State-dependent serial dependence in perceptual decisions

Ayberk Ozkirli, David Pascucci

Laboratory of Psychophysics, Brain Mind Institute, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

Keywords: serial dependence, history biases, perceptual uncertainty, decision-making, optimality, Bayesian perception

Corresponding author: ayberk.ozkirli@gmail.com

Abstract

According to traditional views, perception deals with uncertainty using optimal strategies. For instance, when prior stimuli are more reliable than current ones, perception tends to give greater weight to the past, resulting in stronger serial dependence in perceptual decisions. Here, we challenge this view and reveal a systematic deviation from this strategy. We asked participants to reproduce the average orientation of an ensemble of stimuli, with the level of uncertainty manipulated by varying the ensemble variability. Surprisingly, we found that observers exhibited stronger serial dependence when prior stimuli were uncertain rather than reliable. To explain this pattern, we present an extension of the ideal observer model. In this model, uncertain events trigger internal states of the observer, such as expectations about moments of degraded sensory input, thereby modulating the tendency to integrate prior and current stimuli. The striking finding is that manipulation of these internal states through external feedback can lead to clear effects on serial dependence, even when the level of uncertainty in the stimulus remains constant. We propose that serial dependence reflects a 'state-dependent' mechanism that the brain amplifies to cope with moments of expected uncertainty.

Introduction

No one likes uncertainty. Yet, uncertainty is not only a persistent aspect of our environment, but also an extremely volatile one, evolving swiftly and sometimes with dramatic consequences. Imagine a pilot navigating through rapidly changing weather conditions, facing the risk of making the wrong decisions by relying on suddenly unclear visual cues.

According to traditional models of perception and decision-making, the brain deals with uncertainty in a near-optimal and Bayesian way[1–9]. For example, when a stimulus is unclear or ambiguous, the brain leverages more reliable sources of information, such as prior experiences, to guide its interpretation[9–12].

A central tenet of these models is that uncertainty is computed from each stimulus, and then integrated over multiple stimuli, with a greater weight to the more reliable ones[2,4,10,13,14]. Models of this kind, often described within the framework of the 'ideal observer'[15], can explain performance in a variety of perceptual tasks, suggesting that, in most cases, humans can effectively use all available information to reduce uncertainty in an optimal way. However, not all human behavior conforms to such Bayesian and optimality principles[16–19].

Recent work on serial dependence, a phenomenon where perceptual decisions are biased towards previous stimuli[20,21], has revealed clear deviations from ideal observer strategies. In particular, when integrating current and prior stimuli, humans do not take the relative uncertainty between stimuli into account in the same way as ideal observers[22–24]. This suggests that classic models, which typically assume fixed computations based solely on stimulus uncertainty, may not fully account for other factors. For instance, moment-to-moment fluctuations in uncertainty may trigger internal states and expectations, which eventually interact with the way uncertainty is handled in perceptual decision-making[25,26].

Here, we report a novel phenomenon and accompanying framework to explain how humans integrate sequential events as a function of uncertainty. We used a perceptual averaging task in which participants reproduced the perceived average orientation of an ensemble of stimuli (Figure 1). By manipulating uncertainty in the ensemble orientations, we found that serial dependence —i.e., the bias towards the average orientation of the preceding ensemble— increased when the previous stimulus was uncertain and was the largest when both the current and previous stimuli were uncertain: a pattern nearly opposite from that predicted by the ideal observer. To explain this behavior, we revisited the classical ideal observer[10,27] and introduced an extended version. In this extended model, serial dependence is influenced not only by the stimulus, but also by internal 'states' that dynamically change, based on the level of uncertainty experienced by the observer. We hypothesized that when

high uncertainty persists across consecutive moments, the brain enters a state that promotes serial dependence.

Based on this hypothesis, we predicted that, even under constant stimulus uncertainty, changes in internal states can still affect serial dependence. To test this prediction, we performed a final experiment where we introduced and manipulated performance feedback. Strikingly, serial dependence was observed only when participants believed their performance was poor due to manipulated feedback, despite both the level of stimulus uncertainty and performance remaining constant throughout the experiment.

These findings not only reconcile traditional models of perception with the emerging insights from the domain of serial dependence but also lay the foundation for a new perspective. We propose that phenomena like serial dependence are largely 'state-dependent' and contingent upon internal, dynamic states that the brain visits and transitions between to mitigate moments of uncertainty.

Results

Stimulus uncertainty

In serial dependence, perceptual decisions are biased towards the stimulus shown on the preceding trial. For instance, the reported average orientation of the current ensemble is biased towards the average of the ensemble presented before. This bias is more pronounced when the difference between the previous and current stimulus is relatively small[21,28] (Δ; see Methods). A key aspect of the present paradigm is the effect of uncertainty, that is, whether the bias would increase or decrease as a function of the relative uncertainty between the current and previous ensemble.

The predictions of an ideal observer are shown in Figure 1C. According to this model, the strength of serial dependence should increase when the current stimulus is more uncertain than the previous one (as observed in the condition where uncertainty transitions from *low-to-high*) and decrease when the current stimulus is more reliable (*high-to-low*). This is because, in classic ideal observer models, optimal integration is contingent solely upon the relative uncertainty between two consecutive stimuli, and the more reliable stimulus consistently carries more weight[10].

Sixteen observers reproduced the perceived average orientation of an ensemble of 36 Gabor stimuli (Figure 1A). On randomly interleaved trials, the ensemble orientations were either drawn from a normal distribution with a standard deviation (σ) of 8° (*high* σ) or all identical (*low* σ = 0°). These two conditions carried different degrees of uncertainty, as reflected in performance, with a larger standard deviation of adjustment errors in the *high* σ (12.56 ± 2.69°) compared to the *low* σ condition (9.60 ± 1.92°; *high* vs. *low*: $t(15) = 7.93$, $p < .001$, Cohen's $d' = 1.98$, paired t-test).

