

Serial dependence alters perceived object appearance

Thérèse Collins

Integrative Neuroscience and Cognition Center,
Université de Paris and CNRS, Paris, France



The visual world as it presents itself to our eyes is constantly changing, in contrast with human perceptual experience, which is smooth and stable. One of the posited psychological mechanisms that may contribute to this constructed perceptual stability is the continuity field, a spatiotemporal integration window. The current study examined whether temporal integration, as quantified by serial dependence (SD) between perceived attributes of successive visual stimuli, influenced the subjective appearance of objects or decisional stages in response determination. To do so, an oddball task required participants to directly compare visual objects and decorrelated responses (present/absent) from the visual attribute on which SD may occur (orientation). Results showed that SD could cause a single visual object to appear different from surrounding distractors, leading to modulations of performance. These results argue in favor of a perceptual level of SD, and against decisional accounts.

Introduction

The visual world as it presents itself to our eyes is highly variable: There are multiple sources of external noise (e.g., objects reflect different wavelength compositions depending on lighting conditions) and internal noise (e.g., blinks and eye movements interrupt and displace the retinal image several times per second). Despite this highly variable input, perceptual experience is remarkably stable. Objects do not seem to change positions every time we move our eyes, or to change color each time the sun moves behind a cloud. A mechanism that may contribute to seamless perception is the continuity field, a spatiotemporal window within which sensory/perceptual information is integrated to give rise to the immediate contents of perception. The temporal aspect of the continuity field has been supported empirically by the behavioral phenomenon of serial dependence (SD): When reporting the perceived nature of a visual stimulus (e.g., the orientation of a Gabor patch), participants' responses err toward the previously seen stimulus (Cicchini, Anobile, & Burr, 2014; Fischer & Whitney, 2014). Of course, responses are primarily dependent on the stimulus itself, but

Citation: Collins, T. (2020). Serial dependence alters perceived object appearance. *Journal of Vision*, 20(13):9, 1–8, <https://doi.org/10.1167/jov.20.13.9>.

response errors are not random, suggesting that the immediate content of perception integrates both current input and previous stimulus history. Stimuli from a few seconds ago have an attractive influence, whereas more distant stimuli may exert a repulsive influence (Fritsche, Spaak, & de Lange, 2020; Gekas, McDermott, & Mamassian, 2019). Attractive biases allow the visual system to capitalize on the fact that the visual environment is, in the short-term, rather stable (Dong & Atick, 1995; Schwartz, Hsu, & Dayan, 2007; Simoncelli & Olshausen, 2001).

If SD is a fundamental determinant of immediate perceptual content and evidence that biological perceptual systems exploit temporal continuity to promote visual stability, then the level at which SD arises is of crucial theoretical import: At the level of perception or response determination. The current empirical support for early perceptual or decisional SD is divided.

Most studies on SD use a reproduction task in which participants view a stimulus such as a Gabor patch and report its orientation by adjusting a response cue. This task does not differentiate between perceptual versus decisional levels. A more direct measure of perception is to ask participants to compare two simultaneous stimuli, only one of which has been preceded by another Gabor patch likely to induce SD (Cicchini, Mikellidou, & Burr, 2017; Fischer & Whitney, 2014; Fritsche, Mostert, & de Lange, 2017). Fritsche et al. (2017) showed that there was a repulsion of the current percept from the previous stimulus, and thus that attractive SD must occur at the decisional stage rather than at the level of perception. Cicchini et al. extended Fritsche et al.'s methods to cover a wider range of inducer-stimulus differences and showed that when orientations were similar, perceptual SD was attractive, and that a small repulsive effect occurred only for larger differences (nonsignificantly in Cicchini et al., despite an effect size similar to Fritsche et al.).

Another method for examining whether SD is perceptual or decisional is to look at its spatial specificity, by varying the locations of successive stimuli. Early, perceptual processing should be more spatially restricted than later, response-related processing, owing to the retinotopic organization of human visual cortex (Wandell, Dumoulin, & Brewer, 2007). Behavioral



reports have shown that SD is retinotopic (Collins, 2019). Furthermore, the orientation signal in V1 is biased toward the orientation presented on the previous trial in a spatially specific manner (St. John-Saaltink, Kok, Lau, & de Lange, 2016). However, the spatial tuning is relatively broad (Collins, 2019; Fischer & Whitney, 2014; Fritsche et al., 2017).

