Serial Dependence Predicts Generalization in Perceptual Learning

Noga Pinchuk-Yacobi,[1,2] Dov Sagi,[3] and Yoram S. Bonneh[1,2]
[1]School of Optometry and Vision Science, Faculty of Life Science, Bar-Ilan University, Ramat-Gan 5290002, Israel
[2]The Leslie and Susan Gonda Multidisciplinary Brain Research Center, Bar-Ilan University, Ramat Gan, Israel
[3]Department of Brain Sciences, The Weizmann Institute of Science, Rehovot 7610001, Israel

SUMMARY

Recent visual experiences systematically bias current perceptual reports, a phenomenon known as serial dependence. This study examines how serial dependence interacts with perceptual learning, specifically focusing on its role in learning generalization. We reanalyzed data from 50 observers who practiced the texture discrimination task (TDT) under three conditions designed to modulate learning generalization: consistent target location, alternating two target locations, and mixed with targetless trials (Harris et al., 2012); over 200,000 trials. Our analysis revealed substantial report biases toward the orientation of targets from previous trials, extending up to 10 trials back, much longer than previously reported. These biases, when accumulated across consistent trial sequences, nearly overshadowed responses to the current target. Remarkably, although TDT thresholds improved by 40% over the 8 days of training, the biases persisted despite the randomization of target orientations, with trial history being irrelevant to the current trial. The effect was most pronounced when current targets had low visibility, prior targets were highly visible, and prior targets appeared at the same location as the current target. Training conditions that facilitated learning generalization exhibited stronger and longer serial dependence effects from distant trial histories compared to conditions promoting specific learning. Across observers, stronger serial dependence from distant history was associated with greater learning transfer, suggesting that distant memory traces are essential for efficient learning. We propose that these extended memory traces integrate information across multiple trials to create a generalized learning template, preventing overfitting to specific stimuli thus enhancing learning transfer. In contrast, strong adaptation in the consistent target location condition disrupts these traces, thereby limiting generalization.

Introduction

Serial dependence in vision refers to the systematic bias of current perceptual reports toward recent visual experiences. This phenomenon has been observed across various visual domains, including orientation judgments (Fischer & Whitney, 2014), face recognition (Liberman et al., 2014), numerosity estimation (Fornaciai & Park, 2018) and variance perception (Suárez-Pinilla et al., 2018). By integrating past and present information, serial dependence is thought to enhance perception by leveraging the continuity of the visual world to maintain stable perception and make processing more efficient (Cicchini et al., 2018).

Despite their shared reliance on perceptual memory, serial dependence and perceptual learning have been studied separately. Serial dependence was shown to operate over short timescales, biasing perceptual reports based on recent stimuli, while perceptual learning involves sensory sensitivity improvements that develop over days of training, leading to long-term memory changes. This study explores the interaction between these two processes in the context of the texture discrimination task (TDT).
The effects of serial dependence in the TDT are likely complex, influencing both immediate perceptual performance and long-term learning dynamics. In the short term, biases toward the orientations of previous targets—randomly varying and irrelevant to the current trial—may impair performance by introducing decision noise. However, integrating information across trials may also facilitate decision-making by leveraging temporal variations inherent in the visual system. Over longer timescales, serial dependence effects may evolve within and between daily sessions potentially being learned out to reduce decision noise. Of particular interest is how serial dependence influences learning specificity—a key characteristic of perceptual learning where improvements are closely tied to practiced stimulus features, such as location or orientation (Sagi, 2011).

To investigate this relationship, we reanalyzed data from Harris, Gliksberg, and Sagi (2012), in which observers practiced the TDT under conditions designed to modulate learning generalization. Their findings showed that learning generalized to a new, untrained location when targets were presented randomly across two locations or intermixed with targetless ‘dummy’ trials. Conversely, learning became location-specific when targets consistently appeared in a single location, likely due to increased sensory adaptation, consistent with earlier findings (Karni & Sagi, 1991). The study proposed that spatial invariance depends on the level of adaptation: unadapted low-level networks support space-invariant learning, whereas adaptation induces localized changes that disrupt this invariance. However, the mechanisms through which unadapted networks enable space-invariant learning and how adaptation interferes remain unclear. By reexamining this dataset through the lens of serial dependence, we aim to provide further explanation of these mechanisms and uncover how recent visual experiences influence both immediate perceptual reports and long-term perceptual learning.