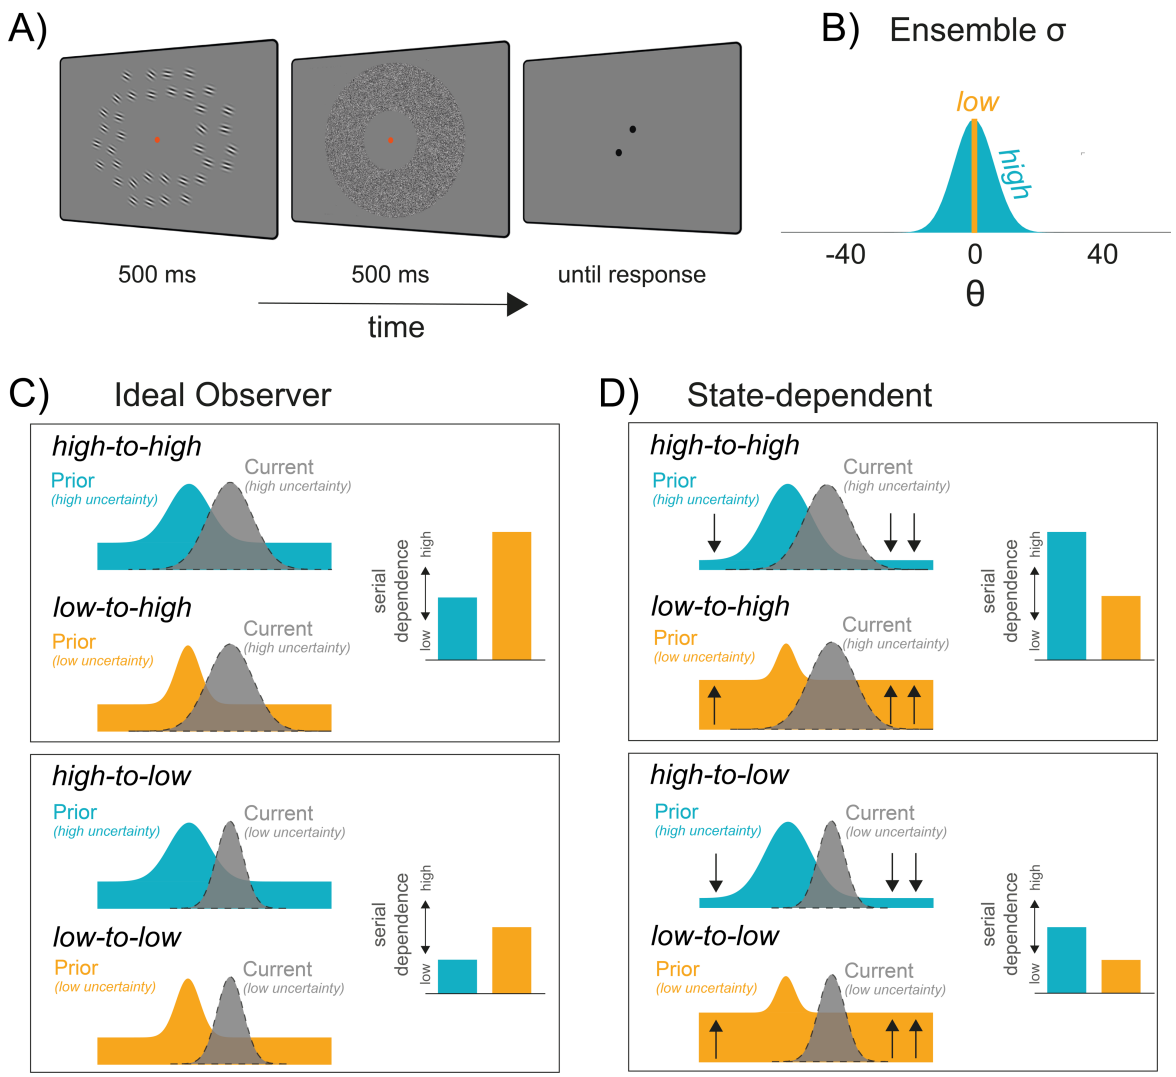

Figure 1. Main paradigm and models of serial dependence. A) Participants reproduced the average orientation of an ensemble of 36 Gabors shown in the periphery of the visual field (see Methods). B) We manipulated uncertainty by intermixing ensembles with *high* and *low* orientation variability (e.g., standard deviation [σ] of the orientation distribution, σ = 8° vs. 0° in Experiment 1a, σ = 15° vs. 5° in Experiment 1b). Color coding: blue for *high* and orange for *low* uncertainty. C-D) Two models of the effects of uncertainty in serial dependence. In the ideal observer model (C), prior and current stimuli are integrated depending on their relative uncertainty, with the more reliable stimulus always receiving larger weights. In such a model, serial dependence is always stronger when the previous stimulus is more reliable than the current one. The prior is here composed of a mixture of a uniform distribution and a normal distribution centered on the previous stimulus (blue and orange distributions, see Methods for further details). Note that the strength of the prior here is only determined by the uncertainty in the stimulus, whereas the uniform component remains fixed. In the 'state-dependent' model (D), stimuli are integrated depending on their relative uncertainty as in the ideal observer, but the overall tendency towards integration (here modeled as the ratio between the uniform and normal prior, see Methods) varies as a function of internal states: when an uncertain stimulus is encountered, the weight of the uniform component decreases and the tendency towards serial dependence increases. These changes are highlighted by the black arrows. The four boxes in C-D) depict all four combinations of prior and current uncertainty tested (see main text).

To test the predictions of the ideal observer, we estimated the strength of serial dependence using the half-amplitude parameter from a δoG fit (α, see Methods), considering all possible

combinations of current and past uncertainty, in a full-factorial design. A repeated-measures ANOVA with factors 'Uncertainty on the preceding trial' (*low* vs. *high*) and 'Uncertainty on the current trial' (*low* vs. *high*) revealed no effect of the current trial uncertainty (F(1, 15) = 1.54, p = .23, η_p^2 = 0.09), but a significant main effect of the uncertainty on the preceding trial (F(1, 15) = 6.86, p = .019, η_p^2 = 0.31) as well as a significant interaction (F(1, 15) = 14.32, p = .002, η_p^2 = 0.49). The interaction was driven by the serial dependence occurring when both the previous and current stimuli were of high uncertainty (the *high-to-high* condition), which was significantly larger when compared to all the other conditions (*high-to-high* vs. average of the other conditions: t(15) = 4.566, p < .001, Cohen's d' = 1.141; see Figure 2A).

This finding was inconsistent and even nearly opposite to what would be expected by the ideal observer with classic optimal integration (see Figure 1C). In other words, observers did not exhibit stronger biases when the prior stimulus was more reliable; instead, they exhibited the strongest bias toward past stimuli when uncertainty was high in both the current and previous trials.