Finally, some researchers examined whether the particular way in which participants produce a response influences SD. If SD acts on perception, then it should be possible to measure it regardless of response modality. Responses can be decorrelated from stimuli either by asking participants to respond on only a subset of trials (Fischer & Whitney, 2014), or by asking them to match the mirror orientation of the Gabor patch (Cicchini, Mikellidou, & Burr, 2018). These studies have shown that SD occurs between stimuli rather than responses, but their interpretation is complicated by the finding that the orientation signal in V1 is biased toward the previous response, even when incorrect, rather than the previous stimulus (St. John-Saaltink, Kok, Lau, & de Lange, 2016).

All of these studies seem to point toward a perceptual level for SD, although each line of evidence has its limitations, as outlined above. In particular, the two alternative forced choice task used by Cicchini et al. (2017) and Fritsche et al. (2017, 2020) may not be the most direct measure of perception, because participants still have to perform an orientation reproduction on every other trial. This process may incite participants to explicitly recall and use orientation information from previous stimuli.

The current experiment uses a visual search task in which the response (to report the presence or absence of a visual oddball among distractors) is orthogonal to the stimulus feature on which SD may occur (orientation), thus bypassing the limitations of previous studies. The hypothesis was that if SD acts directly on appearance, then subjecting only one visual object to SD would make it pop out from visual objects not subject to SD. Likewise, if the object influenced by SD is physically different from surrounding objects, SD may make it appear identical under certain conditions. In other words, performance in the oddball task should depend on the relationship between the oddball and the surrounding distractors, and on the relative orientation between the oddball and the previous trial.

On each trial, observers saw several Gabor patches, one of which could differ from the others. The oddball, when there was one, was always at the same test location. An assumption was that the test patch would be subject to SD, whereas the distractors would not (or to a lesser extent), because distractors changed locations on every trial and attention was oriented toward the test location (because observers were told to monitor that location for a potential oddball, and likely also learned that this was indeed the oddball location

across trials). This is a reasonable assumption given evidence that SD is spatially selective (Collins, 2019; Fischer & Whitney, 2014) and enhanced at the location of attention (Fischer & Whitney, 2014).

The hypothesis of SD between successive test patches makes different predictions for oddball versus no-oddball trials. For oddball trials, if SD alters the appearance of the test patch (but not, or to a lesser extent, that of the distractors) such that the test looks more like the distractors, the proportion of “oddball” responses should decrease. If SD pulls the appearance of the test patch away from the distractor orientation, the proportion of “oddball” responses should increase. These modulations should occur when the relative orientation between previous and current test patches is small. (Previous work has determined that SD is tuned in feature space and emerges most strongly for relative orientations around $\pm 10\text{--}20^\circ$). The hypotheses are summarized in Figure 1.

Methods

Subjects

Twenty subjects participated in this experiment (9 women, aged 21–39 years old). The experiment was run during the first lockdown due to the novel coronavirus disease-2019 pandemic, which in France was from March 17 to May 11, 2020, with research laboratories remaining closed until the end of June 2020. Therefore, it was transferred from a laboratory setting to an online platform (testable.org). Participation was voluntary and anonymous. The link to the online experiment was posted on student forums at the University of Paris. All participants self-reported normal or corrected-to-normal vision and no known neurologic deficits. The relevant ethical information pertaining to the laboratory version of the study and approved by the local ethics committee (Conseil d’Evaluation Ethique pour les Recherches en Santé, Paris Descartes University) and in accordance with the declaration of Helsinki was provided at the onset of the experiment, and participants had to click to accept the consent form before proceeding with the experiment. Sex, age, and screen characteristics were recorded for each participant.