Methods

We reanalyzed data from 50 observers who participated in the texture discrimination task (TDT) as described by Harris, Gliksberg, and Sagi (2012). In this dual-task experiment, observers identified the orientation (vertical or horizontal) of a target composed of three peripheral diagonal lines embedded within a uniform background of horizontal lines while simultaneously performing a forced-choice letter discrimination task (T vs L at the center of the stimulus) to maintained fixation (Figure 1A). The experiment consisted of four daily sessions with the target presented at one location (or two fixed locations in the 2loc condition), followed by four additional daily sessions at a second location (or a second pair of fixed locations in the 2loc condition).

Participants were assigned to one of three experimental conditions:

- **1loc condition**: The target consistently appeared at the same fixed location across all trials.
- **2loc condition**: The target randomly appeared at one of two diagonally opposite locations with equal eccentricity, alternating between trials.
- **Dummy condition**: The target appeared at a fixed location, but genuine target trials were randomly interleaved with 'dummy' trials, where no target was present (replaced by background elements).

For detailed specifications of the groups assigned to each condition, see Supplementary Table S1. Our analysis quantified serial dependence effects (SDE) to examine how prior visual experiences influenced current perceptual reports (Figure 1B) and whether these serial dependencies affected learning specificity across the different experimental conditions. To quantify SDE, we used a linear mixed effects (LME) model (Figure 1C), which evaluated the influence of prior target stimuli (up to 10 trials back) on current trial reports. The model's weights (W_n, referred to as W-bias) represented the magnitude of the serial dependence effect for each N-back stimulus. This relationship was modeled by the following equation:

$$P(T_0|T_0) = P_0 + \Sigma (W_n \cdot sT_n) \qquad \text{for } n=1 \text{ to } 10$$

Where:

- $P(T_0|T_0)$: The probability of reporting target T_0 (=±1) when the target is T_0.
- P_0: The history independent report probability.
- sT_n: Equals 1 if the orientation of the n-back trial matches the current trial orientation, and -1 if not.
- W_n: The added report probability (bias) associated with the N-back trial.

Fixed effects accounted for the orientations of these stimuli, while a random intercept captured individual differences across observers. Interaction terms were excluded, as statistically significant interactions were small and inconsistent.

For analyses involving individual observers, weights were calculated using an equalized number of visible history trials, accounting for individual differences in threshold by limiting the analysis to trials with SOA > threshold and SOA < threshold + 140 ms.

