 2019b; Cicchini, Anobile, & Burr, 2014; Corbett, Fischer, & Whitney, 2011), to more complex attributes, such as the ensemble representation of a visual scene (Manassi, Liberman, Chaney, & Whitney, 2017), the perception of visual variance (Suárez-Pinilla, Seth, & Roseboom, 2018), and face perception (Liberman, Manassi, & Whitney, 2018; Taubert, Van der Burg, & Alais, 2016; Xia, Leib, & Whitney, 2016; Liberman, Fischer, & Whitney, 2014). This suggests that serial dependence represents a global brain process affecting all aspects of perception. Second, serial dependencies are dependent on attention, mostly occurring when the object (Fischer & Whitney, 2014) or at least its spatial position (Fornaciai & Park, 2018a, 2018b) is attended by participants and/or relevant for the task. Finally, although the nature of serial dependence has been subject to debate (e.g., Alais et al., 2017; Bliss, Sun, & D’Esposito, 2017; Fritsche et al., 2017), evidence suggests that it is a bias arising during perceptual processing and not at later working memory encoding or decision stages (Pascucci et al., 2019; Fornaciai & Park, 2018a, 2018b; Manassi et al., 2018; Cicchini et al., 2017). However, this does not exclude the possibility that the initial perceptual bias interacts with subsequent stages and gets amplified (Fornaciai & Park, 2018a; Fritsche et al., 2017; Cicchini et al., 2017).

Although the behavioral signature of serial dependence has been thoroughly investigated, much less is known about the neural mechanisms underlying the attractive bias. Results using fMRI have shown that, in the context of orientation perception, the bias arises in the primary visual cortex (St. John-Saaltink, Kok, Lau, & de Lange, 2016)—possibly implicating a very early sensory/perceptual mechanism. Although due to the nature of the fMRI technique, such an activity in early visual cortex may alternatively reflect reentrant signals. Furthermore, in a previous study from our group using EEG (Fornaciai & Park, 2018a), we have demonstrated that a neural signature of serial dependence—in the context of numerosity perception—emerges very early (~50 msec) after stimulus onset, thus supporting the idea that the bias operates on the initial perceptual representation of the stimulus. Moreover, in that study, participants passively watched a sequence of dot arrays while paying no attention to their numerosity, only making a response to occasional oddball stimuli, which provides strong evidence that serial dependence occurs independently from a decision process. However, the rapid nature of such a paradigm did not allow us to track how the “biased” neural representation of numerosity evolves over time. Indeed, there is evidence that the behavioral effect of serial dependence seems to increase over time if a forced pause is introduced between the stimulus and response (Bliss et al., 2017; Fritsche et al., 2017). Such results were interpreted as an increasingly distorted stimulus representation during working memory storage, although it is not clear what process may be involved in that effect.

Here, we aim to investigate how a visual representation affected by serial dependence evolves over time during an active perceptual decision-making task. More specifically, we tested whether or not the biasedness of a representation induced by the recent history of stimulation is retained throughout the time course of a trial. Participants performed a numerosity discrimination task during EEG recording, judging which of the two sequentially presented dot arrays (“reference” followed by “probe”) contained more items. To induce serial dependence, a task-irrelevant “inducer” stimulus was presented before the two task-relevant arrays, following the procedure used in previous studies (e.g., Fornaciai & Park, 2018a, 2018b). Additionally, we took advantage of the recent findings, which demonstrated that activating a neuronal network by presenting a brief, high-energy neutral stimulus evokes a brain signal carrying the signature of information stored in such a network (Wolff, Jochim, Akyürek, & Stokes, 2017; Wolff, Ding, Myers, & Stokes, 2015). Following this technique, we presented two bright flashes between the reference and probe, serving as “pings,” to reactivate the working memory representation of the reference array affected by serial dependence. We reasoned that activity evoked by pings should reflect a stimulus representation carrying the signature of serial dependence. In other words, the activity evoked by the

pings should reflect the current content of visual working/short-term memory (Wolff et al., 2015, 2017), including the stored perceptual representation of the reference stimulus biased by serial dependence. We complemented this pinging approach with a multivariate neural decoding technique to assess how patterns of brain activity evolve over time.

Our design and approach entertained the following hypothesis. If memory traces of the visual representation biased by serial dependence are preserved after the initial perceptual processing stage, then the neural pattern associated with this initial visual representation should be detectable (or decodable) at a later time point, when this trace of representation is reactivated by the pings. To preview, our results first show that serial dependence starts very early after stimulus onset, in line with previous findings and suggesting that the effect starts at the early level of perceptual processing. Second, the results show that the neural representation associated with this initial perceptual processing, when pinged, is detectable at a later time point, providing evidence that the biased representation remains stored throughout a relatively long retention period.

METHODS

Participants

Thirty-five individuals took part in the study. Participation in the study was rewarded with course credits or monetary compensation (10 USD/hr). All participants had normal or corrected-to-normal vision and were naïve to the purpose of the study. All participants signed an informed consent form before participation in the study. Experimental procedures were approved by the University of Massachusetts institutional review board and are in line with the Declaration of Helsinki. Eight participants were excluded from data analysis (see Exclusion Criteria section below), leaving 27 participants in the final group (20 women, mean age = 22.5 years, $SD = 3.6$ years).

Apparatus and Stimuli

Stimuli were generated using the Psychophysics Toolbox (Kleiner et al., 2007; Pelli, 1997) in MATLAB (Version r2013b; The MathWorks, Inc.) and were presented on an LCD screen located approximately 90 cm from the participant. The screen encompassed approximately $34^\circ \times 19^\circ$ at the viewing distance of 90 cm and was running at 144 Hz.

Stimuli were dot arrays, including black and white dots (50%/50% in most of the cases; in case of odd numerosities, the color of the exceeding dot was randomly assigned) presented on a gray background. A mixture of black and white dots was employed to keep the global luminance across the array approximately similar to the luminance of the gray background. Three different dot

array stimuli were presented on each trial in the main conditions of the study (in this order): a task-irrelevant “inducer” comprising either 11 or 23 dots, a reference stimulus always comprising 16 dots, and a probe stimulus with a variable numerosity across trials (8, 10, 13, 16, 20, 25, or 32 dots). The arrays were built to span similar ranges in three orthogonal dimensions: numerosity, size, and spacing, following the design previously used by DeWind, Adams, Platt, and Brannon (2015) and Park, DeWind, Woldorff, and Brannon (2016). As the purpose of the study was to investigate serial dependence in numerosity perception and because the effect of other non-numerical dimensions has been addressed elsewhere (Fornaciai, Farrell, & Park, 2019; Fornaciai, Brannon, Woldorff, & Park, 2017; Fornaciai & Park, 2017; Park et al., 2016), the different levels of nonnumerical dimensions were collapsed together during data analysis. See DeWind et al. (2015) and Park et al. (2016) for more details about this stimulus construction procedure.

Regarding the stimulus parameters, the minimum individual area of the dots was set to 113 pixel², equal to a diameter of 0.21° (12 pixel); the maximum value was 452 pixel², corresponding to a diameter of 0.42° (24 pixel). Regarding the field area, which refers to the virtual invisible circular area within which the dots were drawn, the minimum value was set to 70,685 pixel², corresponding to a diameter of 5.16° (300 pixel), and the maximum value was 223,960 pixel², equal to a diameter of 10.32° (600 pixel).

The “ping” stimulus (see Procedure section below) that was flashed during the interval between the reference and probe stimulus was a white circle covering an area equal to the field area of the preceding reference stimulus.

Procedure

The experiment took place in a quiet and dimly illuminated room. Throughout each block, three trial types were intermixed: “no-ping” trials, “ping” trials, and “catch” trials. The sequence of stimuli presented in the no-ping and ping trials is shown in Figure 1. In all cases, participants fixated on the central fixation point, and the sequence of stimuli was presented at the center of the screen. In the no-ping trials (Figure 1A), a sequence of three stimuli was presented on the screen: First, a task-irrelevant inducer stimulus (11 or 23 dots), followed by the reference stimulus (16 dots) after 750 msec, and then a probe stimulus (8–32 dots) after 2450 msec. All the stimuli were presented on the screen for 250 msec each. At the end of the trial, participants were instructed to report whether the reference or the probe stimulus contained more dots. Response time was limited to 1250 msec after probe onset and was marked by the fixation cross turning red but with no corrective feedback. Participants were instructed to respond as fast as they could. After the end of the response period, the fixation cross returned to black, and the next trial started after 1500 msec.