Stimuli

Stimuli were full-contrast Gabor patches generated with Psychtoolbox for Matlab (Brainard, 1997; Kleiner, Brainard, Pelli, Ingling, Murray, & Broussard, 2007; Pelli, 1997). Because screen size depended on the participant’s personal computer, the actual size of the



Figure 1. Procedure. (Left) On each trial, a Gabor patch appeared at the test location (pink circle). (The patches and masks are stylized here for ease of viewing, and the difference between test and distractor patches is exaggerated for visibility). Eleven distractor patches appeared at random locations that changed on each trial. The patches stayed on for 500 ms and were then replaced by a mask until participants responded as to whether an oddball was present or absent. The insets on the right illustrate the predictions and hypothetical results. Between trials 1 and 2, there is a positive relative orientation between test Gabor patches (i.e., trial 1 Gabor is more counterclockwise than trial 2 Gabor). Because distractors are more clockwise than the test, any SD between test patches will pull the test in the counterclockwise direction, away from the distractors. The physical distance between test and distractors is smaller than the perceived distance, thus the proportion of oddball responses should increase. Between trials 2 and 3, the relative orientation between test patches is negative, the physical discrepancy would therefore be larger than the perceived discrepancy (if the appearance is influenced by SD) and the proportion oddball trials should decrease. The hypothesis that performance will depend on relative orientation is illustrated by the orange curve in the cartoon results. The blue curve illustrates the alternative hypothesis of no effect of the previous trial on perceived orientation.

Gabors varied slightly between participants. They were instructed to view the stimuli full screen and sit 60 cm from the screen, which means that the Gabor patches were 1.5 ± 0.4 dva in diameter. The test Gabor was always located to the left of the fixation spot (dashed pink location in Figure 1), and subjects were alerted

to this fact. They were instructed to maintain fixation on the fixation spot (a small dot at screen center) throughout the experiment. On each trial, 12 Gabor patches appeared, one always at the test location, and 11 others (the distractors) at a random location within an approximately 16×16 dva window around the fixation

spot, the only constraint being that Gabor patches could not overlap. In addition to the Gabor patches, Perlin-noise mask stimuli with the same dimensions and locations were generated using Psychtoolbox's CreateProceduralNoise function. All stimuli were generated using Psychtoolbox and saved as image files that were then presented via the online platform.

Procedure

Gabor patches appeared for 500 ms, followed immediately by the masks. Participants were instructed to respond by pressing on one key if they saw an oddball and another key if they did not see an oddball. Oddballs were present on two-thirds of the trials. All distractors on any given trial had the same orientation, and on each trial, distractor orientation was chosen randomly from 1° to 180° . The oddball differed from the distractors by 5° (clockwise rotation for one-half of the participants, counterclockwise for the other one-half). This value was chosen after pretests determined that it should be near threshold for most participants (individual staircases were not feasible via the online platform). Each participant performed 720 trials, resulting from four repetitions of each distractor orientation. The entire study took between 20 and 30 minutes to complete.

Data analyses

Data analysis was performed with Matlab (MathWorks Inc, Natick, MA) using the CircStat toolbox for circular statistics (Berens, 2009), and R (R Core Team, 2019). Reaction time data was recorded and used to eliminate trials in which participants took longer than 2 standard deviations from their own average RT (0.15%–0.45% of the data). The relative orientation between test patches was quantified as the circular distance between the current and previous test patch orientations. Performance was quantified by the probability of responding “oddball present” for different relative orientations. The individual raw data was smoothed by calculating running averages across windows of relative orientations in each condition (oddball, non-oddball). Each averaging window was 22.5° wide and successive windows were stepped by 1° , starting at a -90° relative orientation. The mean number of trials for each averaging window per participant is illustrated in Figures 3 and 4 (bottom). Grand averages were calculated on the smoothed individual data. Significance was assessed with permutation tests, in which the x-labels (relative orientation) were shuffled for 1000 iterations and the mean and 95% confidence intervals calculated on the resulting distribution. This process approximates a