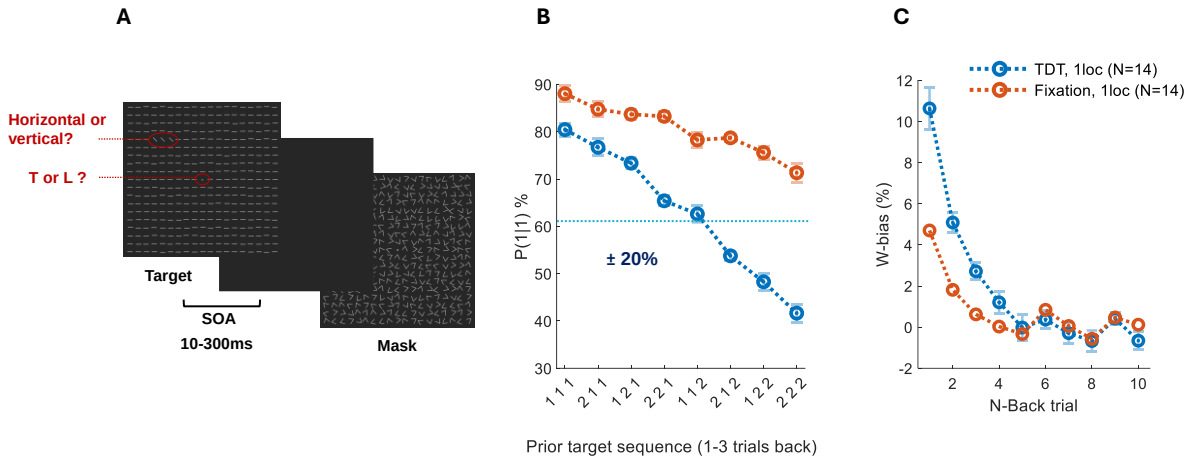

Figure 1: Texture Discrimination Task (TDT) and Serial Dependence Effect (SDE) (A) TDT Trial Sequence: Observers identified the orientation (vertical or horizontal) of a target composed of three diagonal lines embedded within a background of horizontal lines while simultaneously performing a forced-choice letter discrimination task (T vs L) to maintain fixation. The 10 ms target frame was followed by a 100 ms patterned mask, with the stimulus onset asynchrony (SOA) between target and mask (10 to 300 ms) randomized across trials. (B) SDE by performance: Correct responses ('Hits') increased by approximately 20% when the current target orientation matched the preceding three orientations (e.g., '1' preceded by '111') and decreased by about 20% with mismatching orientations (e.g., '1' preceded by '222'), relative to average performance (dotted light blue line). '1' and '2' represent the two possible orientations—either vertical and horizontal, or vice versa. The fixation task (red) showed much smaller biases, likely due to its higher overall performance. (C) SDE via Linear Mixed Effects (LME) Weights (W-bias, %): Influence of 1-10 back trials on current report. The ±20% bias for 1-3 back trials shown in panel B corresponds to the sum of W-bias (%) from the 1st, 2nd, and 3rd prior trials. Panels B and C include data from the 1loc condition (N=14) pooled across all training days, including only trials with low-visibility current targets (SOA < threshold + 20 ms). Blue represents the texture discrimination task, and red indicates the letter discrimination (fixation control) task.

Results

We found a significant 15% bias in observers' reports toward the orientation of the immediately preceding target (1-back), with cumulative bias reaching approximately 40% when considering up to 10 prior targets (SDE-10 = sum of 1-10 W-bias weights, expressed as percentages) under specific filtering conditions that amplified bias (Figure 2A). Importantly, these biases were not attributable to motor responses, as they persisted in a dual-task setup and were influenced by the spatial location and visibility of target stimuli, independent of motor actions (Figure 2).

Biases were most pronounced under conditions where the current target had low visibility (SOA < observer's threshold + 20 ms) and prior targets were highly visible (SOA > threshold) with SDE-10 reaching 40 ± 3 (Figure 2A). Biases were significantly reduced when the current targets were highly visible (SDE-10 = 5 ± 1; reduction: 36 ± 3; $t(49) = 11.81$, $p < 0.0001$, Cohen's d =1.67) or when prior targets had low visibility

(SDE-10 = 5 ± 2; reduction: 36 ± 3; $t(49) = 13.38$, $p < 0.0001$, Cohen's d = 1.89). Additionally, in the 2loc condition, biases were predominantly location-selective, with significantly stronger effects when trial history originated from the same location (SDE-10 = 41 ± 3) compared to a diagonally opposite location (SDE-10 = 15 ± 2; reduction: 26 ± 4; $t(13) = 7.19$, $p < 0.0001$, Cohen's d = 1.92; Figure 2B).

Biases gradually decayed over successive trials but remained substantial, extending far into trial history (Figure 4A). In the 1loc condition, biases were significant up to 4 trials back ($p < 0.001$ for $N \leq 4$). In the 2loc condition, biases persisted longer, remaining significant up to 8 trials back ($p < 0.001$ for $N \leq 7$; $p < 0.05$ for $N = 8$). The dummy condition showed the most prolonged biases, with significant effects extending up to 10 trials back ($p < 0.001$ for $N \leq 7$; $p \leq 0.01$ for $N = 9, 10$), although the 8-back trial was not significant ($p = 0.77$). Our analysis focused on decay across trials rather than elapsed time, as doubling inter-trial durations had no impact on the pronounced short-history effects but only attenuated the much smaller long-history effects (as observed when comparing the two groups in the 1loc condition; see Supplementary Material Table S1).