In the ping trials, the sequence of inducer, reference, and probe stimuli was identical to the no-ping trials. However, two bright white circles were additionally presented on the screen (duration = 100 msec) between the reference and probe. In those cases, the sequence was as follows: First, the inducer was presented on the screen, followed by the reference stimulus after 750 msec. Then, the reference stimulus was followed by a sequence of two white circle (ping) stimuli and finally the probe stimulus. The ISI between different stimuli was in all cases 750 msec (Figure 1B).

To avoid presenting the same sequence of stimuli all the time we included some catch trials intermixed with the other types of trial. In the catch trials, only the inducer stimulus was presented on the screen, and the fixation cross turned orange right after it. In those cases, participants were instructed to press a different button as fast as they could. Doing so, participants thus had to pay attention to the inducer to promptly provide a response in the case of a catch trial.

Overall, the study comprised 10 blocks of trials. Each block included 24 no-ping trials, 24 ping trials, and 6 catch trials, randomly ordered. Participants were free to take breaks between blocks, and the experiment took about 1 hr.

Electrophysiological Recording and Analysis

EEG Recording

While participants performed the task described above, the EEG was continuously recorded for the entire duration of the experiment. We used a 64-channel, extended coverage, triangulated equidistance cap (M10, EasyCap, GmbH), with a sampling rate of 1000 Hz (actiCHamp, Brain Products). During the experiment, the channels were referenced to the vertex (Cz). The EOG was monitored by means of electrodes positioned below the left eye and lateral to the left and right canthi to monitor artifacts due to blinks and eye movements. Across all the channels, impedance was usually kept below 15 k Ω most of the time, but impedances up to 35 k Ω were tolerated.

Preprocessing and ERPs

EEG data were analyzed offline using MATLAB (Version R2015b, The MathWorks, Inc.), using the functions provided by the EEGLAB package (Delorme & Makeig, 2004) and the ERPLAB toolbox (Lopez-Calderon & Luck, 2014). Continuous data were first visually inspected to remove noisy segments before proceeding with preprocessing. During preprocessing, the EEG signals were high-pass filtered (0.1 Hz) and were rereferenced to the average value of all the 64 channels. The continuous EEG data were then segmented into epochs from –300 msec before to 2700 msec after the onset of the reference stimulus, with a baseline correction using the prestimulus interval.

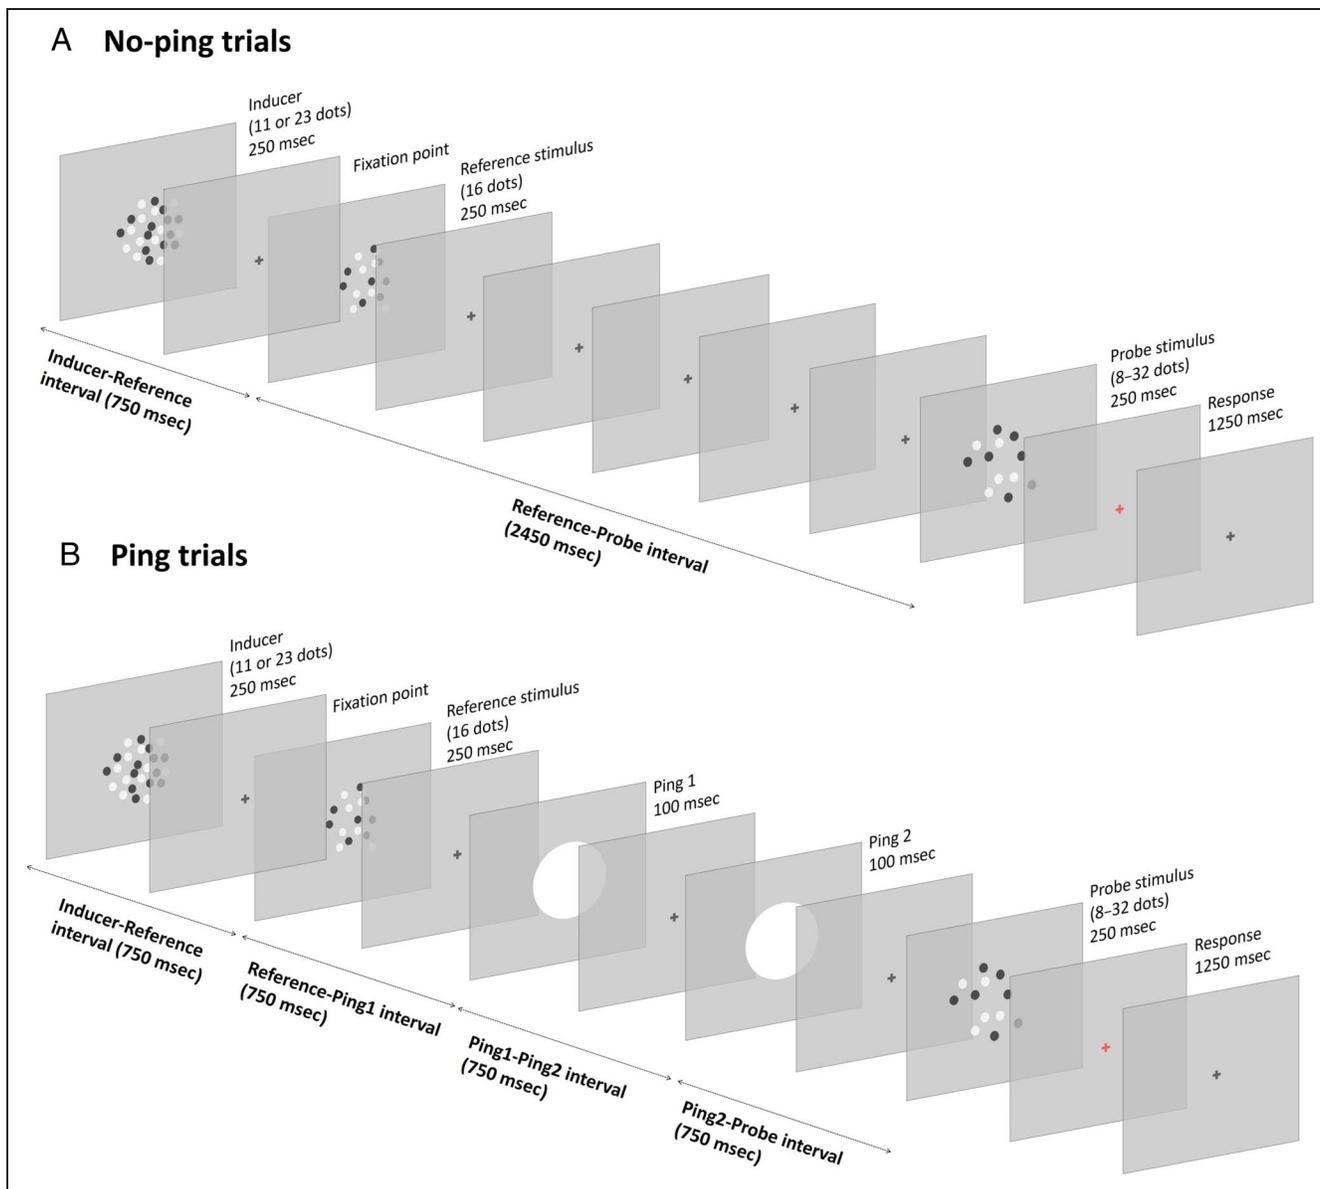


Figure 1. Experimental procedure. (A) Presentation procedure in the “no-ping trials.” While participants kept their gaze on a central fixation cross, a sequence of three dot arrays was presented at the center of the screen. First, a task-irrelevant inducer stimulus (11 or 23 dots) was presented, followed by the reference stimulus (16 dots) after 750 msec, and then a probe stimulus (8–32 dots) after 2450 msec. The stimulus duration was 250 msec each for all the stimuli. At the end of the trial, participants were instructed to report whether the second (reference) or third (probe) array in the sequence contained more dots by pressing the appropriate key on a standard keyboard. The period during which a response was accepted was marked by the fixation cross turning red and was limited to 1250 msec after probe onset, and participants were told to respond as fast as they can. After the end of the response period, the fixation cross returned black, and the next trial started after 1500 msec. Before starting the experiment, participants were told that the first stimulus was not relevant for the task but anyway to pay attention to the entire sequence of the stimuli. (B) Presentation procedure in the “ping trials.” The sequence of inducer, reference, and probe dot arrays was identical to the no-ping trials, with the exception that two bright white circles were presented on the screen (duration = 100 msec) between the reference and probe. Finally, on some occasions only, the inducer stimulus was presented on the screen (“catch” trials), and the fixation cross turned orange right after it (not shown in the figure). In those cases, participants were instructed to press the spacebar as fast as they can. Note that the stimuli depicted in the figure are not to scale. The timing measures reported in the figure refer to the ISI (not SOA) between two successive stimuli.