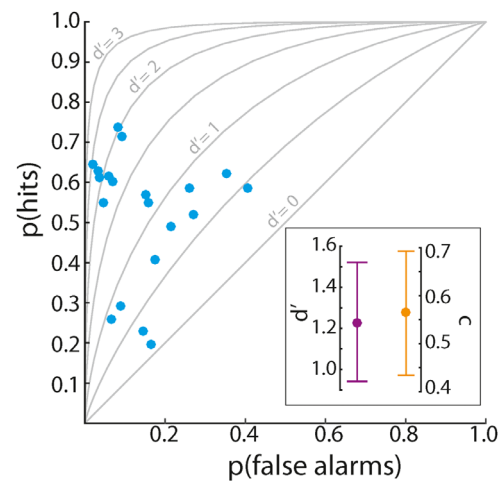


Figure 2. Overall performance. The proportion of hits and false alarms with theoretical receiver operating characteristic curves. Each point is an individual subject. Inset: Median d' (purple) and criterion c (orange). Error bars correspond to 95% bootstrapped confidence intervals.

null distribution of no SD that has the same response bias (height) as the empirical data. First derivatives of Gaussians (DoG) further quantified the data. The DoG is given by $y = h + (x + b)awce - (w(x + b))^2$, where x is the relative orientation between successive test patches, a the amplitude of the curve, w its width, h its height, b the intercept, and c the constant $\sqrt{2/e^{-0.5}}$. Fits were constrained to have positive a and w parameters, meaning that the DoG decreased then increased (and not the other way around). These parameters were constrained in this way because a negative-then-positive inflection of the DoG corresponds to a specific hypothesis about the way in which performance and relative orientation should covary.

Results

Oddball detection performance varied between participants (Figure 2), with d' ranging from 0.12 to 2.4, and criteria (c) from 0.01 (no preference for a particular response) to 1.08 (preference for replying “no oddball”).

The mean reaction time was 1054 ± 239 ms and 1057 ± 256 ms for oddball and no oddball trials, respectively. Given stimulus presentation times, this means that approximately 1.5 seconds elapsed between successive stimuli.

Figure 3 shows the probability of responding “oddball present” as a function of the relative orientation between successive patches at the test location, for trials with and without oddballs. Negative relative orientations are cases in which the previous test

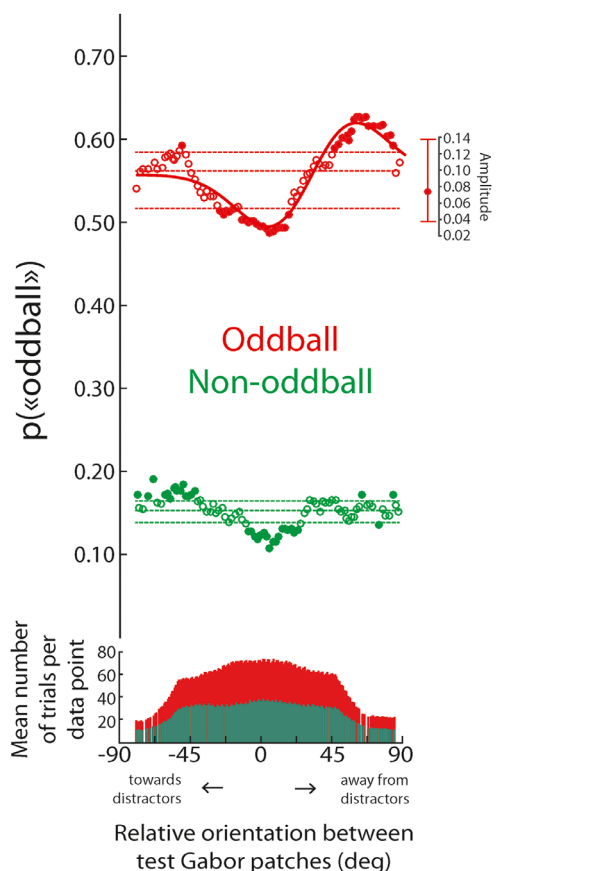


Figure 3. (Top) Grand averages. Proportion of “oddball” responses as a function of the relative orientation between test Gabor patches, for oddball (red) and no-oddball (green) trials. Dashed lines represent bootstrapped median and 95% confidence intervals for each condition, and modulations that differ significantly from the mean are represented as filled dots. The inset shows the median and 95% confidence intervals for the amplitude of individual DoG fits. (Bottom) The mean number of trials per data point.

orientation was in the same direction as the distance between the current test and distractors. In oddball trials, SD should make the test patch appear more like the distractors, and indeed the proportion of oddball responses decreased. This decrease in responses was maximal for relative orientations of approximately -10° to -20° , coherent with feature tuning of previous reports. Positive relative orientations are cases in which the previous test orientation was opposite in direction to the distance between current test and distractors. In line with the prediction, the proportion of oddball responses increased. Interestingly, this increase in oddball responses occurred for relative orientations starting at approximately 40° , and continued for even larger relative orientations. This range of SD is wider than previously reported.