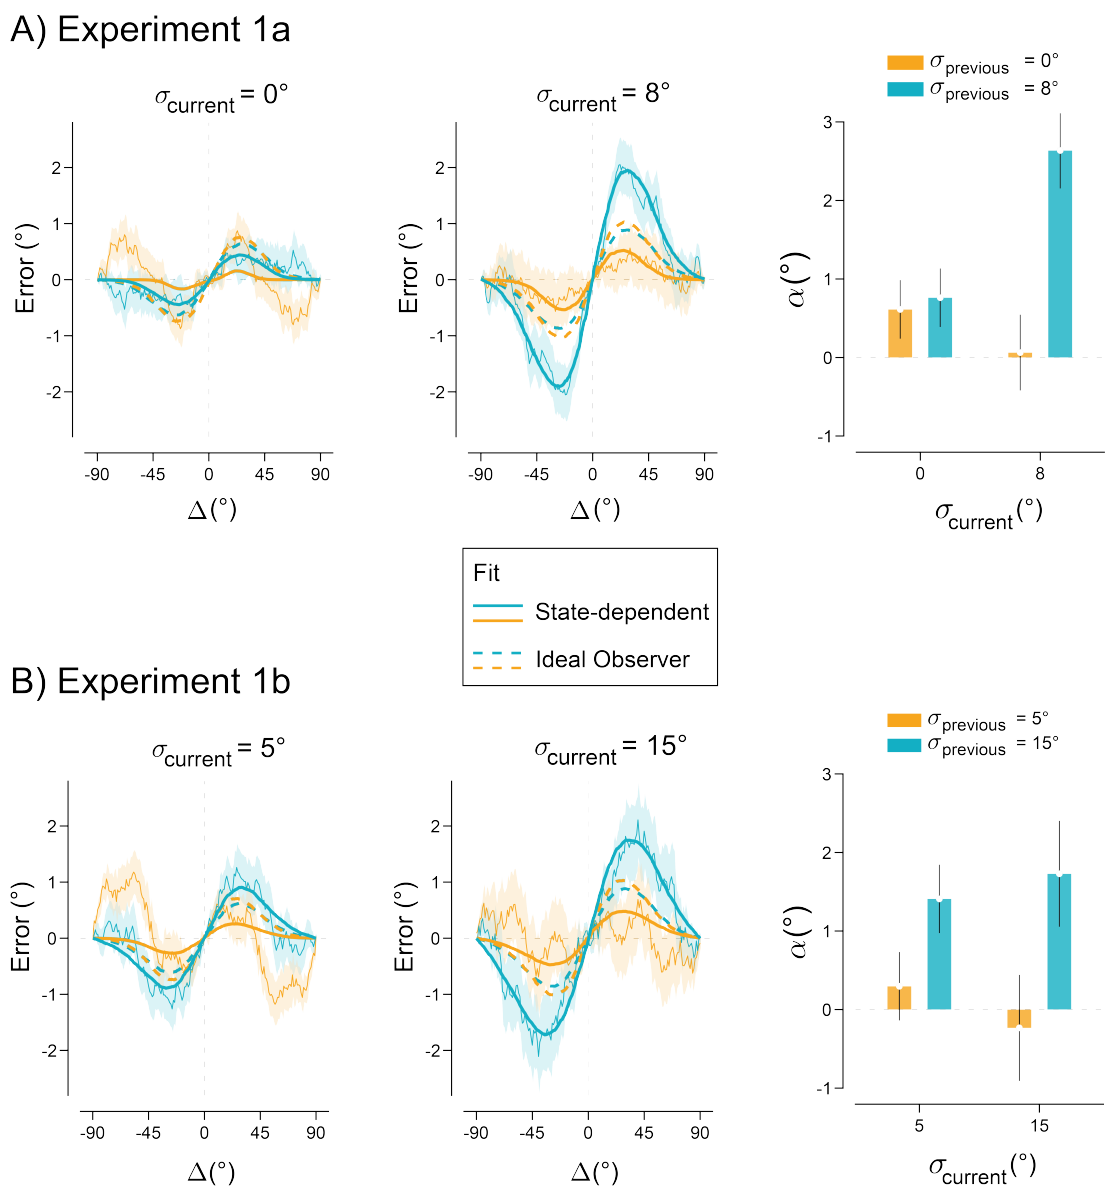

Figure 2. Results of Experiment 1a-b. A) In Experiment 1a, *high* and *low* uncertainty trials contained ensembles where the standard deviation of the orientation distribution (σ) was 8° and 0°, respectively. The first panel shows serial dependence in the *high-to-low* (in blue) and *low-to-low* conditions (in orange). The second panel shows serial dependence in the *high-to-high* (in blue) and *low-to-high* conditions (in orange). A circular running average (sliding window size of 35°) is computed on the adjustment errors of all participants as a function of the difference in average orientation between the previous and current ensemble (Δ; previous minus current). Shaded areas are 1 standard deviation of the running average. The curves represent the fit of two models: dashed lines are the predictions of the ideal observer; continuous lines are the predictions of the state-dependent model. The bar plots in the third panel show the average half-amplitude parameter of δoG fits across participants, for each combination of prior and current uncertainty. Error bars represent standard errors of the mean, corrected for within-subject repeated-measures designs[29]. B) Results from Experiment 1b where the two σ used were 15° and 5°. The color coding is the same as in A.

One aspect worth noting is that we did not find any modulation of serial dependence due to different levels of uncertainty on the previous trial, when the uncertainty on the current trial was low (Figure 2A, first panel). This could potentially be due to ceiling performance and weak serial dependence, as the *low* σ condition had minimal uncertainty due to all stimuli having the same orientation.

We therefore performed an additional experiment with 16 new participants and increased the standard deviations (σ) in both the *high* (σ = 15°) and *low* (σ = 5°) uncertainty conditions. As in the first experiment, observers were less precise in the *high* σ (standard deviation of adjustment errors: 16.54 ± 4.90°) compared to the *low* σ condition (11.05 ± 2.78°; *high* vs. *low*: t(15) = 7.37, p < .001, *d'* = 1.84, paired t-test).

Upon increasing σ in both conditions, a repeated-measures ANOVA revealed a significant effect of the preceding trial (F(1, 15) = 8.318, $p = 0.011$, $\eta_p^2 = 0.357$), but no effect of the current trial uncertainty (F(1, 15) = 0.017, $p = 0.898$, $\eta_p^2 = 0.001$) and no significant interaction (F(1, 15) = 0.499, $p = .49$, $\eta_p^2 = 0.032$). Therefore, when uncertainty was increased in both conditions, serial dependence was overall stronger when the preceding stimulus was more uncertain, regardless of the current trial condition, indicating that past uncertainty plays a key role.