We then performed an independent component analysis (ICA) to remove eye movement and other muscular artifacts. After ICA, we further excluded trials containing eye-blink artifacts or other large artifacts by applying the step-like artifact rejection tool in ERPLAB. Trials were rejected whenever activity from the eye channel and a set

of posterior channels of potential interest (i.e., based on previous studies; Iz, O9', O10', Oz, O1', O2', PO7', PO8'; Fornaciai & Park, 2017, 2018a) exceeded a threshold equal to 30 μ V (in a time window spanning 400 msec, with 20-msec steps). This procedure led to an average (\pm SD) rejection rate of 8.55% \pm 10.94%. To create ERPs, the

epochs were selectively averaged and low-pass filtered (30 Hz). Finally, the ERPs were smoothed with a 100-msec sliding window to remove alpha noise before computing the grand average.

Time Window of Interest Selection for the ERP Analysis

Time windows of interest for the ERP analyses and statistical tests were chosen based on previous results. Previous ERP studies of numerosity perception have shown that the brain waves around 150–200 msec after stimulus onset are most sensitive to numerosity and to modulations affecting perceived numerosity (Fornaciai & Park, 2018c; Fornaciai et al., 2017; Park et al., 2016). Furthermore, previous results also show that a signature of serial dependence in numerosity perception emerges from ERPs around 200 msec. Thus, when assessing the serial dependence effect on ERPs evoked by the reference stimulus, our statistical analysis focused on a broad 100-msec time window centered at 200 msec after reference onset. Unlike the ERPs evoked by the reference stimulus, we did not have specific predictions concerning precisely when the ERPs evoked by the neutral “ping” stimulus should carry information about the representation in working memory (Wolff et al., 2017). Thus, the time window of interest was chosen as a 100-msec window around the peak of the collapsed ERP evoked by the pings, including the two inducer conditions, separately for the first and second ping. Single-subject ERPs were then averaged across these time windows, separately for the two inducer conditions, and tested with a paired *t* test.

Multivariate Pattern Analysis in the Time Domain

To achieve a better description of the temporal dynamics of serial dependence during the numerosity discrimination task, we applied a multivariate pattern analysis in the time domain (King & Dehaene, 2014) using the Neural Decoding Toolbox (Meyers, 2013). This method allows us to evaluate how neural activity patterns coming from multiple sensors across the scalp differ between different experimental conditions, and how such patterns of activity generalize across time. We used this analysis to assess how serial dependence affects the pattern of brain activity evoked by the reference (or the ping) stimulus as a function of the preceding inducer stimulus, and whether and to what extent the serial dependence effect changes over time. This neural decoding analysis involved the training of a support vector machine (SVM) classifier on a subset of the data corresponding to specific conditions and its predictions about which stimulus was presented in the remaining subset of data and across the entire time course of activity. Conditions were based on the inducer stimulus presented on each trial, so that the activity to be decoded always corresponded to the reference, but with trials divided as a function of the inducer numerosity.

Using this neural decoding approach, we evaluated whether and to what extent the pattern of brain activity evoked by the reference stimulus differs as a function of the inducer numerosity, and whether the activity evoked by the ping stimulus carries the signature of the biased reference representation. Brain responses were evaluated across the full epoch spanning from –300 to 2700 msec. Different comparisons were tested individually for each participant, training the SVM classifier with the responses (i.e., activity recorded at all the channels with the exception of the EOG channels) to the two class of stimuli at hand (i.e., trials divided as a function of inducer numerosity) and testing it on another subset of trials not used in the training phase (using a leave-one-trial-out cross-validation). To optimize the analysis, we followed a series of practices suggested by Grootswagers, Wardle, and Carlson (2017) to account for high noise in single-trial EEG data and improve the decoding procedure. First, we created “pseudotrials,” taking the average of randomly chosen groups of 12 trials to improve signal-to-noise ratio. Second, to avoid overfitting, the number of features (i.e., channels) included in the analysis was limited to the five most significant ones, determined using a univariate ANOVA. Moreover, we averaged activity across large time windows to smooth out noise such as alpha oscillations (i.e., 100-msec time windows with 25-msec step). Finally, the decoding procedure was repeated 30 times for each participant using different subsets of data for training and testing (as well as different groups of trials to generate pseudotrials), and the average of the 30 runs was taken as the final estimate of the classification performance. The outcome of the decoding analysis was a temporal generalization plot showing the performance of the classifier at each time point, with classification accuracy reflecting how well the pattern classifier can discriminate two conditions (see Figure 5A). Namely, activity along the diagonal of the temporal generalization matrix reflects the classification accuracy obtained by training and testing the classifier at the same time point, whereas off-diagonal decoding reflects the classification accuracy obtained by performing training at one time point and testing at all the different time points. This latter procedure allows us to assess whether and to what extent a specific pattern of brain activity generalizes to activity at different time points (King & Dehaene, 2014).

Time Window of Interest Selection for the Multivariate Pattern Analysis

Classification accuracies obtained from the multivariate pattern analysis were tested across several time windows of interest to assess the effects of inducer numerosity on the reference and ping stimuli. The analysis in this context was divided in two parts. The first part concerns the patterns of activity along the diagonal of the temporal generalization matrix, where classification is performed

by training and testing the classifier at the same time point; the second part concerns off-diagonal patterns of activity, revealing whether and to what extent a specific pattern generalizes over different times. This latter analysis represents the crucial aspect of our main hypothesis concerning the temporal stability of the biased representation of the reference array.

Time windows of interest corresponding to times around the reference onset were chosen according to our previous study (Fornaciai & Park, 2018a). Namely, two time windows were selected, spanning 0–250 and 250–500 msec after the reference onset. Time windows along the diagonal corresponding to the ping onsets were chosen to span the same duration as the first time window at the reference array (250 msec). Off-diagonal time windows were chosen to span a broad temporal extent to capture how the reference representation gets reactivated by the pings. To this aim, we chose eight time windows (four along the vertical dimension and four along the horizontal dimension). Time windows chosen in this way spanned 0–250 and 250–500 msec corresponding to the reference time windows and from the onset to 500 msec after each of the two ping onsets. Corresponding time windows along the two off-diagonal directions were then collapsed together, leaving a total of four windows of interest. A depiction of all the time windows used for assessing the classification performance is reported in Figure 5A, with windows along the diagonal identified by letters (A–D) and off-diagonal windows identified by numbers (1–4). Classification accuracy across these time windows was averaged separately for each participant. For the two early time windows corresponding to the reference (Time Windows A and B in Figure 5A), the effect of inducer was tested using one-sample *t* tests against chance-level performance (50%). In the later time windows, we used paired *t* tests comparing the classification accuracy in the ping condition against the classification accuracy in the no-ping condition. Indeed, the effect of inducer at later time windows was expected only in the presence of pings, so that the no-ping condition served as a baseline for the ping condition.

Behavioral Data Analysis

Numerosity discrimination performance across the ping and no-ping trials was analyzed separately for each participant by assessing the serial dependence effect by binning trials according to the inducer numerosity. To obtain a measure of accuracy and precision in the task, we fitted a cumulative Gaussian curve to the distribution of response probabilities as a function of probe numerosity, according to the maximum likelihood method (Watson, 1979). The point of subjective equality (PSE), which corresponds to the probe numerosity perceptually matching (i.e., indistinguishable from) the reference numerosity, was taken as a measure of accuracy in the task and of the reference perceived numerosity. The PSE was

defined as the median of the best-fitting cumulative Gaussian curve to all the data of each participant in each condition. Each participant's performance was quantified by the just-noticeable difference (JND), which was calculated as the difference in numerosity between chance-level (50%) responses and 75% “probe more numerous” responses. A finger error rate correction (5%) was applied to account for lapses of attention or random response errors independent of the stimuli (Wichmann & Hill, 2001). To assess serial dependence effects at the behavioral level, a paired *t* test was performed comparing the distribution of PSEs corresponding to different inducer numerosity conditions separately for no-ping and ping trials to assess whether presenting the pings affected the effect. To compare the two conditions, we calculated a serial dependence effect index, which represents the difference in PSE between the 23-dot inducer condition and the 11-dot inducer condition. The distributions of serial dependence effect indexes across the two conditions were compared using a paired *t* test.