Despite randomized target orientations across trials—rendering past orientations irrelevant to current judgments and introducing decision noise— daily biases remained strong and highly significant ($p < 0.0001$ for each day; Figure 3). Across the 8 days, biases showed no significant overall change ($F(7, 343) = 1.99$, $p = 0.066$, repeated measures ANOVA, Greenhouse-Geisser corrected), even as TDT thresholds improved significantly from Day 1 (126 ± 6 ms) to Day 8 (80 ± 2 ms; $F(7, 343) = 41.77$, $p < 0.0001$; Figure 3). No significant correlation was observed between biases and thresholds across days ($r = 0.28$, $p = 0.50$), suggesting that improvements in performance were independent of biases. Biases were consistent across locations, with comparable magnitudes at Location 1 (SDE-10 = 41 ± 3) and Location 2 (SDE-10 = 40 ± 3; $t(49) = 0.14$, $p = 0.89$). At the individual level, biases were highly correlated between the two locations ($r = 0.62$, $p < 0.0001$), suggesting stable observer-specific traits.

Within-session analyses across all days showed that biases remained significant throughout sessions but decreased by ~20% from the beginning to the end of sessions. Biases in the first third of trials (SDE-10 = 43 ± 3, $t(49) = 14.4$, $p < 0.0001$) were significantly higher than in the last third (SDE-10 = 36 ± 4, $t(49) = 9.9$, $p < 0.0001$; reduction: 7 ± 3, $t(49) = 2.23$, $p < 0.05$). This reduction was primarily driven by SDE-distant (sum of 4–6 W-bias weights, expressed as percentages), which decreased significantly from start to end of sessions (10 ± 1 to 7 ± 1; $t(49) = 2.42$, $p < 0.05$), while SDE-recent (sum of 1–3 W-bias weights, expressed as percentages) remained stable (28 ± 2 to 28 ± 2; $t(49) = 0.09$, $p = 0.93$). Such behavior is consistent with sensory adaptation developing during the testing session.

A comparison of biases across the three experimental conditions showed similar magnitudes for recent history (1–3 back trials) across all conditions (SDE-recent: dummy: 28 ± 3; 1loc: 26 ± 2; 2loc: 25± 2; $F(2,47) = 0.18$, $p = 0.83$). In contrast, biases from distant history (4–6 back trials) differed significantly between conditions (SDE-distant: dummy: 12 ± 2; 2loc: 10 ± 1; 1loc: 1 ± 2; $F(2,47) = 13.58$, $p < 0.0001$). Post-hoc analyses indicated that distant history biases were significantly lower in the 1loc

condition compared to both the dummy condition ($p < 0.0001$) and the 2loc condition ($p = 0.001$), likely due to enhanced sensory adaptation from consistent target presentation in the 1loc setup. In the letter discrimination (fixation control) task, where identical foveal stimuli were presented across all conditions, no significant differences were observed between conditions for either recent history biases (dummy: 6 ± 2; 1loc: 8 ± 2; 2loc: 8 ± 2; $F(2,47) = 0.40$, $p = 0.674$) or distant history biases (dummy: 2 ± 1; 1loc: 2 ± 1; 2loc: 3 ± 1; $F(2,47) = 0.59$, $p = 0.56$). The smaller biases observed in this task likely resulted from higher overall performance levels and ceiling effects.

We propose that stronger distant serial dependence in the '2loc' and' dummy' conditions facilitated greater generalization to a new location, consistent with findings by Harris et al. (2012). This hypothesis is further supported by a significant correlation between stronger distant serial dependence and increased perceptual learning transfer across observers ($r = 0.55$, $p = 0.0001$, $n = 48$, two outliers exceeding 3 SD were excluded), while no significant correlation was found for recent serial dependence ($r = -0.19$, $p = 0.19$, $n = 50$; Figure 5).