State-dependent serial dependence

As mentioned, within the framework of the ideal observer, the bias towards prior stimuli depends exclusively on the relative uncertainty between the current and previous stimuli[10]. However, this model falls short in explaining the results of Experiments 1a and 1b (Figure 2), as well as similar patterns that are qualitatively evident in other studies[22–24,30].

To account for these deviations from the ideal observer, we formulated an extended version that incorporates a dynamic component (see Methods). We reasoned that the history of uncertainty in a series of stimuli can lead to changes in internal states of the observer, such as expectations about consecutive moments of high or low uncertainty, which eventually interact with the strength of serial dependence, beyond the effects of the stimulus itself (see Figure 1D). For instance, when an uncertain stimulus is encountered, there may be an increased tendency to integrate prior stimuli into current perceptual decisions due to the expectation of a transition into a more uncertain state of the environment. This, in turn, leads to a transitory state that promotes serial dependence, thereby increasing the bias towards the past when a new stimulus is encountered.

As a proof of concept, we compared this model, which we refer to as the 'state-dependent' model of serial dependence, against the classic ideal observer (see Methods). The patterns of serial

dependence predicted by the two models are shown in Figure 2 (solid lines: state-dependent model, dashed lines: ideal observer). As evident, the state-dependent model provides a more accurate representation of the actual patterns of serial dependence, outperforming the classic ideal observer model (difference in Akaike Information Criterion between the state-dependent and ideal observer models [ΔAIC_{Exp1a}] = 8.28 and [ΔAIC_{Exp1b}] = 4.26, where $\Delta AIC > 2$ favors the state-dependent model).

In addition to its superior fit, especially in the *high-to-high* and *low-to-high* conditions (see Figure 2), the proposed model yields a clear prediction: even when the stimulus uncertainty remains constant, internal states are determinants of the strength of serial dependence. To test this prediction, we conducted a final experiment.

Internal states modulate serial dependence

We tested 24 new participants with a paradigm in which we fixed the ensemble σ to 15° and introduced trial-by-trial feedback. Our goal was to manipulate the observers' internal states and their tendency to integrate prior and current stimuli by influencing their beliefs about their task performance in different blocks of trials (see Methods and Figure 3A). This manipulation resulted in two types of blocks of interest: the *low feedback* blocks, in which observers received a higher frequency of feedback indicating poor performance, and the *high feedback* blocks, where they received a higher frequency of feedback indicating good performance (see Methods). In both types of blocks, poor or good performance feedback was intermixed with neutral feedback (see Figure 3A). We hypothesized that these distinct conditions might lead participants to perceive one type of block as more uncertain or difficult than the other, even though the actual uncertainty in the stimulus remained constant.

According to our model, experiencing more uncertainty in an entire sequence of trials would lead to an internal state that promotes integration, thus, we expected overall larger serial dependence in the *low* compared to the *high feedback* blocks. Our analysis of serial dependence focused on trials following average performance feedback (e.g., neutral, see Methods), aiming to isolate the effect of internal states by controlling for external factors, such as preceding feedback and stimulus uncertainty while maintaining an overall difference in the level of believed performance between the *low* and *high feedback* blocks. Importantly, in this experiment, we found comparable performance between blocks (error standard deviation in *low*: 13.73 ± 3.59°, *high*: 14.20 ± 3.44°, *low* minus *high*: $t(23) = -1.65$, $p > .05$, $d' = -0.34$).

In line with our model predictions, serial dependence was significantly larger in the *low feedback* compared to the *high feedback* blocks (*low* minus *high*: t(23) = 2.26, p = .03, d' = 0.46). Thus, even though the stimulus uncertainty was fixed, and performance was comparable between conditions, manipulating internal states of the observer led to clear differences in serial dependence.

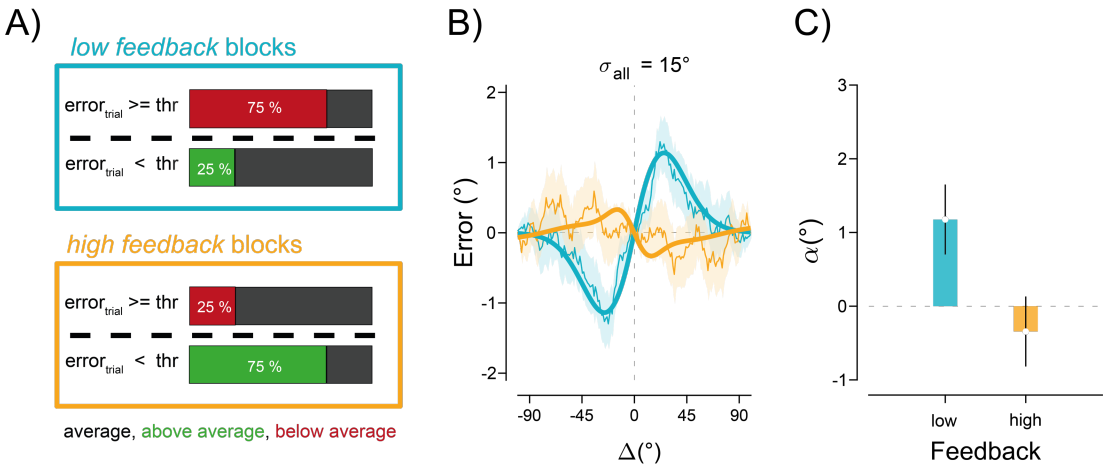

Figure 3. Results of Experiment 2. A) A schematic of the feedback manipulation. In *low feedback* blocks, participants received poor performance feedback more often (red) than good performance feedback (green). In *high feedback* blocks, the ratio was reversed. In both types of blocks, participants received about the same amount of neutral feedback (dark grey). Note that feedback was rigged and based on a pre-established threshold (thr, see Methods). B) Serial dependence in the two types of blocks (blue: *low feedback*, orange: *high feedback*). Curves are the best fitting δoG. C) Average of the δoG half amplitude across participants in the two conditions. Error bars represent standard errors.

Discussion

We investigated how fluctuations in uncertainty influence serial dependence in human perceptual decisions. Classic models of perceptual decision-making predict stronger serial dependence when prior stimuli are more reliable than current ones[10]. Contrary to these predictions, we found that serial dependence increases when prior stimuli are more uncertain. To account for this pattern, we introduced the 'state-dependent' model of serial dependence. In this model, we propose that perceptual decisions are influenced by adaptable and dynamic internal states of the observer. These internal states reflect the observer's expectations and beliefs about current and future states of uncertainty. As uncertainty increases, the brain transitions into a state where prior stimuli have a stronger influence on current perceptual decisions, resulting in increased serial dependence.