Exclusion Criteria

Because of the nature of the effect investigated in this work, we thoroughly checked both the behavioral and EEG data to exclude participants showing an insufficient level of performance and/or EEG data of insufficient quality. To this aim we excluded participants showing either very high JND in the behavioral task ($JND > 10$) or other indications that they may have not understood or complied with the task (i.e., opposite responses). Regarding the EEG data, participants were excluded during preprocessing either based on excessive number of rejected trials ($> 45\%$) or based on visual inspection of continuous data and/or ICA components. Seven participants were excluded before data analysis based on behavioral (4) and EEG (3) data quality. One more participant was excluded after data analysis due to the failure of the decoding procedure, which showed systematically below-chance decoding accuracy in the majority of the tested time points.

RESULTS

Behavioral Results

We analyzed the behavioral results to confirm that our specific design successfully induced attractive serial dependence biases separately for the no-ping and ping conditions. The reason to separate the two conditions is to assess whether reactivating the reference representation affected by serial dependence could modulate behavior by increasing or decreasing the effect. As shown in Figure 2, we found strong attractive serial dependence effects in both conditions. The perceived numerosity of a 16-dot reference appeared to be significantly underestimated when it was preceded by an 11-dot inducer, compared with when it was preceded by a 23-dot inducer

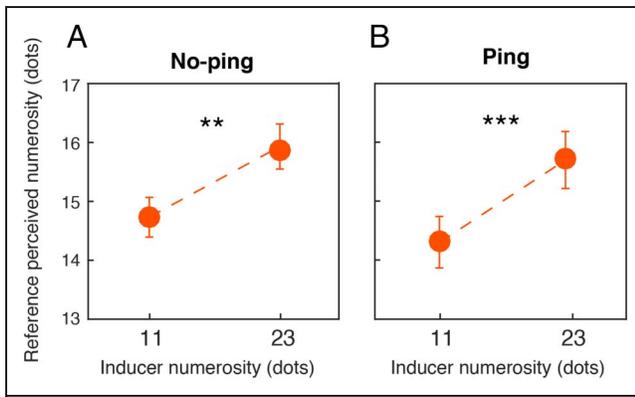


Figure 2. Behavioral results. (A) Average PSEs in the two inducer conditions in the no-ping trials. (B) Average PSEs in the ping trials. In both cases, strong and significant attractive serial dependence effects were observed. $**p < .01$, $***p < .001$. Error bars are *SEM*.

($t(26) = 3.3266$, $p = .0026$, $d = 0.64$, and $t(26) = 4.1796$, $p < .001$, $d = 0.80$, respectively, for the no-ping and ping conditions). Even if the effect seems slightly stronger in the ping condition (Figure 2B), comparing the serial dependence effects across the two conditions revealed no significant difference, $t(26) = 0.4076$, $p = .6869$. This shows that, the reference was strongly biased by the inducer in our paradigm, the presence of the pings did not systematically modulate that effect. This is in line with a previous observation showing that reactivating a representation by means of pings does not affect behavior (Wolff et al., 2017).

ERP Results

After confirming that our paradigm successfully induced serial dependence in numerosity perception, we analyzed the brain responses to the reference stimulus. First, we looked at the topographic distribution of brain activity, collapsing no-ping and ping trials together, in a relatively early time window following previous studies (Fornaciai & Park, 2018a). Figure 3 shows the distribution of activity representing the contrast between the two inducer conditions across posterior channels. There was a relatively

large difference in the responses to the same 16-dot reference as a function of the inducer numerosity, emerging at relatively early latencies over posterior–inferior scalp locations. Such a difference—reflecting higher positive amplitude in the more numerous inducer condition—peaked at around 250–300 msec after stimulus onset, with activity centered at channel Iz (peak = 261 msec, contrast amplitude = 0.484 μV).

Then, to characterize the temporal effect of serial dependence on the ERPs, we plotted and analyzed the brainwaves at channel Iz, which appeared as the center of the occipital focus of activity. The brainwaves corresponding to different stimuli and latency windows are shown in Figure 4. Note that, for the purpose of ERP analysis at latencies around the reference onset, we collapsed together the ping and no-ping trials, as no difference is expected between the two conditions. Relatively early ERPs evoked by the reference stimulus (Figure 4A) showed a clear attractive modulation, peaking at around 240 msec after stimulus onset. A larger inducer numerosity (23 dots) resulted in a more positive amplitude of the brainwaves evoked by the reference, compared with the less numerous inducer (11 dots), which was consistent with what we observed at the behavioral level. The ERP effect of serial dependence was quantified in a 100-msec time window around 200 msec after reference onset (marked with a shaded area in Figure 4A) based on our previous finding (Fornaciai & Park, 2018a) and was tested using a paired t test (between the two inducer conditions across participants). The results showed a significant difference between ERPs corresponding to different inducer numerosities, $t(26) = 3.434$, $p < .001$.

Figure 4B, C shows the brainwaves corresponding to the two pings presented between the reference and the probe in the ping condition. As illustrated, the bright flashes evoked a strong visual response starting early after stimulus onset. However, the difference between different inducer conditions appeared to be very small at the time of the first ping while slightly stronger at the time of the second ping. A 100-msec time window around the peak of the activity of the inducer conditions collapsed together (peak = 125 and 137 msec after the onset of the first and second ping, respectively) did not

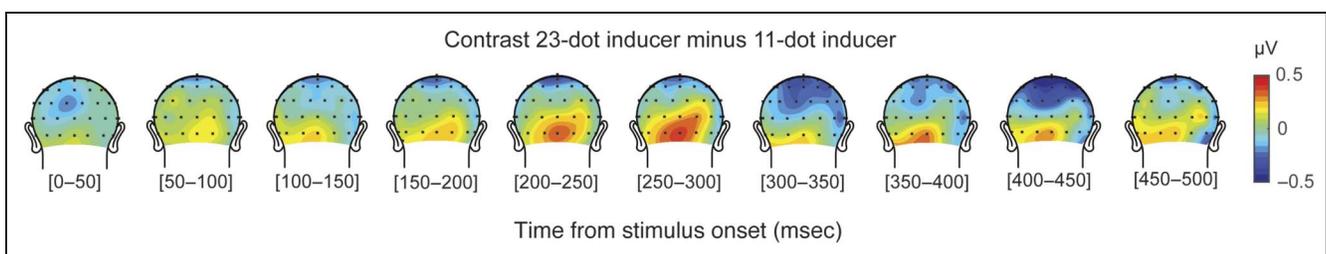


Figure 3. Topographic distribution of ERPs across posterior channels, corresponding to the contrast between the 23-dot inducer condition and the 11-dot inducer condition. This topomap reveals a relatively strong effect of inducer starting from early latencies and peaking at around 250–300 msec after stimulus onset. The most involved channels appear to be a group of posterior channels centered around Iz.