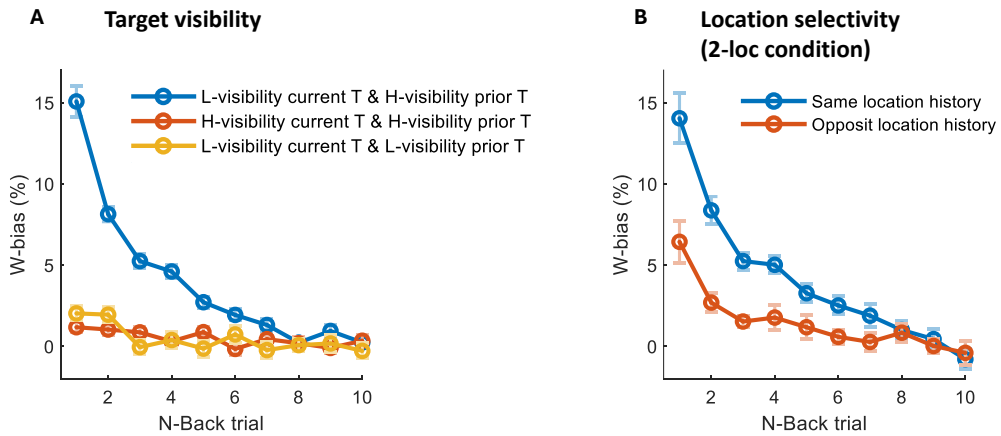

Figure 2: Conditions Enhancing SDE

(A) Target Visibility: SDEs across trial history (W-bias) were stronger when current targets had low visibility (SOA < threshold + 20 ms) and prior targets had high visibility (SOA > threshold; blue line). Effects significantly decreased when current targets were highly visible (SOA > threshold + 20 ms; red line) or when prior targets had low visibility (SOA < threshold; yellow line)(N=50). (B) Location Specificity: SDEs were significantly stronger when the trial history originated from the same screen location as the current target (blue line) compared to a diagonally opposite location (red line) (N=14). Error bars indicate the standard error of the mean.

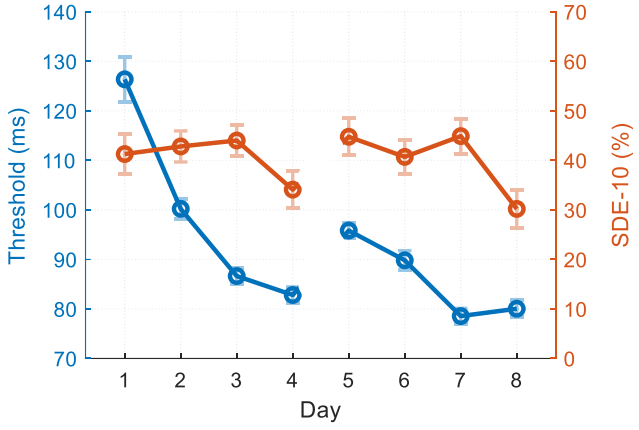

Figure 3: Changes in SDEs and Thresholds Across Days

Daily serial dependence effects (SDE-10, red) remained strong and significant throughout all eight days, with a reduction on the final training days at each location (days 4 and 8) that did not reach statistical significance, despite the significant improvements in texture discrimination task (TDT) thresholds (blue). Target location was changed after day 4.