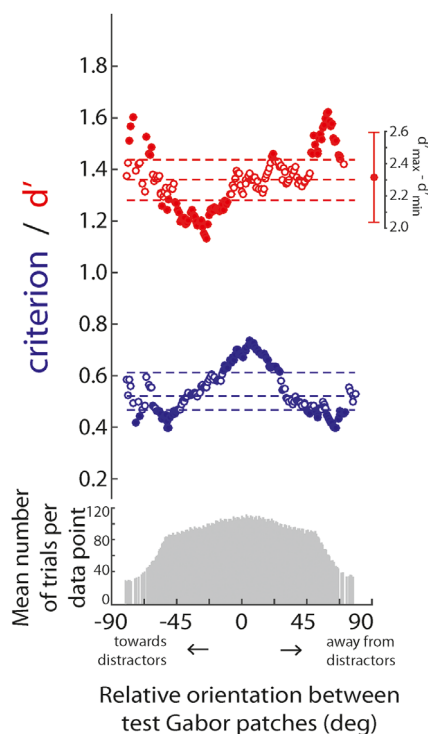


Figure 4. (Top) Grand averages. Sensitivity (d') and criterion (c) as a function of the relative orientation between test Gabor patches. Dashed lines represent bootstrapped median and 95% confidence intervals for each condition, and modulations that differ significantly from the mean are represented as filled dots. The inset shows the median and 95% between-subject confidence intervals for the difference between maximum and minimum d' values. (Bottom) Mean number of trials per data point.

For no-oddball trials, if SD alters the appearance of the test patch (but not, or less than, the distractors), then there should be more erroneous oddball reports in these trials for nearby relative orientations ($\pm 10^\circ$ – 20°), whatever the sign of the relative orientation (positive or negative). “Toward” and “away from” distractors are thus misleading labels here. The data supported this conclusion; there was a lower proportion of oddball responses when the relative orientation was near zero, and a slight increase of oddball responses for larger relative orientations, in particular near -45° . For positive relative orientations, the increase was maintained for such a wide range of relative orientations that the modulation did not reach significance (except for two data points). The difference between negative and positive relative orientations may be due to the fact that toward distractor relative orientations are opposite to those experienced in oddball trials (in which the orientation difference between test and distractor patches was fixed). Because this direction of difference

between test and distractors was unusual, it may have been particularly salient.

A permutation test assessed the significance of performance modulations as a function of relative orientation. The median and 95% confidence intervals are plotted as thin dashed lines in Figure 3, and data points that differ significantly from the null distribution are highlighted as filled circles. They corroborate the descriptive analysis presented above.

To further quantify the modulation of performance as a function of relative orientation in the oddball trials, the data was fit with a first DoG. The fit shown in Figure 3 is for illustrative purposes; DoGs were fit on individual data, and the mean number of trials per individual data point is illustrated in the bottom panel. The inset shows the median amplitude and 95% confidence intervals of individual fits. (Note that the data from 18 of 20 subjects was fit by a DoG but that 2 of 20 datasets could not be fit adequately.) The average amplitude of the DoG (i.e., the difference between the maximum and minimum) was 0.074 ± 0.027 , meaning that variations in relative orientation led to fluctuations of approximately 7 percentage points in the proportion of oddball responses (the range of amplitude across participants was 0.038–0.138).