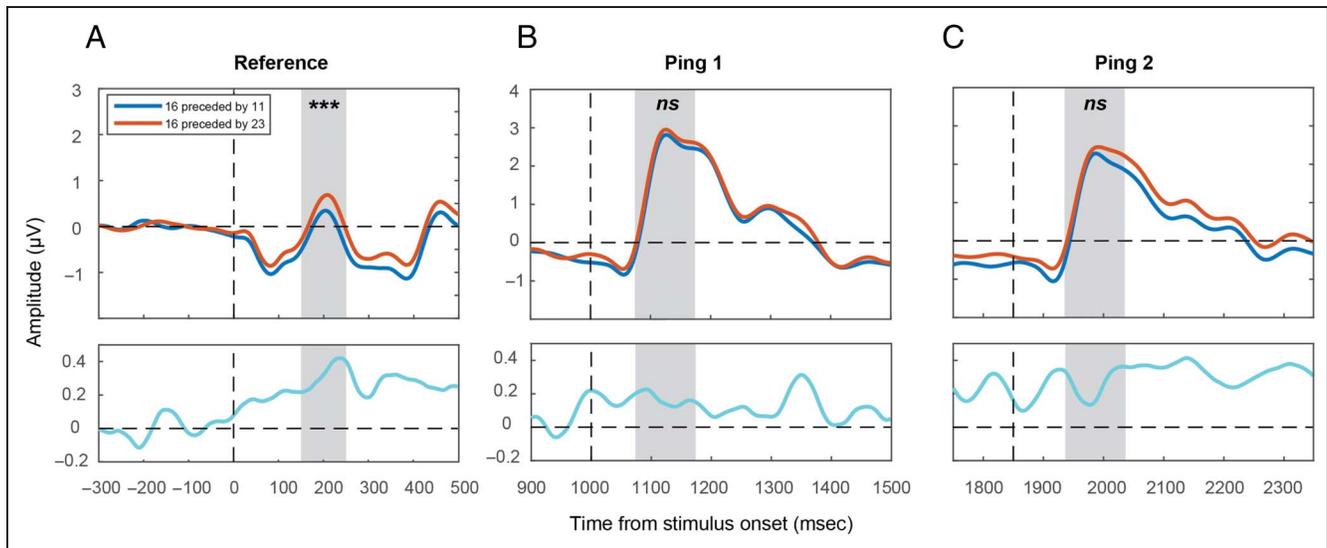


Figure 4. ERPs at channel Iz. (A) ERPs time-locked to the onset of the reference stimulus (collapsing together ping and no-ping trials), sorted by the numerosity of the preceding inducer stimulus (top) and the contrast between the two waveforms (bottom). (B) ERPs time-locked to the onset of the first ping in the ping condition (top) and their contrast (bottom). (C) ERPs time-locked to the onset of the second ping in the ping condition (top) and their contrast (bottom). Shaded areas represent the time windows used for a statistical test. In the case of brainwaves corresponding to the reference stimulus (A), the time window of interest was defined based on previous studies (see Methods). In the case of the activity evoked by the pings (B, C), time windows were chosen according to the peak of the two brainwaves collapsed together. *ns* = not significant. *** $p < .001$.

show any significant difference between the two inducer conditions, neither after the first ping nor after the second ping (paired t test, $t(26) = 0.9467$, $p = .1762$, and $t(26) = 1.1617$, $p = .1280$).

Multivariate Neural Decoding Results

To assess whether the initial neural representation of the reference numerosity is retained in the memory traces (which are reactivated by the pings), we performed a multivariate pattern analysis in the time domain (e.g., see King & Dehaene, 2014, for a review). Although ERP analysis can show a signature of serial dependence, it is not sensitive enough to pick up subtle differences in the stimulus representation and does not provide a full characterization of the time course of the effect. A multivariate analysis, instead, has been shown to be much more sensitive to serial dependence biases and allows a better characterization of the time course of the effect (Fornaciai & Park, 2018a). Such a multivariate analysis involves training and testing a pattern classifier (i.e., SVM) to assess whether and to what extent it can predict the pattern of brain activity related to a particular condition. Moreover, assessing the temporal generalization of specific patterns of brain activity across different time windows allows us to assess how the—biased—representation of the reference stimulus evolves over time during the course of a trial. More specifically, we compared the activity time-locked to the onset of the reference, dividing the trials according to the preceding inducer stimulus, resulting in two categories: 16 preceded by 11 and 16

preceded by 23. The analysis was performed separately for no-ping and ping trials, and the no-ping condition was used as a comparison to assess the effect related to the pings.

Figure 5A illustrates the results of the neural decoding analysis in terms of a temporal generalization matrix, which characterizes the degree to which patterns of neural activity evoked by the identical reference numerosity (16 dots) and the two pings can be dissociated as a function of two inducer numerosities (11 or 23 dots). The results are quantified in various time windows of interest. We first set two time windows corresponding to 0–250 (labeled as A in Figure 5A) and 250–500 msec (labeled as B) after reference onset. Across these two time windows, we found significantly above chance decoding in the earlier time window (0–250 msec, average classification accuracy = 0.5291; one-sample t test against chance level, $t(26) = 3.0669$, $p = .005$), but not in the later time window (250–500 msec, average classification accuracy = 0.5069, $t(26) = 0.5918$, $p = .5591$). However, although the decoding accuracy in the later window (B) was not significant, it was much stronger in the off-diagonal directions adjacent to such a later window (i.e., the areas simultaneously adjacent to Windows A and B). We thus also tested the average decoding performance in these two windows. Doing so, we indeed found above chance classification accuracy ($t(26) = 2.26$, $p = .033$, and $t(26) = 2.44$, $p = .022$, respectively for the off-diagonal windows adjacent to the reference window in the vertical and horizontal direction), suggesting that activity from the earlier window generalizes to later time points

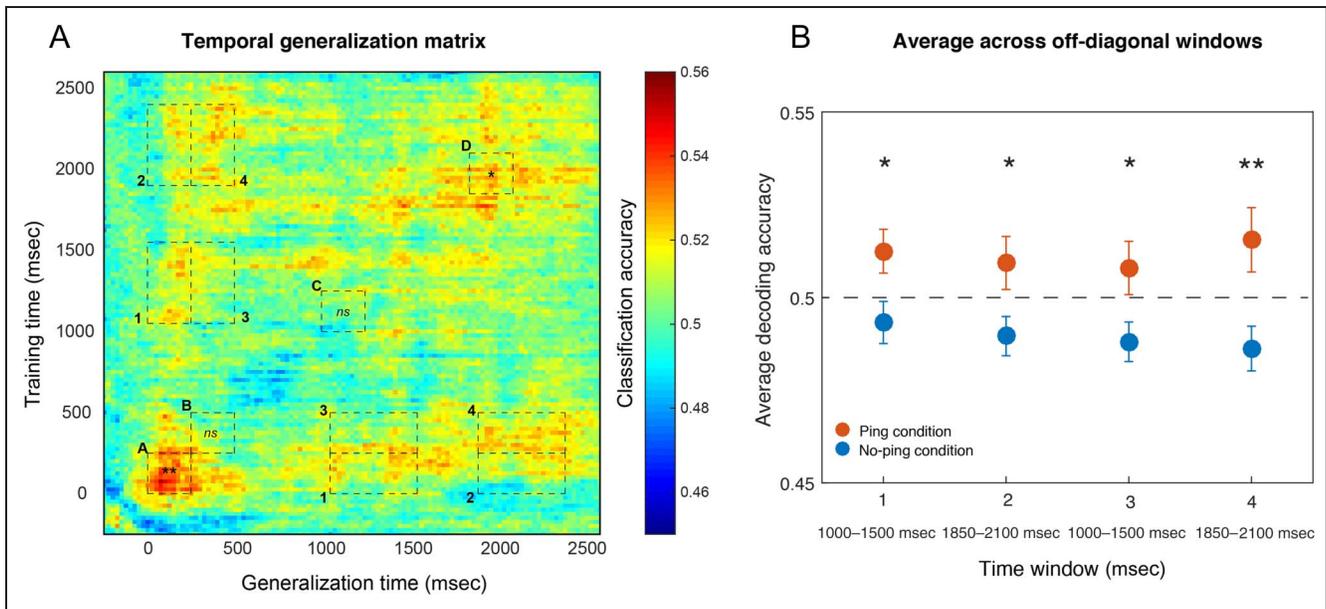


Figure 5. Decoding results. (A) Temporal generalization matrix, representing classification accuracies obtained by training the classifier at one time point and testing at all the time points. The temporal generalization shows how the specific pattern of brain activity at one time point generalizes at different time points, possibly showing the reactivation of the same representation and/or the same neural generator at different times. The dashed lines delimit the different temporal windows of interest used to perform statistical tests on the classification performance. Time windows of interest along the diagonal are marked with letters, whereas off-diagonal time windows are labeled with numbers. (B) Average classification accuracy across the off-diagonal windows, for the ping and no-ping conditions. Error bars are *SEM*. Number labels on the x axis refer to the numbering of time windows reported in A. *ns* = not significant. **p* < .05, ***p* < .01.

corresponding to the late reference window. We then examined the time windows 1000–1250 and 1850–2100 msec, which correspond to a 250-msec window after the first and second ping, respectively. The decoding classification accuracy after the first ping (labeled as C in Figure 5A) was negligible (average classification accuracy = 0.5036, paired *t* test against the no-ping condition: $t(26) = 0.5027, p = .6194$); however, the classification accuracy after the second ping (labeled as D) was significantly higher than that in the absence of a ping stimulus (average classification accuracy = 0.5183, paired *t* test against the no-ping condition: $t(26) = 2.6667, p = .013$).