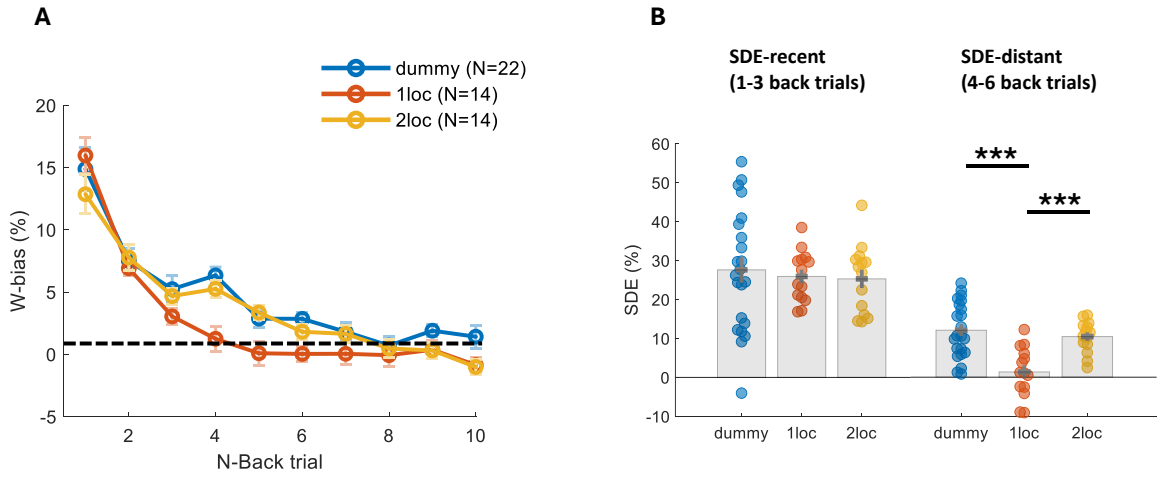

Figure 4: SDEs Across Experimental Conditions and Trial History

(A) Mean SDEs across observers are shown as a function of trial history (N-back) for the three experimental conditions (1loc: red, 2loc: yellow, dummy: blue). Data points on or above the black dashed line indicate statistical significance ($p < 0.001$ for $N \leq 7$; $p < 0.05$ for $N = 8$; $p \leq 0.01$ for $N = 9, 10$).

(B) SDEs from 50 observers are analyzed separately for SDE-recent (left, sum of 1-3 W-bias weights) and SDE-distant (right, sum of 4-6 W-bias weights), both expressed as percentages. Recent SDEs were consistent across conditions, whereas distant SDEs were significantly stronger in the 2loc and dummy conditions compared to the 1loc condition (***$p \leq 0.001$).

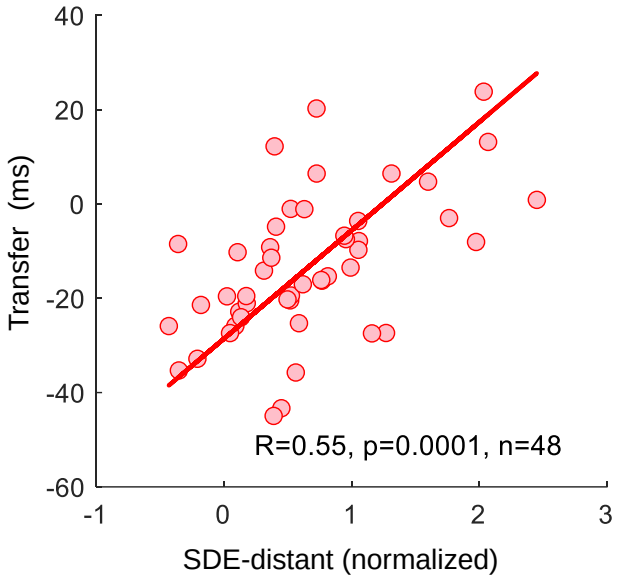

Figure 5: Relationship Between Distant Serial Dependence and Learning Transfer

Observers (N=-48) with stronger distant SDE (4–6 back trials) showed greater learning transfer to a second location (linear correlation, p<0001). Learning transfer was calculated as the difference in thresholds between Day 4 (final day at the first location) and Day 5 (initial day at the second location), with negative values corresponding to performance loss. Distant serial dependence effects were normalized to 1-back trial effects to represent each observer's time constant. Two observers with values exceeding 3 SD were excluded.