Previous research has explained serial dependence leveraging the central assumptions of optimal and Bayesian models of perception[5,31,32]. In these models, serial dependence is a function of internal priors about the stability of objects in the natural world and of the relative reliability between priors and current sensory evidence[27,32,33]. However, these studies have not systematically tested the

effects of fluctuations in uncertainty across all possible scenarios. Thus, evidence supporting these models is based only on a limited range of conditions. By manipulating and testing the effects of uncertainty in a full factorial design, we found a drastic deviation of serial dependence patterns from the predictions of such models: the bias increases as the uncertainty of prior stimuli increases (Experiments 1a-b). These findings suggest that ideal observer models offer only limited and incomplete explanations of serial dependence. In line with this, prior work has used the term 'partial Bayesian inference' when referring to the computations behind serial dependence[24], to explain cases in which, for instance, the effects depended only on the uncertainty in the present, but not in the past[22–24]. It must be noted that even in such cases, existing data contains trends compatible with our results, with peaks of serial dependence when both the past and current stimuli are uncertain, despite the use of diverse stimuli and paradigms[22,23,30,34]. Thus, the pattern observed in the present study appears to be a consistent and general aspect of serial dependence.

In Experiment 1a-b, we found that the condition leading to increased serial dependence was the occurrence of an uncertain stimulus on the preceding trial. Consequently, when the stimulus on the current trial was also uncertain, the bias was most pronounced (see Figure 2A-B). This finding can be understood in the context of temporal expectations and adaptable internal beliefs[25,34]. Even in the absence of regularities and temporal structures, the brain constantly uses recent experiences to infer the current state of the environment[35–38]. These internal beliefs act as downstream, top-down predictions that are fed back and shape processing as early as in primary visual cortex[39]. Crucially, internal beliefs encompass not only specific stimulus features, but also the level of uncertainty associated with encountering a stimulus[25].

While typical models assume a static structure of prior beliefs during sequential decisions (e.g., a mixture of a uniform and a narrow distribution centered on the expected stimulus value)[27,32,33], real beliefs are dynamic and depend on internal states that update and reset upon detecting changes in the environment[34,40]. In contrast to the static structure of ideal observer models (Figure 1C), for instance, our state-dependent model fully embraces this principle. When an uncertain stimulus is encountered, the brain may interpret this event as signaling the beginning of a period of degraded sensory input. This affects the belief about the future environmental state, leading to the recruitment of serial dependence mechanisms and hence stronger temporal integration. From a computational perspective, this can be likened to an adaptive Kalman filter that self-tunes to the environment, learning when to rely more on incoming data or on its predictions, depending on the uncertainty in its measurements[31,41,42].

A key prediction of our model is that internal states modulate serial dependence even when stimulus uncertainty is constant. The results of our Experiment 2 fully confirmed this prediction,

showing that when stimulus uncertainty was fixed, manipulating internal states through biased performance feedback became the determining factor in both the strength and presence of serial dependence. This novel finding clearly demonstrates that internal states of the observer are a key factor in serial dependence, beyond the uncertainty in the stimulus itself. Evidence of modulatory effects of confidence[43], attention[20], memory load[44], time-on-task[45], task demand[36,46–49], and stimulus predictability[50], may all reflect the significant contribution of internal states in determining the tendency to integrate or disregard prior stimuli in current decisions.

Systematic deviations from ideal observer/optimal models, like the one reported here, do not imply that such models are incorrect but highlight the need to move beyond simple and classic models that rely solely on basic aspects of the stimulus[25,26]. In the present context, differences in serial dependence can be found even without manipulating the stimulus parameters: what matters is the internal state of the observer. Similarly, in other domains like spatial vision, classic models that rely solely on local stimulus properties and feedforward processing may fail to explain phenomena like crowding and ensemble perception[51–54].

In summary, here we demonstrate a novel aspect of serial dependence which highlights the fundamental role of flexible, adaptable internal states of the observer. We show that a state-dependent model taking into account these flexible states outperforms classic ‘static’ ideal observer models. We propose that flexible and 'state-dependent' mechanisms of serial dependence may be more efficient than hard-wired and fixed ones. Integrating prior and current stimuli can be advantageous, particularly when anticipating moments of uncertainty in perception. However, integration may not be compulsory, especially when we can rely on our ability to discriminate between changes.

Methods

Ethics statement

The study was approved by the local ethics committee under the Declaration of Helsinki[55], except for preregistration.

Apparatus

Stimuli were presented on a gamma-corrected VG248QE monitor (resolution: 1920 x 1080 pixels, refresh rate: 120 Hz) and were generated with custom-made scripts written in MATLAB (R2013a) and the Psychophysics Toolbox[56], on a Windows-based machine. Participants sat at approximately 57 cm from the computer screen, with the head held stable on a chin rest. The experiments were performed in a dark room.

Participants

62 healthy participants (20 in Experiment 1a, 17 in Experiment 1b; age range of 18-37 years; 5 females in Experiment 1a, 9 females in Experiment 1b, and 15 females in Experiment 2), from the EPFL and the University of Lausanne, took part in the study for a monetary reward (25 CHF/hour). Participants were naïve to the purpose of the experiments. Written informed consent was collected from all participants beforehand.

Stimuli and procedures

The stimuli and procedure are depicted in Figure 1. Each trial started with a dark red fixation spot appearing for 300-1000 ms (randomly sampled in steps of 100 ms). Participants were then presented with a display of 36 oriented Gabor stimuli for 500 ms (peak contrast: 25% Michelson, spatial frequency: 2 cycles/°, Gaussian envelope: 0.225°, phase = 0.25). Gabor stimuli were presented on a gray background (62.66 cd/m^2) and positioned inside an annulus centered at the fovea (internal radius of 4.825°; external radius of 9.175°). No stimuli were presented in the central part of the display to avoid systematic biases toward the center during mean orientation estimates[54,57]. Inside the annulus, Gabor stimuli were arranged at equidistant locations (approximately 2.2° with a 1° random jittering) inside a 12x12° grid. To prevent after-effects, a noise mask was presented within the same annulus for 500 ms immediately after the Gabor ensemble. Following a blank interval of 1000 ms, participants reproduced the perceived average orientation by adjusting a response tool. The response tool consisted of two symmetric dots, each with a diameter of 0.25°, positioned approximately 1.5° apart, marking the endpoints of an imaginary line. The orientation of the response tool was initialized randomly, and participants were instructed to rotate the imaginary line using the mouse and confirm their response with a left click.