More importantly, the extent to which the neural patterns are generalized across time at later time windows (e.g., from reference to pings) was evaluated by observing off-diagonal time windows in the temporal generalization matrix. Generalized patterns across the off-diagonal directions would indeed show whether the biased representation of the reference stimulus exists in the perceptual memory traces after the initial processing of it. To this aim, we focused on broad 500-msec time windows corresponding to the reference time windows but shifted in time to the onset of the two pings (time windows labeled 1, 2, 3, and 4 in Figure 5A) to capture the time windows in which the neural decoding classifier is trained with the neural activity pattern evoked by the reference numerosity and tested with the pattern evoked by the pings (and vice versa). In other words, these off-diagonal

time windows allowed us to evaluate whether the pattern of serial dependence effects observed at the time of reference processing generalizes to activity evoked by the pings. As the temporal generalization matrix (Figure 5A) is two-dimensional, there are two directions in which we could look for the generalization of patterns of brain activity: “horizontal” and “vertical.” Corresponding time windows along the two directions were collapsed together for statistical analyses.

Comparisons concerning off-diagonal windows (ping vs. no ping condition) are shown in Figure 5B. As shown, average classification accuracies obtained in the ping condition were greater than those in the no-ping condition. The inducer-dependent pattern of neural activation evoked by the reference in the 0–250 msec time window was decodable by the neural activation pattern evoked by the first ping (paired *t* test of ping vs. no-ping condition), $t(26) = 2.10$, one-sided $p = .0228$, and by the pattern evoked by the second ping, $t(26) = 2.24, p = .0171$ (see the time windows labeled 1 and 2 in Figure 5). Similarly, the inducer-dependent pattern of neural activation evoked by the reference in the 250–500 msec time window was decodable by the pattern evoked by the first ping, $t(26) = 2.16, p = .0202$, and by the pattern evoked by the second ping, $t(26) = 2.57, p = .0081$ (see the time windows labeled 3 and 4 in Figure 5). In other words, the same differences between the neural activity patterns evoked by the reference under the two inducer

conditions could be used to distinguish the activity patterns evoked by the pings under the two inducer conditions.

These results provide evidence for the reactivation of a stored representation of the initial reference stimulus. Although these results show a significant difference between ping and no-ping activities, the average classification accuracy in the no-ping condition appeared to be below the theoretical chance level of 50%. Although puzzling at first glance, below-chance classification results have been interpreted in terms of dependence from subsample means in the context of cross-validation procedures (Jamalabadi, Alizadeh, Schönauer, Leibold, & Gais, 2016). Namely, in the presence of a small effect size and a relatively small number of observations (trials), pulling out the subset of data used for testing the classifier from the training set during cross-validation could shift the means of the remaining training set to the opposite direction, making the trained classifier to predict the test set to be the other class. This, in turn, results in a classification error and hence in below-chance classification accuracy. The possibility of these spurious effects thus makes it very important to use the “baseline” classification accuracy obtained in the no-ping condition to assess the effect provided by the pings, instead of using a fixed 50% chance level. Nevertheless, we performed an additional series of one-sample *t* tests against the theoretical chance level of 50%. The results were significant in Window 1, $t(26) = 2.12$, $p = .022$, and Window 4, $t(26) = 1.79$, $p = .043$.

DISCUSSION

In this study, we investigated the neural signature and dynamics of serial dependence across the time course of perceptual decision-making. Specifically, we questioned whether the biased representation of the visual stimulus (due to serial dependence) is transient or whether it is retained in visual memory traces. To do so, we exploited a recently developed technique aimed to reactivate hidden memory representations (Wolff et al., 2015, 2017). Namely, the neural signal evoked by a brief, high-energy “ping” reflects the signature of underlying representations currently encoded in the neural network (i.e., for instance stored in working memory). Using a relatively long time course in an active task, incorporating this pingging technique in combination with a neural decoding analysis, we characterized how serial dependence emerges and evolves throughout the time course of a trial.

Our results first show that the specific paradigm used in this study successfully induces strong and systematic attractive serial dependence, virtually identical to a previous study employing a similar paradigm but in a faster rate of stimulus presentation (Fornaciai & Park, 2018b). Specifically, we observed a relative under- or overesti-

mation of the reference stimulus as a function of the inducer numerosity, in line with the attractive effect observed in previous studies (Fornaciai & Park, 2018a, 2018b, 2019a, 2019b; Cicchini et al., 2014; Corbett et al., 2011).

As in the behavioral results, the ERPs evoked by the reference stimulus were modulated by the numerosity of the preceding stimulus, showing a signature of serial dependence consistent with our previous study (Fornaciai & Park, 2018a). Such a signature likely reflects a distorted encoding of stimulus numerosity according to the numerosity of the preceding inducer stimulus, with a clear attractive pattern—that is, the amplitude of brain signals evoked by the reference was higher (more positive) following an inducer containing a larger numerosity, compared with an inducer containing fewer items. Such an effect peaked at around 200–250 msec, consistently with recent studies investigating the dynamics of numerosity processing (i.e., Fornaciai & Park, 2018c; Fornaciai et al., 2017; Park et al., 2016; Hyde & Spelke, 2009) and with the timing of serial dependence in our previous study (Fornaciai & Park, 2018a). ERPs evoked by the two pings, on the other hand, showed little or no effect (i.e., did not reliably distinguish the two inducer conditions), despite the large visually evoked potentials induced by the bright flashes (Figure 4B and C).

To test our central hypothesis concerning the temporal evolution of neural representation of these numerosity stimuli, we employed a multivariate neural decoding analysis, which has previously proven to be sensitive in distinguishing subtle experimental manipulations (Fornaciai & Park, 2018a, 2018c). Such an analysis allowed us to assess the pattern of brain activity corresponding to a specific stimulation condition and how such a pattern can successfully predict the stimulus presented in another specific condition. Looking at a brief epoch around the onset of the reference stimulus (Figure 5A, Time Windows A and B), it is clear that a signature of serial dependence emerges very early after stimulus onset, consistent with previous results (Fornaciai & Park, 2018a). Comparing the present results with our previous study, however, the decoding accuracy remained high only at relatively early latencies after the reference onset (0–250 msec), whereas our previous data showed that the effect is amplified at later latencies. Instead, in the current data, there was an increase in decoding accuracy in the off-diagonal direction (see the areas simultaneously adjacent to Windows A and B in Figure 5A), illustrating the generalization of the classification across time points close to the reference representation. This result suggests that the pattern of activity corresponding to the initial (0–250 msec) processing stage (and likely the same neural generator) remains sustained at subsequent latencies up until 500 msec after stimulus onset. In other words, although in our previous work (Fornaciai & Park, 2018a) we observed a series of two potentially distinct processing stages rapidly unfolding after stimulus

presentation (0–200 and 200–450 msec), here we observed just one of them spanning the first 250 msec and then lingering up until 500 msec. Such a difference between the current and the previous results may be due to the differences in the paradigms employed. The two paradigms differed in terms of stimulus timing (much longer ISI in the present work), task (oddball detection vs. numerosity discrimination), and the attended feature (color of the stimuli in the oddball task vs. numerosity in the present task), which by themselves can modulate the brain responses recorded with EEG (e.g., Koivisto, Kainulainen, & Revonsuo, 2009; Zhang & Luck, 2009; Nakata et al., 2005; Potts, 2004; Senkowski & Herrmann, 2002; Strüber & Polich, 2002; Zani & Proverbio, 1995). In particular, a longer ISI can significantly delay evoked responses linked to stimulus processing (Nakata et al., 2005) or even affect the pattern of responses evoked by a stimulus (Strüber & Polich, 2002). Although previous results on the effect of ISI concern very different stimuli and tasks, some of the differences may be driven by the much slower paradigm employed in this study as opposed to the rapid stimulus presentation in Fornaciai and Park (2018a). In other words, the slower presentation rate in the current paradigm may have made activity at subsequent processing stages (for instance related to the reactivation of working memory traces to perform the task) to be delayed at later latencies closer to the presentation of the second stimulus. However, due to the large differences in both the stimulus dynamics and the task used, it is difficult to precisely pinpoint the reason for such different results.