The analysis on the proportion of oddball responses is corroborated by a signal detection theory analysis. Figure 4 shows how sensitivity (d') and criterion (c) progress as a function of the relative orientation between Gabor patches. When SD pulls the oddball in the direction of the distractor, sensitivity to oddballs decreases, whereas when SD pulls the oddball away from distractors, sensitivity increases. Criterion also changes slightly depending on relative orientation, increasing when relative to orientation is small compared to when the relative orientation is larger ($\pm 45^\circ$). In other words, when the appearance of the target patch is not within the range of relative orientations for which SD is usually observed, participants have a tendency to respond no oddball more often. The significance of these modulations was assessed as above by a permutation test (see dashed lines in Figure 4). Individual differences between minimum and maximum d' values are illustrated in the inset, and the bottom panel illustrates the mean number of trials per individual data point.

Finally, the effect of trials at difference distances in time was also examined. Figure 5 shows DoG fits to data in which relative orientation was defined as the difference in orientation between the current trials and trials farther and farther into the past (negative distances). The reddest DoG is the same as the fit shown in Figure 3; fits for trials 1, 2, 3, 4, 5, 10, 25, and 50 in the past are shown in increasingly light shades of red. The amplitude of the DoG decreases with distance in the past, as shown in the inset at the bottom. As a control analysis, modulations in performance

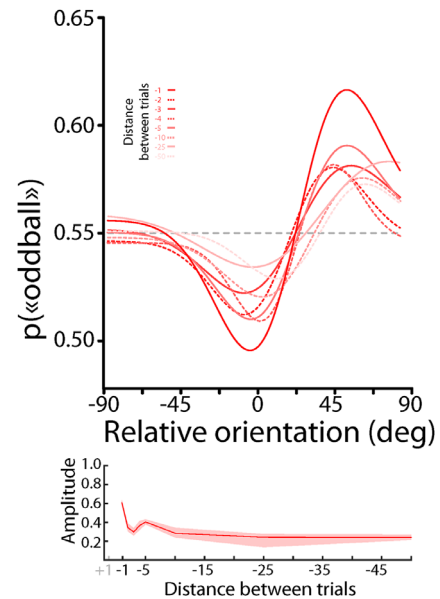


Figure 5. DoG fits of the proportion oddball responses as a function of the relative orientation between the current test Gabor patch and the Gabor patch n trials in past (increasingly pale shades of red). The grey line is for illustrative purposes: the proportion of oddball responses as a function of relative orientation between future and current trials could not be fit with a DoG. The bottom inset shows median (and 95% confidence intervals) as a function number of trials into the past.

as a function of the relative orientation between the current and next trial was examined; of course, there is no expectation that the future trial will influence the current trial. The grey line in Figure 5 is for illustrative purposes only; the data could not be fit by a DoG.

Discussion

Observers performed an oddball task and indicated whether the Gabor patch at the test location differed or was identical to surrounding distractors. Because distractor patches changed locations on each trial, and because attention was oriented toward the test location, it is likely that SD acted more on the test location than the distractor locations.

Results showed that performance depended on the relative orientation between successive test patches. When SD attracted the test patch away from the distractor orientation, observers tended to report seeing an oddball more often than when SD attracted the oddball toward the distractor orientation. The present results fit well with a recent report showing that SD can influence shape perception in a visual search task (Manassi, Kristjánsson, & Whitney, 2019). A signal detection analysis of the data further

revealed that sensitivity decreased when SD pulled the test patch toward the distractor orientation and increased when SD pulled the test patch away from distractors. These results argue in favor of the idea that SD occurs at the level of perception, rather than at response determination stages. Indeed, reporting the presence or absence of an oddball relies on a perceptual comparison between nearby stimuli, and the yes/no response is orthogonal to the visual attribute causing SD (orientation). Determining the level of SD is a theoretically central argument in favor of the continuity field as a basic perceptual mechanism that exploits short-term regularities in the visual environment to smooth perception over time. Such a continuity field may, therefore, be one of the mechanisms that construct our visual stability despite the frequent changes in input. Of course, the current data show only that SD occurs at a perceptual level, which does not eliminate the possibility that SD may occur at later processing stages, including response determination, as well. However, these levels of processing are unlikely to have contributed to the pattern of behavioral results observed here.