On each trial, the orientation of the Gabor stimuli was drawn from a normal distribution with the mean randomly selected over the full 0-179° orientation range, in steps of 1°. Experiment 1a involved two randomly intermixed conditions in which the standard deviation (σ) of the ensemble's orientations was either 0° (*low* σ) or 8° (*high* σ). In Experiment 1b, σ was increased to 5° (*low* σ) and 15° (*high* σ). In Experiment 2, it was fixed at 15°.

In Experiment 2, each trial was followed by written feedback, shown at the center after 1 s from the response: "below average", "average", or "above average". To establish the criteria for assigning feedback, we first determined an error threshold based on the mean absolute error (11°) estimated from the group of participants in Experiment 1b, in the condition with σ = 15°. The frequency was manipulated across different blocks. In the *low feedback* blocks, we assigned "below average" feedback on 75% of the trials where the error exceeded the threshold, and "average"

feedback on the remaining 25% of those trials. Conversely, for trials in which the error was below the threshold, we assigned "above average" feedback on only 25% and "average" feedback to 75% of them. The opposite assignment pattern was applied in the *high feedback* blocks. In other words, "below average" feedback was given on only 25% of trials where the error was above the threshold, while "average" feedback was assigned to the remaining 75% of those trials. For trials with errors below the threshold, "above average" feedback was assigned to 75% and "average" feedback to 25% of them. As a result of this manipulation, participants were more likely to receive "below average" (indicating poorer performance) in the *low feedback* blocks and "above average" (indicating better performance) in the *high feedback* blocks, despite their actual performance potentially remaining consistent (see Figure 3).

An important aspect of this manipulation was that the number of trials with “average” feedback was kept the same in both blocks. This allowed us to examine serial dependence based on the overall block condition (*low* vs. *high feedback*) while controlling for the specific feedback received on each trial. This deliberate design ensured that any differences in serial dependence between conditions could only be attributed to the false belief induced by the frequency of feedback in the respective blocks, rather than the individual trials themselves. According to our debriefing with participants at the end of Experiment 2, 23 out of 25 of them reported having had poorer performance in some blocks, and 22 of those reported that this was due to difficulty difference across blocks (even though the stimulus σ was constant).

All participants received verbal and written instructions at the beginning of each experiment and completed a series of practice trials under the experimenter's guidance. The data from the practice trials were not included in the analysis. In total, participants completed 400 trials in all three experiments. The trials were divided into 10 blocks for Experiment 1a and 1b, and 4 for Experiment 2. The *high* and *low feedback* conditions alternated across the blocks, with the order of the conditions counterbalanced among participants. Participants were instructed to keep their gaze fixed at the center of the screen throughout each block. All experiments lasted about 1 hour.

Analysis

Adjustment errors were computed as the acute angle between the response and the stimulus on each trial. We first excluded trials where errors were larger than 1.5 interquartile ranges above the upper quartile or below the lower quartile (separately for each experimental condition) and reaction times larger than 10 seconds. The first trial of each block, as well as trials following a removed trial, were also excluded. Overall, less than 10% of trials were removed in all experiments. In both Experiment 1a and 2, one participant was excluded due to a large overall negative bias (< -5°). In addition,

participants were excluded if their biases in a specific condition of interest were 1.5 interquartile ranges larger than the upper quartile or smaller than the lower quartile within that condition (3 in Experiment 1a and 1 in Experiment 1b). The final analyses included 16 participants in Experiment 1a, an independent group of 16 participants in Experiment 1b, and 24 participants in Experiment 2. While we present the results after excluding outlier participants, the findings remain consistent when considering all participants.

In Experiment 1a and 1b, we separated the conditions based on the preceding and current stimulus uncertainty, which resulted in four categories: *low-to-low*, *high-to-low*, *low-to-high*, and *high-to-high*. We fitted the first derivative of a Gaussian function (δoG) to the errors as a function of the difference in orientation between the previous and current average orientation of the ensemble (Δ). The δoG has the following form:

$$error = \Delta \alpha w c e^{-(w\Delta)^2} \quad [1]$$

where c is a constant $c = \frac{\sqrt{2}}{e^{-0.5}}$ and w is the inverse of the curve width. The half-amplitude parameter α quantifies the deviation of the errors, in degrees, from the actual orientation, as a function of Δ: positive values of α indicate a systematic deviation of errors towards the orientation of the preceding stimulus, and negative values indicate a deviation away —i.e., repulsion. The δoG was fitted to each condition and for each participant's data separately. Then, repeated-measures ANOVA and t-tests were performed on the estimated α, depending on the experimental design and comparison of interest.

Model

To account for the patterns of serial dependence observed in Experiment 1a and 1b, we compared two models: the classic ideal observer model and an extended version, which we called the 'state-dependent model' (see Figure 1).

In the ideal observer model, serial dependence results from how prior and current stimuli are relatively weighted.[10,21,27,32,33] Weights depend on the relative uncertainty of stimuli and on prior expectations. For example, priors may be composed of a mixture of a Gaussian-like probability distribution centered on the previous stimulus orientation, and a uniform distribution covering the whole orientation space.[32] The ratio between these two distributions determines the overall tendency to weight prior stimuli —i.e., when the uniform is the dominant component, the prior on past stimuli is flat, in a Bayesian sense, and there is no effect of prior stimuli.

We implemented the ideal observer model using a population coding approach.[49] In this approach, the representation of the average orientation is approximated via a probabilistic distribution over the orientation space. The distribution peak reflects the mean orientation of the ensemble, and

the width is related to the uncertainty, parametrized by the mean and concentration parameter of a Von Mises distribution:

$$f_i(\theta) = e^{(k(cos\,(\theta-\varphi_i)\,-1))}, \qquad [2]$$

where $f_i(\theta)$ is the response associated with the i-th orientation θ in the 0-179° orientation space, k is the concentration parameter related to the width and φ_i is the preferred orientation of each coding unit of the population. By normalizing $f_i(\theta)$ by its sum, the population response is represented as a probability distribution. In a Bayesian framework, eq. [2] and the resulting distribution describes the stimulus likelihood (e.g., $p(m|\theta)$, with m = measurement).[27]

In a straightforward implementation of the ideal observer, the posterior $p(\theta|m)$ on each trail t is the result of the combination between the likelihood and the prior —i.e., a representation of previous events (posterior) propagating from the trial $t-1$:

$$p(\theta^t|m^t) \propto p(m^t|\theta^t) * p(\theta^{t-1}|m^{t-1}). \qquad [3]$$

In serial dependence, changes in the uncertainty associated with each stimulus can be incorporated by varying k in eq. [2] as a function of the ensemble standard deviation. For ease of interpretation, we define σ_{low} and σ_{high} as the two standard deviations of a wrapped normal distribution corresponding to two Von Mises concentration parameters k, and reflecting the ensemble's σ under the two levels of uncertainty. In line with previous work[32], p_{past} is modeled as a weighted combination of the posterior distribution on trial $t-1$ and a uniform distribution:

$$p_{past}(\theta^{t-1}|m^{t-1}) = w * \left(p(\theta^{t-1}|m^{t-1})\right) + (1-w) * U(0,179), \qquad [4]$$

where $U(0,179)$ is a uniform prior with equal probability at each orientation and w is a parameter determining the relative weight of the two components. It is straightforward to note that w regulates the strength of the effect of prior stimuli: when $w = 0$ the prior is uniform, and history has no effect.