Similarly to the ERP analysis, the first ping was ineffective in distinguishing the two inducer conditions (Figure 5A, Time Window C). This lack of significant decoding along the diagonal shows that, at least for the first ping, the stimulation did not evoke any distinctive pattern of activity carrying the serial dependence effect. In contrast, the second ping did result in successful decoding (Figure 5A, at Time Window D), suggesting that the pattern of activity evoked by this second ping carries information about the reference affected by the serial dependence effect. However, the pattern of effects seems to spread across a relatively broad area around the ping onset. The fact that we observe a signature of serial dependence only at the time of the second (but not the first) ping and with activity spread around this presentation time can be interpreted in several ways. First, predictability and expectations concerning the timing of the second ping and the end of the trial may explain the spread of the effect to time points even before the ping presentation and a stronger effect of the second ping. Indeed, the timing of the pings was consistent across trials, which made it possible to predict the presentation of the second ping after seeing the first one; in turn, those expectations may have caused a sharpening of the pattern of brain activity (e.g., see de Lange, Heilbron, & Kok, 2018, for a review), making the underlying serial dependence effect more easily decodable.

Second, the effect emerging exclusively at the second ping may reflect the increase of the serial dependence effect over time observed in previous studies (Bliss et al., 2017; Fritsche et al., 2017), although the magnitude of the effect in our behavioral results is comparable to previous results employing much shorter presentation times (Fornaciai & Park, 2018b). Additionally, increased attention toward the end of the trial may have played a role in amplifying the responses to the second ping (e.g., Müller et al., 2006), resulting in stronger signals related to the serial dependence effect. Finally, considering that the second ping is closer to the end of the trial, neural activity evoked by the second ping may reflect a decision stage representation (i.e., the activity of perceptual decision circuits) in line with a recently proposed framework of serial dependence (Pascucci et al., 2019). According to this proposition, serial dependence would emerge because of “perceptual decision” or “read-out” templates persisting from the processing of past stimuli to the representation of current ones. Such decision templates are for instance represented by the set of read-out weights of low-level activity at a higher order decision stage, with the lingering template from a previous stimulus biasing the representation of a current one. In the context of the present results, the activity found at the time of the second ping may thus reflect activity arising at such perceptual decision stage, related to the reactivated decision template. Following the previous points, the fact that we observe this effect only at the time of the second ping may be related to predictions and expectations about the timing of the incoming probe stimulus. Indeed, the predictability of the stimulus timing may have induced an anticipated activation of the perceptual decision circuits to speed up and facilitate the representation of the incoming probe stimulus, in line with the idea that expectations can speed up perception (e.g., Pinto, van Gaal, de Lange, Lamme, & Seth, 2015).

The crucial point for testing our main hypothesis, however, concerns the “reactivation” of the reference representation induced by the pings. Such a reactivation has to be searched in the off-diagonal direction in the temporal generalization matrix (Figure 5A). Our results show that the pings successfully reactivate the reference representation, carrying the signature of serial dependence. Interestingly, significant effects emerge not only at off-diagonal times corresponding to the early window of the reference stimulus (off-diagonal Windows 1 and 2, which are aligned with the early Window A in Figure 5A), but also at times corresponding to the later window (off-diagonal Windows 3 and 4, which are aligned with the early Window B in Figure 5A), where there is actually little or no effect at the time of reference processing along the diagonal. Even if activity in the later window after reference onset provides a relatively poor classification of inducer conditions, the effect is amplified at later latencies, especially when such a pattern gets reactivated by the pings.

In general, however, caution is in order when interpreting these results, as our main decoding analysis shows relatively low classification accuracies, especially at off-diagonal windows corresponding to activity reactivated by the pings. Considering the nature of the effect tested and the specific analysis used, such low classification accuracies, however, are not surprising. First, we need to consider that serial dependence is a relatively subtle behavioral effect, causing in our specific case an under-/overestimation of about 1 dot in a dot array of 16 dots. It is reasonable to hypothesize that any difference at the neural level would be more subtle or at best similar. Moreover, we used large temporal windows to average classification accuracies at different time points, especially in the off-diagonal direction, which may have diluted the effect. Future studies with more specific temporal predictions may achieve better results by using narrower windows. Nevertheless, in the absence of the pings the serial dependence effect is not decodable from neural signals at those later latencies, suggesting that the effects evoked by the pings is genuine rather than an artifact.

In summary, the current findings suggest that there are multiple patterns of neural activity involved in the serial dependence effect. First, consistently with previous reports (Fornaciai & Park, 2018a), the effect of serial dependence develops during a relatively short interval immediately after the onset of the stimulus, suggesting that it may arise directly at the earliest level of perception. Regarding this point, however, the fact that the decoding analysis does not provide information about the direction of the effect (i.e., attractive vs. repulsive) leaves open the possibility that the very early activity may be driven by other sensory/perceptual processes, like adaptation. Adaptation resulting from very short visual stimulation has been indeed observed when serial dependence is suppressed by means of visual backward masking (Fornaciai & Park, 2019a). However, the fact that ERP results show an attractive effect within the early window of activity (around 200 msec after stimulus onset) suggests that the decoding results within this same window may similarly reflect an attractive effect.

After the relatively early processing, then, the effect is evident in possibly two distinct processes. On the one hand, the off-diagonal pattern of activity indicates that—borrowing the interpretation of King and Dehaene (2014)—the same neural generators yield the pattern of activation for the initial biased representation and the pattern of activation evoked by the pings. One parsimonious interpretation is that the exact same trace of the initial biased representation gets stored throughout a trial (and is reactivated by the pings later). On the other hand, the diagonal pattern of activity especially around and after the second ping indicates—again according to King and Dehaene (2014)—a distinct set of neural generators decoding the two inducer conditions. The functional role of this second mechanism remains to

be studied, as we can only speculate as discussed above. Nevertheless, both diagonal and off-diagonal patterns found in the neural decoding analyses lend support for at least two distinct processing stages involved in the representation of serial dependence in the time course of a trial, which serves as an important starting point for future research.

Conclusion

The present results advance our knowledge about the neural underpinnings of the serial dependence by showing the temporal dynamics of its neural signature over a relatively long time interval. We show that serial dependence starts very early after stimulus onset, similarly to what has been demonstrated before with a passive-viewing paradigm. We further show that the representation of a stimulus affected by serial dependence can be reactivated by task-irrelevant pings. The effect of pinging memory has two results: On the one hand, it reactivates a memory representation that remains similar to the stimulus representation generated immediately after its onset. This suggests that once the bias is induced during perceptual processing, a trace of the biased stimulus representation is stored throughout a relatively long period. On the other hand, it highlights a different, later processing stage, which likely reflects the involvement of possibly a different neural generator yielding a distinctive pattern of activity but still carrying the signature of serial dependence. The present results thus show that serial dependence directly alters the perceptual representation of a stimulus and propagates throughout subsequent processing stages.

Acknowledgments

We thank Jihyun Hwang, Camryn Boutin, and Brianna Chamberlain for their assistance in EEG data collection. This study was supported by National Science Foundation CAREER award BCS1654089 to J. P. M. F. and J. P. designed the study. M. F. collected the data. M. F. and J. P. analyzed the data, interpreted the results, and wrote and revised the manuscript.

Reprint requests should be sent to Michele Fornaciai, International School for Advanced Studies (SISSA), Via Bonomea, 265, 34136 Trieste (TS), Italy, or via e-mail: michele.fornaciai@gmail.com.

REFERENCES

- Alais, D., Leung, J., & Van der Burg, E. (2017). Linear summation of repulsive and attractive serial dependencies: Orientation and motion dependencies sum in motion perception. *Journal of Neuroscience*, *37*, 4381–4390.
- Bliss, D. P., Sun, J. J., & D'Esposito, M. (2017). Serial dependence is absent at the time of perception but increases in visual working memory. *Scientific Reports*, *7*, 14739.
- Burr, D., & Cicchini, G. M. (2014). Vision: Efficient adaptive coding. *Current Biology*, *24*, R1096–R1098.