Discussion

Our investigation into serial dependence in the texture discrimination task (TDT) reveals significant biases toward prior perceptual reports. Despite improvements in perceptual thresholds with training, these biases remained stable, even though target orientations were randomized across trials and trial history was irrelevant to the current trial. This finding points to a robust, hard-wired feature of visual perception that is unmodifiable by training and independent of current stimulus statistics.

Importantly, our study shows that serial dependence extends beyond short-term biases to influence long-term learning processes, particularly affecting how learning generalizes to new, untrained locations. Stronger distant biases were found in experimental conditions that exhibited learning generalization (‘2loc’, ‘dummy’), with observers showing stronger biases from distant history (4-6 back) demonstrating greater learning generalization. Compared to previous studies, our data reveal an exceptionally long memory trace, extending up to 8-back stimuli in the 2loc condition and to 10-back stimuli in the dummy condition, which exceeds the typical 3-back limit reported in previous studies (Fischer & Whitney, 2014; John-Saaltink et al., 2016; Lau & Maus, 2019; Manassi et al., 2019, 2023). This extended trace may be due to reduced sensory adaptation in these conditions.

Consistent with previous research, we found that the reliability of both current and prior targets affected bias magnitude, but in opposing directions. Biases were stronger with less reliable (low-visibility) current targets (Ceylan et al., 2021; Cicchini et al., 2017, 2018; Manassi et al., 2018) and with highly reliable (high-visibility) prior targets (Pascucci et al., 2019; Van Bergen & Jehee, 2019), aligning with Bayesian perceptual models (Kersten et al., 2004; Knill & Pouget, 2004). Additionally, we observed location selectivity in the 2loc condition, where biases were stronger when prior targets appeared at the same location as the current target (Collins, 2019; Fischer & Whitney, 2014; Fornaciai & Park, 2018; John-Saaltink et al., 2016; Manassi et al., 2019).

Our results uncover a link between serial dependence and perceptual learning. We show that consistent stimulus repetition ('1loc' condition), which induces sensory adaptation and location-specific learning (Censor & Sagi, 2009; Harris et al., 2012; Ofen et al., 2007), limits the temporal span available for serial dependence, likely due to sensory inhibition. In contrast, experimental conditions that promote learning transfer ('2loc', 'dummy'; Harris et al., 2012) showed an extended temporal span of serial dependence. These results offer a new insight into the mechanisms of perceptual learning, suggesting a single mechanism that controls learning specificity. Previous accounts of learning specificity assumed learning to be local to the cortical site encoding the target (Karni & Sagi, 1991), or to location specific perceptual templates used for decision (Dosher & Lu, 1999), or to both (Watanabe & Sasaki, 2015). Here we suggest an account based on a single decision template that learns the discrimination task by classifying neuronal response features as signaling vertical or horizontal texture targets. These decision templates operate on eccentricity-dependent cortical representations and generalize across equal-eccentricity retinal locations but not across different retinal eccentricities (Harris & Sagi, 2018). However, when trained with targets in a fixed location, the decision template may become overfitted to specific characteristics useful only within the trained location, thus limiting generalizability (Sagi, 2011). For such a classification process to be robust, and of general use, many samples (trials) need to be integrated efficiently thus removing noisy response features. Our results show that decision biases are integrated linearly over trials, suggesting efficient feature integration over many trials in experimental conditions yielding learning generalization. Reduced temporal integration may generate a classifier that relies on spurious local features, possibly tabulating local features that are specific to the trained location (Mollon & Danilova, 1996).

We suggest that previous experimental results showing learning generalization can be explained by a modulation of temporal integration. This includes short training phases that are stopped before adaptation takes over, showing generalization to other retinal locations (Censor & Sagi, 2009; Karni & Sagi, 1993), short pre-training phases enabling generalization across visual tasks (Zhang et al., 2010) and other paradigms that effectively reduce sensory adaptation and by that serial dependence (reviewed in Sagi, 2011). Our results have strong implications for current theories of perceptual learning. These theories explain learning specificity by modulation of specific neuronal wiring, either at the encoding stage (visual cortex) or at a higher processing level between encoding and decision. Here we suggest a single mechanism that can either be specific or general, regardless of the anatomical site where neuronal modifications take place.