When $w > 0$, and σ_{low} and σ_{high} are incorporated, eq. [3-4] shifts the mean of the posterior towards the prior, by an amount that depends on the relative σ of the prior and the likelihood, as in classic optimal integration schemes:[10,32] the smaller the σ on the preceding trial, the larger the shift. Using the circular mean of the posterior as the final estimate of the orientation on each trial, one can

observe systematic biases, compatible with the predictions of an ideal observer model of serial dependence (Figure 1).

In the model described above, two factors modulate serial dependence: 1) the ratio between the prior centered on past stimuli and the uniform one; 2) the relative uncertainty (width) of past and current stimulus representations. Standard models assume that the former, the ratio, is fixed and therefore serial dependence is only determined by the relative uncertainty between stimuli, which is computed independently for each stimulus. By relaxing this assumption, we let the parameter w vary as a function of the σ on the preceding trial —i.e., the ratio of the prior about previous stimuli over the uniform prior changed depending on the stimulus history and the associated uncertainty. In the state-dependent model, w was considered a free parameter, functioning as an internal-state-dependent weight that was updated on each trial based on the value of σ.

We compared the ideal observer and the dynamic observer models by feeding them with the single-trial data of Experiment 1a and 1b. Note that the ideal observer has three free parameters (σ_{low}, σ_{high}, w) whereas the dynamic observer model has four (σ_{low}, σ_{high}, w_{low}, w_{high}). To account for this difference, we compared the two models using the Akaike Information Criterion (AIC)[58] which penalizes the number of parameters and is estimated using the following formula:

$$AIC = n_{trials} * log\left(\frac{\sum\left(errors_{predicted} - errors_{observed}\right)^2}{n_{trials}}\right) + 2 * n_{parameters} \quad [5]$$

A difference in AIC between two models (ΔAIC) larger than 2 is typically considered strong evidence in favor of the model with the lower AIC[59].

Acknowledgments

This research was supported by funding from the Swiss National Science Foundation (grant no. PZ00P1_179988 and PZ00P1_179988/2 to DP). The funders had no role in study design, data collection, and analysis. The authors would like to thank Bianca Ribeiro Riccioppo and Zehra Merchant for helping with the collection of the data.

Data availability

The data will be made available in online repositories upon publication.

Author contributions

The authors contributed equally to this work.

References

1. Diamond, M. E. Perceptual uncertainty. *PLOS Biology* **17**, e3000430 (2019).
2. Ernst, M. O. & Banks, M. S. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* **415**, 429–433 (2002).
3. Gardner, J. L. Optimality and heuristics in perceptual neuroscience. *Nat Neurosci* **22**, 514–523 (2019).
4. Kersten, D., Mamassian, P. & Yuille, A. Object perception as Bayesian inference. *Annu. Rev. Psychol.* **55**, 271–304 (2004).
5. Knill, D. C. & Pouget, A. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences* **27**, 712–719 (2004).
6. Körding, K. P. & Wolpert, D. M. Bayesian integration in sensorimotor learning. *Nature* **427**, 244–247 (2004).
7. Maloney, L. T. & Zhang, H. Decision-theoretic models of visual perception and action. *Vision Research* **50**, 2362–2374 (2010).
8. Mamassian, P. Uncertainty and Confidence in Visual Perception. *i-Perception* **2**, 183–183 (2011).
9. Stocker, A. A. & Simoncelli, E. P. Noise characteristics and prior expectations in human visual speed perception. *Nat Neurosci* **9**, 578–585 (2006).
10. Cicchini, G. M., Mikellidou, K. & Burr, D. The functional role of serial dependence. *Proceedings of the Royal Society B* **285**, 20181722 (2018).
11. Feldman, J. Tuning your priors to the world. *Topics in cognitive science* **5**, 13–34 (2013).
12. Stone, J. v., Kerrigan, I. s. & Porrill, J. Where is the light? Bayesian perceptual priors for lighting direction. *Proceedings of the Royal Society B: Biological Sciences* **276**, 1797–1804 (2009).
13. Alais, D. & Burr, D. The Ventriloquist Effect Results from Near-Optimal Bimodal Integration. *Current Biology* **14**, 257–262 (2004).
14. Cicchini, G. M., Arrighi, R., Cecchetti, L., Giusti, M. & Burr, D. C. Optimal encoding of interval timing in expert percussionists. *J Neurosci* **32**, 1056–1060 (2012).

15. Geisler, W. S. Contributions of ideal observer theory to vision research. *Vision Research* **51**, 771–781 (2011).
16. Adler, W. T. & Ma, W. J. Comparing Bayesian and non-Bayesian accounts of human confidence reports. *PLOS Computational Biology* **14**, e1006572 (2018).
17. Morales, J. *et al.* Low attention impairs optimal incorporation of prior knowledge in perceptual decisions. *Atten Percept Psychophys* **77**, 2021–2036 (2015).
18. Rahnev, D. & Denison, R. N. Suboptimality in Perceptual Decision Making. *Behavioral and Brain Sciences* 1–107 (2018).
19. Stengård, E. & Berg, R. van den. Imperfect Bayesian inference in visual perception. *PLOS Computational Biology* **15**, e1006465 (2019).
20. Fischer, J. & Whitney, D. Serial dependence in visual perception. *Nature neuroscience* **17**, 738–743 (2014).
21. Pascucci, D. *et al.* Serial dependence in visual perception: A review. *Journal of Vision* **23**, 9 (2023).
22. Ceylan, G., Herzog, M. H. & Pascucci, D. Serial dependence does not originate from low-level visual processing. *Cognition* **212**, 104709 (2021).
23. Gallagher, G. K. & Benton, C. P. Stimulus uncertainty predicts serial dependence in orientation judgements. *Journal of Vision* **22**, 6–6 (2022).
24. You, F.-H., Gong, X.-M. & Sun, Q. Serial dependencies between form orientation and motion direction are asymmetric. *Frontiers in Psychology* **14**, (2023).
25. Beierholm, U., Rohe, T., Ferrari, A., Stegle, O. & Noppeney, U. Using the past to estimate sensory uncertainty. *eLife* **9**, e54172 (2020).
26. Denison, R. N., Adler, W. T., Carrasco, M. & Ma, W. J. Humans incorporate attention-dependent uncertainty into perceptual decisions and confidence. *Proceedings of the National Academy of Sciences* **115**, 11090–11095 (2018).