- Cicchini, G. M., Anobile, G., & Burr, D. C. (2014). Compressive mapping of number to space reflects dynamic encoding mechanisms, not static logarithmic transform. *Proceedings of the National Academy of Sciences, U.S.A.*, *111*, 7867–7872.
- Cicchini, G. M., Mikellidou, K., & Burr, D. (2017). Serial dependencies act directly on perception. *Journal of Vision*, *17*, 6.
- Cicchini, G. M., Mikellidou, K., & Burr, D. C. (2018). The functional role of serial dependence. *Proceedings of the Royal Society B: Biological Sciences*, *285*, 20181722.
- Corbett, J. E., Fischer, J., & Whitney, D. (2011). Facilitating stable representations: Serial dependence in vision. *PLoS One*, *6*, e16701.
- de Lange, F. P., Heilbron, M., & Kok, P. (2018). How do expectations shape perception? *Trends in Cognitive Sciences*, *22*, 764–779.
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*, 9–21.
- DeWind, N. K., Adams, G. K., Platt, M. L., & Brannon, E. M. (2015). Modeling the approximate number system to quantify the contribution of visual stimulus features. *Cognition*, *142*, 247–265.
- Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, *17*, 738–743.
- Fornaciai, M., Brannon, E. M., Woldorff, M. G., & Park, J. (2017). Numerosity processing in early visual cortex. *Neuroimage*, *157*, 429–438.
- Fornaciai, M., Farrell, A., & Park, J. (2019). Looking for more food or more people? Task context influences basic numerosity perception. *Cortex*, *114*, 67–75.
- Fornaciai, M., & Park, J. (2017). Distinct neural signatures for very small and very large numerosities. *Frontiers in Human Neuroscience*, *11*, 21.
- Fornaciai, M., & Park, J. (2018a). Attractive serial dependence in the absence of an explicit task. *Psychological Science*, *29*, 437–446.
- Fornaciai, M., & Park, J. (2018b). Serial dependence in numerosity perception. *Journal of Vision*, *18*, 15.
- Fornaciai, M., & Park, J. (2018c). Early numerosity encoding in visual cortex is not sufficient for the representation of numerical magnitude. *Journal of Cognitive Neuroscience*, *30*, 1788–1802.
- Fornaciai, M., & Park, J. (2019a). Spontaneous repulsive adaptation in the absence of attractive serial dependence. *Journal of Vision*, *19*, 21.
- Fornaciai, M., & Park, J. (2019b). Serial dependence generalizes across different stimulus formats, but not different sensory modalities. *Vision Research*, *160*, 108–115.
- Fritsche, M., Mostert, P., & de Lange, F. P. (2017). Opposite effects of recent history on perception and decision. *Current Biology*, *27*, 590–595.
- Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2017). Decoding dynamic brain patterns from evoked responses: A tutorial on multivariate pattern analysis applied to time series neuroimaging data. *Journal of Cognitive Neuroscience*, *29*, 677–697.
- Hyde, D. C., & Spelke, E. S. (2009). All numbers are not equal: An electrophysiological investigation of small and large number representations. *Journal of Cognitive Neuroscience*, *21*, 1039–1053.
- Jamalabadi, H., Alizadeh, S., Schönauer, M., Leibold, C., & Gais, S. (2016). Classification based hypothesis testing in neuroscience: Below-chance level classification rates and overlooked statistical properties of linear parametric classifiers. *Human Brain Mapping*, *37*, 1842–1855.
- King, J.-R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: The temporal generalization method. *Trends in Cognitive Sciences*, *18*, 203–210.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, *36*, 1–16.
- Koivisto, M., Kainulainen, P., & Revonsuo, A. (2009). The relationship between awareness and attention: Evidence from ERP responses. *Neuropsychologia*, *47*, 2891–2899.
- Lieberman, A., Fischer, J., & Whitney, D. (2014). Serial dependence in the perception of faces. *Current Biology*, *24*, 2569–2574.
- Lieberman, A., Manassi, M., & Whitney, D. (2018). Serial dependence promotes the stability of perceived emotional expression depending on face similarity. *Attention, Perception, & Psychophysics*, *80*, 1461–1473.
- Lopez-Calderon, J., & Luck, S. J. (2014). ERPLAB: An open-source toolbox for the analysis of event-related potentials. *Frontiers in Human Neuroscience*, *8*, 213.
- Manassi, M., Lieberman, A., Chaney, W., & Whitney, D. (2017). The perceived stability of scenes: Serial dependence in ensemble representations. *Scientific Reports*, *7*, 1971.
- Manassi, M., Lieberman, A., Kosovicheva, A., Zhang, K., & Whitney, D. (2018). Serial dependence in position occurs at the time of perception. *Psychonomic Bulletin & Review*, *25*, 2245–2253.
- Meyers, E. M. (2013). The neural decoding toolbox. *Frontiers in Neuroinformatics*, *7*, 8.
- Müller, M. M., Andersen, S., Trujillo, N. J., Valdés-Sosa, P., Malinowski, P., & Hillyard, S. A. (2006). Feature-selective attention enhances color signals in early visual areas of the human brain. *Proceedings of the National Academy of Sciences, U.S.A.*, *103*, 14250–14254.
- Nakata, H., Inui, K., Wasaka, T., Tamura, Y., Kida, T., & Kakigi, R. (2005). Effects of ISI and stimulus probability on event-related go/no-go potentials after somatosensory stimulation. *Experimental Brain Research*, *162*, 293–299.
- Park, J., DeWind, N. K., Woldorff, M. G., & Brannon, E. M. (2016). Rapid and direct encoding of numerosity in the visual stream. *Cerebral Cortex*, *26*, 748–763.
- Pascucci, D., Mancuso, G., Santandrea, E., Della Libera, C., Plomp, G., & Chelazzi, L. (2019). Laws of concatenated perception: Vision goes for novelty, decisions for perseverance. *PLoS Biology*, *17*, e3000144.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Pinto, Y., van Gaal, S., de Lange, F. P., Lamme, V. A. F., & Seth, A. K. (2015). Expectations accelerate entry of visual stimuli into awareness. *Journal of Vision*, *15*, 13.
- Potts, G. F. (2004). An ERP index of task relevance evaluation of visual stimuli. *Brain and Cognition*, *56*, 5–13.
- Senkowski, D., & Herrmann, C. S. (2002). Effects of task difficulty on evoked gamma activity and ERPs in a visual discrimination task. *Clinical Neurophysiology*, *113*, 1742–1753.
- St. John-Saaltink, E., Kok, P., Lau, H. C., & de Lange, F. P. (2016). Serial dependence in perceptual decisions is reflected in activity patterns in primary visual cortex. *Journal of Neuroscience*, *36*, 6186–6192.
- Strüber, D., & Polich, J. (2002). P300 and slow wave from oddball and single-stimulus visual tasks: Inter-stimulus interval effects. *International Journal of Psychophysiology*, *45*, 187–196.
- Suárez-Pinilla, M., Seth, A. K., & Roseboom, W. (2018). Serial dependence in the perception of visual variance. *Journal of Vision*, *18*, 4.

- Taubert, J., Van der Burg, E., & Alais, D. (2016). Love at second sight: Sequential dependence of facial attractiveness in an on-line dating paradigm. *Scientific Reports*, *6*, 22740.
- Watson, A. B. (1979). Probability summation over time. *Vision Research*, *19*, 515–522.
- Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics*, *63*, 1293–1313.
- Wolff, M. J., Ding, J., Myers, N. E., & Stokes, M. G. (2015). Revealing hidden states in visual working memory using electroencephalography. *Frontiers in Systems Neuroscience*, *9*, 123.
- Wolff, M. J., Jochim, J., Akyürek, E. G., & Stokes, M. G. (2017). Dynamic hidden states underlying working-memory-guided behavior. *Nature Neuroscience*, *20*, 864–871.
- Xia, Y., Leib, A. Y., & Whitney, D. (2016). Serial dependence in the perception of attractiveness. *Journal of Vision*, *16*, 28.
- Zani, A., & Proverbio, A. M. (1995). ERP signs of early selective attention effects to check size. *Electroencephalography and Clinical Neurophysiology*, *95*, 277–292.
- Zhang, W., & Luck, S. J. (2009). Feature-based attention modulates feedforward visual processing. *Nature Neuroscience*, *12*, 24–25.