Temporal dynamics of implicit memory underlying serial dependence

Authors:
Cristiano Moraes Bilacchi [1], Esaú Ventura Pupo Sirius [1], André M. Cravo [1], Raymundo Machado de Azevedo Neto [2]

[1] Center for Mathematics, Computing, and Cognition, Federal University of ABC (UFABC), São Paulo, Brazil
[2] Brain Institute, Hospital Israelita Albert Einstein, São Paulo, Brazil

Corresponding author:
Raymundo Machado de Azevedo Neto, email: raymundo.neto@einstein.br

Author contributions statement: RMAN and AMC conceived the experiment. CMBA performed the experiments and analysed the data. CMBA, ESVP, AMC, RMAN wrote, reviewed the manuscript, and approved the final version of the manuscript.

Acknowledgements: We thank Matthias Nau, Gustavo B. Azevedo and Mateus Silvestrin for their insightful comments on earlier versions of the manuscript. AMC was supported by grant #2017/25161-8, São Paulo Research Foundation (FAPESP). ESVP was supported by grant #2019/06423-7, São Paulo Research Foundation (FAPESP). RMAN was supported by grant #2018/08318-3, São Paulo Research Foundation (FAPESP). CMBA was supported by grant #2018/800254, National Council for Scientific and Technological Development (CNPq).

Conflict of interest: The authors declare no competing financial interests.

Abstract

Serial dependence is the effect in which the immediately preceding trial influences participants' responses to the current stimulus. But for how long does this bias last in the absence of interference from other stimuli? Here, we had 20 healthy young adult participants (12 women) perform a coincident timing task using different inter-trial intervals to characterize the serial dependence effect as the time between trials increases. Our results show that serial dependence abruptly decreases after 1 s inter-trial interval, but it remains pronounced after that for up to 8 s. In addition, participants' response variability slightly decreases over longer intervals. We discuss these results in light of recent models suggesting that serial dependence might rely on a short-term memory trace kept through changes in synaptic weights, which might explain its long duration and apparent stability over time.

Keywords: serial dependence, short-term memory, implicit memory, coincident timing task, visuomotor integration.

Statement of Relevance:

Recent perceptual and motor experiences bias human behavior. For this serial bias to take place, the brain must keep information for at least the time between events to blend past and current information. Understanding the temporal dynamics of such memory traces might shed light into the short-term memory mechanism and integration of prior and current information. Here, we characterized the temporal dynamics of the serial biases that emerge in a visuomotor task by varying the length of the interval between successive events. Our results show response biases are still present even after intervals as long as 8 s and that participants' response variability decreases over time. Serial dependence thus seems to rely on a memory mechanism that is both long lasting in the absence of interference and stable.

Introduction

Perception and action can be biased due to the characteristics of recently presented stimuli. This serial dependence effect has been shown in the visual domain for a wide range of features, from low-level orientation (e.g., Liberman et al. 2016, Fischer and Whitney 2014) and speed perception (Makin et al. 2008; Kwon and Knill, 2013), to high-level face perception (Liberman et al. 2014). To exert its effect on the current trial, information from the previous trial's stimulus should be stored in memory. However, it is unclear how long this memory trace lasts.

Previous investigations on the temporal dynamics of stimulus information kept in short-term memory have shown that it can be stored with good levels of detail for up to 30 s (Magnussen & Greenlee 1992; Blake 1997), but that its precision decreases over time (Rademaker et al. 2018; Shin et al. 2017) or its representation drifts (Schneegans and Bays, 2018; Wolff et al. 2020). However, these studies have investigated short-term memory effects when keeping stimulus information in memory was both task-relevant and explicit. On the other hand, serial dependence emerges even when information from previous trials is task-irrelevant (Lieberman et al. 2016), suggesting that this bias does not depend on the active maintenance of information in memory. Critically, the temporal dynamics of such implicit memory traces are still mostly unknown.

Recent studies have shown that events from up to 4 trials into the past can modulate performance in the current trial (Kalm and Norris, 2018; Fischer and Whitney, 2014). Nevertheless, these experiments could not evaluate the pure progress of memory, given that the serial dependence effect was evaluated over intervening trials, which might have overwritten memory traces (Kalm and Norris 2018; Lewandowsky, Oberauer, Brown, 2009). When investigating the sole effect of time on serial dependence, Bliss et al. (2017) verified that after 10 s the previous stimulus would no longer bias current responses. In contrast, Papadimitrou et al. (2015) have shown that serial dependence persists for up to 6 s intervals. In both these studies, participants had to explicitly encode the current trial stimuli in short-term memory to perform the task. To characterize the effect and temporal dynamics of implicitly encoded information on serial dependence, we had participants perform a coincident timing task, in which they only had to act upon the current stimulus, without requiring them to encode items in short-term memory explicitly.