REFERENCES

Censor, N., & Sagi, D. (2009). Global resistance to local perceptual adaptation in texture discrimination. *Vision Research*, *49*(21), 2550–2556.

Ceylan, G., Herzog, M. H., & Pascucci, D. (2021). Serial dependence does not originate from low-level visual processing. *Cognition*, *212*, 104709.

Cicchini, G. M., Mikellidou, K., & Burr, D. (2017). Serial dependencies act directly on perception. *Journal of Vision*, *17*(14), 6.

Cicchini, G. M., Mikellidou, K., & Burr, D. C. (2018). The functional role of serial dependence. *Proceedings of the Royal Society B*, *285*(1890), 20181722.

Collins, T. (2019). The perceptual continuity field is retinotopic. *Scientific Reports*, *9*(1), 18841.

Dosher, B. A., & Lu, Z.-L. (1999). Mechanisms of perceptual learning. *Vision Research*, *39*(19), 3197–3221.

Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, *17*(5), 738–743.

Fornaciai, M., & Park, J. (2018). Serial dependence in numerosity perception. *Journal of Vision*, *18*(9), 15.

Harris, H., Gliksberg, M., & Sagi, D. (2012). Generalized perceptual learning in the absence of sensory adaptation. *Current Biology*, *22*(19), 1813–1817.

Harris, H., & Sagi, D. (2018). Visual learning with reduced adaptation is eccentricity-specific. *Scientific Reports*, *8*(1), 608.

John-Saaltink, E. S., 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.

Karni, A., & Sagi, D. (1991). Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. *Proceedings of the National Academy of Sciences*, *88*(11), 4966–4970.

Karni, A., & Sagi, D. (1993). The time course of learning a visual skill. *Nature*, *365*(6443), 250–252.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annu. Rev. Psychol.*, *55*(1), 271–304.

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, *27*(12), 712–719.

Lau, W. K., & Maus, G. W. (2019). Visual serial dependence in an audiovisual stimulus. *Journal of Vision*, *19*(13), 20.

Liberman, A., Fischer, J., & Whitney, D. (2014). Serial dependence in the perception of faces. *Current Biology*, *24*(21), 2569–2574.

Manassi, M., Kristjánsson, Á., & Whitney, D. (2019). Serial dependence in a simulated clinical visual search task. *Scientific Reports*, *9*(1), 19937.

Manassi, M., Liberman, A., Kosovicheva, A., Zhang, K., & Whitney, D. (2018). Serial dependence in position occurs at the time of perception. *Psychonomic Bulletin & Review*, *25*, 2245–2253.

Manassi, M., Murai, Y., & Whitney, D. (2023). Serial dependence in visual perception: A meta-analysis and review. *Journal of Vision*, *23*(8), 18.

Mollon, J. D., & Danilova, M. V. (1996). Three remarks on perceptual learning. *Spatial Vision*, *10*(1), 51–58.

Ofen, N., Moran, A., & Sagi, D. (2007). Effects of trial repetition in texture discrimination. *Vision Research*, *47*(8), 1094–1102.

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*(3), e3000144.

Sagi, D. (2011). Perceptual learning in vision research. *Vision Research*, *51*(13), 1552–1566.

Suárez-Pinilla, M., Seth, A. K., & Roseboom, W. (2018). Serial dependence in the perception of visual variance. *Journal of Vision*, *18*(7), 4.

Van Bergen, R. S., & Jehee, J. F. M. (2019). Probabilistic representation in human visual cortex reflects uncertainty in serial decisions. *Journal of Neuroscience*, *39*(41), 8164–8176.

Watanabe, T., & Sasaki, Y. (2015). Perceptual learning: toward a comprehensive theory. *Annual Review of Psychology*, *66*(1), 197–221.

Zhang, T., Xiao, L.-Q., Klein, S. A., Levi, D. M., & Yu, C. (2010). Decoupling location specificity from perceptual learning of orientation discrimination. *Vision Research*, *50*(4), 368–374.