The influence of previous stimuli on current percepts extends several trials into the past (e.g., Fischer & Whitney, 2014), and such effects were also found here, with (albeit small) performance modulations for trials up to 10 trials back. The modulations for trials further back in time, although significant, do not seem to have the same feature tuning characteristics, making it unsure that the fits are psychologically meaningful.

A final important aspect of the current data is that despite the fact that the continuity field is wide (Collins, 2019; Fischer & Whitney, 2014; Fritsche et al., 2017), SD seems to differentially affect separate objects within it. On successive trials, test and distractor patches were often in close spatial proximity and fell well within the boundaries of the continuity field. SD may, therefore, be selective to specific objects. Quantifying this would require comparing SD as a function of the spatial proximity between test and distractors, an analysis which would necessitate more data than that available here. Nevertheless, the relationship between SD and object identity may be an interesting avenue for future research.

Keywords: serial dependence, perception, visual search

Acknowledgments

Commercial relationships: none.
 Corresponding author: Thérèse Collins.
 Email: therese.collins@u-paris.fr.
 Address: Integrative Neuroscience and Cognition Center, Université de Paris and CNRS, 45 rue des Saints-Pères 75006 Paris, France.

References

- Berens, P. (2009). CircStat: A *MATLAB* toolbox for circular statistics. *Journal of Statistical Software*, *31*(10), <https://doi.org/10.18637/jss.v031.i10>.
- Brainard, D. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.
- 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*, *111*(21), 7867–7872, <https://doi.org/10.1073/pnas.1402785111>.
- Cicchini, G.M., Mikellidou, K., & Burr, D. (2017). Serial dependencies act directly on perception. *Journal of Vision*, *17*(14), 6, <https://doi.org/10.1167/17.14.6>.
- Cicchini, Guido Marco, Mikellidou, K., & Burr, D. C. (2018). The functional role of serial dependence. *Proceedings of the Royal Society B: Biological Sciences*, *285*(1890), 20181722, <https://doi.org/10.1098/rspb.2018.1722>.
- Collins, T. (2019). The perceptual continuity field is retinotopic. *Scientific Reports*, *9*(1), 18841, <https://doi.org/10.1038/s41598-019-55134-6>.
- Dong, D. W., & Atick, J. J. (1995). Statistics of natural time-varying images. *Network: Computation in Neural Systems*, *6*(3), 345–358.
- Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, *17*(5), 738–743, <https://doi.org/10.1038/nn.3689>.
- Fritsche, M., Mostert, P., & de Lange, F. P. (2017). Opposite effects of recent history on perception and decision. *Current Biology*, *27*(4), 590–595, <https://doi.org/10.1016/j.cub.2017.01.006>.
- Fritsche, M., Spaak, E., & de Lange, F. P. (2020). A Bayesian and efficient observer model explains concurrent attractive and repulsive history biases in visual perception. *eLife*, *9*, e55389, <https://doi.org/10.7554/eLife.55389>.
- Gekas, N., McDermott, K. C., & Mamassian, P. (2019). Disambiguating serial effects of multiple timescales. *Journal of Vision*, *19*(6), 24, <https://doi.org/10.1167/19.6.24>.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, *36*(14), 1.
- Manassi, M., Kristjánsson, Á., & Whitney, D. (2019). Serial dependence in a simulated clinical visual search task. *Scientific Reports*, *9*(1), 19937, <https://doi.org/10.1038/s41598-019-56315-z>.
- Pelli, Denis G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.

- R Core Team. (2019). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing, <https://www.R-project.org/>.
- Schwartz, O., Hsu, A., & Dayan, P. (2007). Space and time in visual context. *Nature Reviews Neuroscience*, 8(7), 522–535, <https://doi.org/10.1038/nrn2155>.
- Simoncelli, E. P., & Olshausen, B. A. (2001). Natural Image Statistics and Neural Representation. *Annual Review of Neuroscience*, 24(1), 1193–1216, <https://doi.org/10.1146/annurev.neuro.24.1.1193>.
- 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(23), 6186–6192, <https://doi.org/10.1523/JNEUROSCI.4390-15.2016>.
- Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, 56(2), 366–383, <https://doi.org/10.1016/j.neuron.2007.10.012>.