Materials and Methods

Transparency and data availability

Task and analysis code and raw data will be made openly available, following acceptance of the manuscript. The procedures and analysis were not pre-registered prior to the research being conducted.

Participants

Twenty participants (12 women, 21±4 years old, mean ± standard deviation) were informed about the experimental procedures and signed a Consent Form. We recruited a convenience sample of college students. The experimental protocol was approved by The Research Ethics Committee of the Federal University of ABC. All experiments were

performed following the approved guidelines and regulations. All participants had normal or corrected to normal vision.

Stimulus and task

Participants performed a coincident timing task in which they pressed a button when a rightward moving target hit an interception zone. Participants sat in a quiet, dark room, at a distance of approximately 65 cm from a ViewPixx display monitor (120 Hz refresh rate, resolution of 1920 × 1080 pixels), resting their heads on a chin rest. Stimuli and tasks were generated and controlled using GNU Octave (https://www.gnu.org/software/octave/) and Psychtoolbox (Brainard, 1997; Pelli, 1997; Kleiner et al., 2007), running on a Linux operating system. On each trial, participants saw a red square target (0.5 degrees of visual angle [dva]) that moved from left to right of the screen at a constant speed (20, 22, 24, 26, 28 dva/s) on a gray background (Figure 1A). The beginning of each trial was cued by presenting the start zone (squared black letter "C") on the left side of the screen and the interception zone (vertical black line fitting 30% of the screen vertically), separated by 10 dva. Simultaneously, the target was flashed for 200 ms to warn participants the trial was about to begin. After 400 ms, the target was presented and started moving. Participants were instructed to press a button at the same time the target hit the interception zone (button box sampling frequency 1000 Hz). After 500 ms of target arrival, the target, start and interception zones were removed from the screen. Throughout the experiment, participants were instructed to fixate their gaze at a cross (0.5 dva) on the center of the screen, 0.35 dva below the target's path. The fixation cross was present on the screen for the entire experiment.

Experimental design and procedures

Before the main experiment, participants were instructed and familiarized with the task by performing 20 trials. Each target speed was randomly presented four times during practice. To evaluate the dynamics of the serial dependence effect, we had participants perform the coincident timing task in five separate blocks with different inter-trial intervals (ITI): 0.1, 1, 2, 4, and 8 s. The ITI was defined as the interval between the disappearance of the start and interception zones and its reappearance on the next trial.

Each block consisted of 5 mini-blocks containing 51 trials each. Within each mini-block, the target speeds were counterbalanced, such that each target speed was preceded by every other speed the same number of trials (Brooks, J.L., 2012). This sequence of trials assured that the current and previous target speed were independent of each other. Participants had 30 s rest periods between mini-blocks. Participants were allowed to take longer breaks between blocks. The experimental session took approximately 2 hours and comprised a total of 1275 trials.