27. Fritsche, M., Spaak, E. & de Lange, F. P. A Bayesian and efficient observer model explains concurrent attractive and repulsive history biases in visual perception. *eLife* **9**, e55389 (2020).

28. Manassi, M., Murai, Y. & Whitney, D. Serial dependence in visual perception: a meta-analysis and review. *Journal of Vision* (in press).

29. Cousineau, D. Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorials in quantitative methods for psychology* **1**, 42–45 (2005).

30. Fritsche, M. To Smooth or not to Smooth: Investigating the Role of Serial Dependence in Stabilising Visual Perception [dissertation]. *Radboud University* (2016).

31. Burr, D. & Cicchini, G. M. Vision: Efficient Adaptive Coding. *Current Biology* **24**, R1096–R1098 (2014).

32. van Bergen, R. S. & Jehee, J. F. Probabilistic representation in human visual cortex reflects uncertainty in serial decisions. *Journal of Neuroscience* **39**, 8164–8176 (2019).

33. Kalm, K. & Norris, D. Visual recency bias is explained by a mixture model of internal representations. *Journal of Vision* **18**, 1 (2018).

34. Akaishi, R., Umeda, K., Nagase, A. & Sakai, K. Autonomous mechanism of internal choice estimate underlies decision inertia. *Neuron* **81**, 195–206 (2014).

35. Kok, P. & de Lange, F. P. Predictive coding in sensory cortex. in *An introduction to model-based cognitive neuroscience* 221–244 (Springer, 2015).

36. Pascucci, D. & Plomp, G. Serial dependence and representational momentum in single-trial perceptual decisions. *Sci Rep* **11**, 9910 (2021).

37. Rosen, R. Anticipatory Systems. in *Anticipatory Systems: Philosophical, Mathematical, and Methodological Foundations* (ed. Rosen, R.) 313–370 (Springer New York, 2012). doi:10.1007/978-1-4614-1269-4_6.

38. Trapp, S., Pascucci, D. & Chelazzi, L. Predictive brain: Addressing the level of representation by reviewing perceptual hysteresis. *Cortex* (2021).

39. Yon, D. *et al.* Stubborn predictions in primary visual cortex. *Journal of Cognitive Neuroscience* **35**, 1133–1143 (2023).

40. Nassar, M. R., Wilson, R. C., Heasly, B. & Gold, J. I. An Approximately Bayesian Delta-Rule Model Explains the Dynamics of Belief Updating in a Changing Environment. *J. Neurosci.* **30**, 12366–12378 (2010).

41. Pascucci, D., Rubega, M. & Plomp, G. Modeling time-varying brain networks with a self-tuning optimized Kalman filter. *PLOS Computational Biology* **16**, e1007566 (2020).

42. Gekas, N., McDermott, K. C. & Mamassian, P. Disambiguating serial effects of multiple timescales. *Journal of vision* **19**, 24–24 (2019).

43. Samaha, J., Switzky, M. & Postle, B. R. Confidence boosts serial dependence in orientation estimation. *Journal of vision* **19**, 25–25 (2019).

44. Markov, Y., Tiurina, N. & Pascucci, D. A matter of load: prior history effects scale with visual working memory load. Preprint at https://doi.org/10.31234/osf.io/6h2cy (2023).

45. Barbosa, J. & Compte, A. Build-up of serial dependence in color working memory. *Sci Rep* **10**, 10959 (2020).

46. Ceylan, G. & Pascucci, D. Attractive and repulsive serial dependence: The role of task relevance, the passage of time, and the number of stimuli. *J Vis* **23**, 8 (2023).

47. Fischer, C. *et al.* Context information supports serial dependence of multiple visual objects across memory episodes. *Nature communications* **11**, 1–11 (2020).

48. Houborg, C., Kristjánsson, Á., Tanrıkulu, Ö. D. & Pascucci, D. The role of secondary features in serial dependence. *Journal of Vision* **23**, 21 (2023).

49. Pascucci, D. *et al.* Laws of concatenated perception: Vision goes for novelty, decisions for perseverance. *PLOS Biology* **17**, e3000144 (2019).

50. Blondé, P., Kristjánsson, Á. & Pascucci, D. Tuning perception and decisions to temporal context. *iScience* 108008 (2023) doi:10.1016/j.isci.2023.108008.

51. Doerig, A. *et al.* Beyond Bouma's window: How to explain global aspects of crowding? *PLoS computational biology* **15**, e1006580 (2019).

52. Doerig, A., Schmittwilken, L., Sayim, B., Manassi, M. & Herzog, M. H. Capsule networks as recurrent models of grouping and segmentation. *PLoS computational biology* **16**, e1008017 (2020).

53. Tiurina, N., Markov, Y., Choung, O.-H., Herzog, M. H. & Pascucci, D. Unlocking crowding by ensemble statistics. *Current Biology* **32**, 4975-4981.e3 (2022).

54. Tiurina, N., Markov, Y., Whitney, D. & Pascucci, D. The functional role of spatial anisotropies in ensemble perception. Preprint at https://doi.org/10.31234/osf.io/yd3n8 (2023).

55. World Medical Association. World Medical Association Declaration of Helsinki: Ethical Principles for Medical Research Involving Human Subjects. *JAMA* **310**, 2191–2194 (2013).

56. Brainard, D. H. The psychophysics toolbox. *Spatial vision* **10**, 433–436 (1997).

57. Pascucci, D., Ruethemann, N. & Plomp, G. The anisotropic field of ensemble coding. *Scientific Reports* **11**, 1–10 (2021).

58. Akaike, H. A new look at the statistical model identification. *IEEE transactions on automatic control* **19**, 716–723 (1974).

59. Burnham, K. P. & Anderson, D. R. Multimodel inference: understanding AIC and BIC in model selection. *Sociological methods & research* **33**, 261–304 (2004).