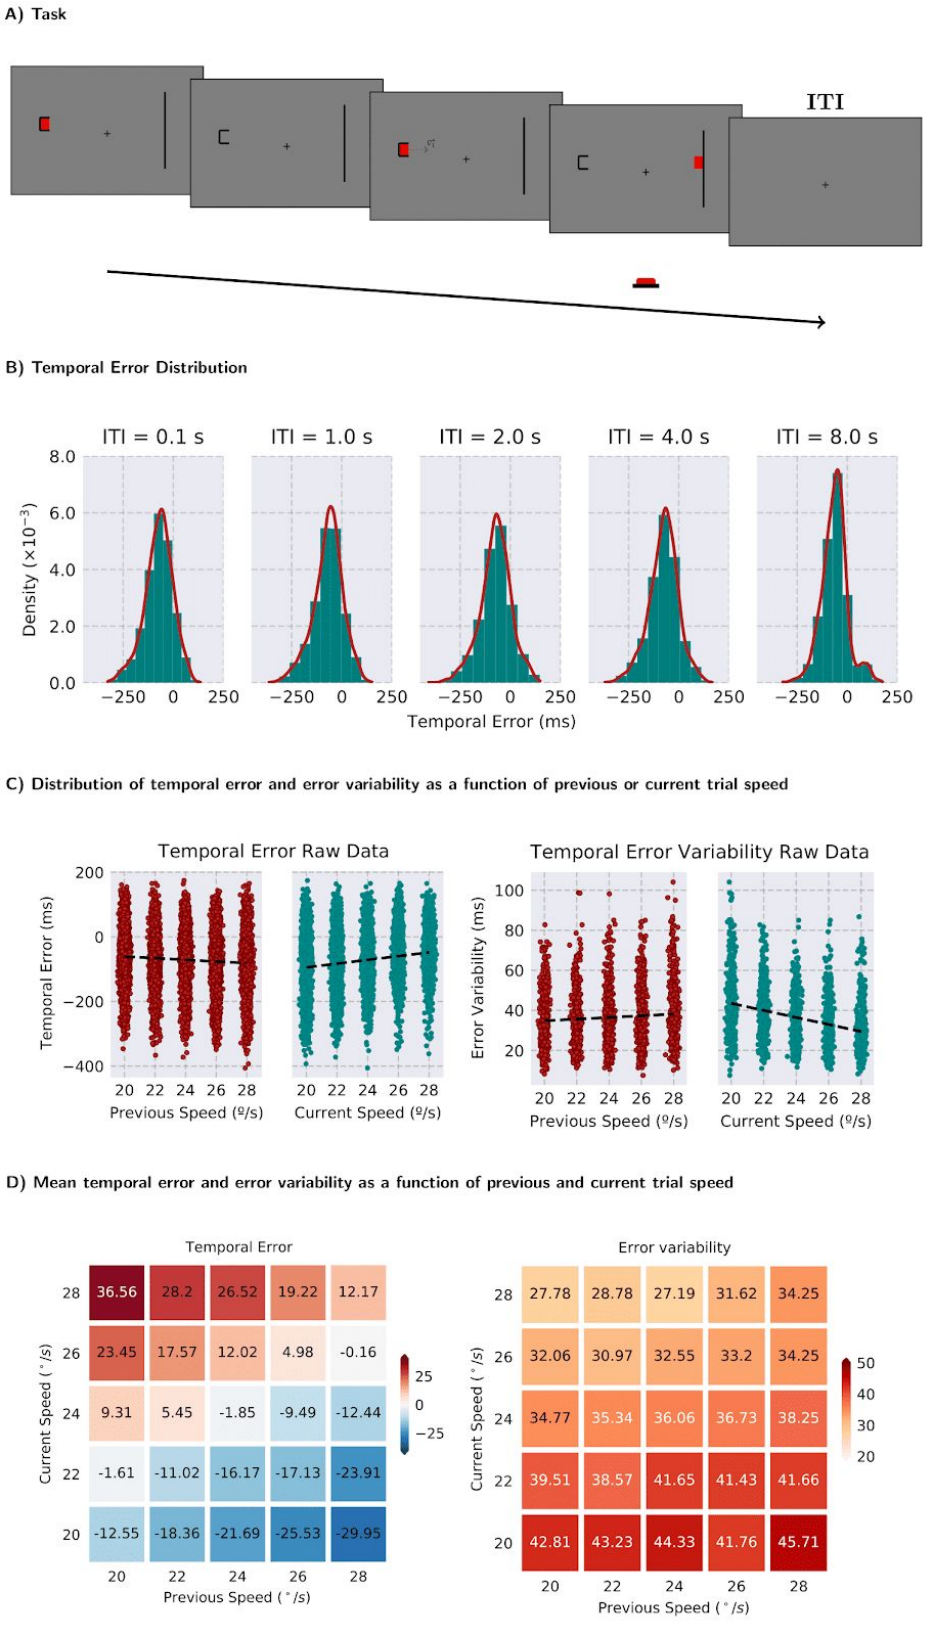

Figure 1. (A) Phases of a trial in the experiment. Participants were instructed to press the response button at the same time the target reaches the vertical bar. (B) Temporal error density for all participants collapsed by ITI. Values centered at zero and ranging between -200 and 200 milliseconds indicate that participants were performing a coincident timing task accordingly and not waiting for the target to reach the vertical bar to press the button. (C) Distribution of temporal error and error variability across participants as a function of previous (red) and current (blue) target speeds. Despite variability across participants, temporal errors

tend to be more negative when the previous trial is faster and more positive when the current trial is faster. The error variability follows an opposite pattern: the greater the current trial speed, the lower the error variability across trials, and the lower the previous trial speed, the lower the error variability across trials. (D) Mean temporal error (left) and error variability (right) for all volunteers as a function of previous (horizontal axis) and current (vertical axis) target speed. The temporal error becomes more positive as the current target is faster (cells get redder from bottom to top), but also becomes more negative as the previous speed is faster (cells get bluer from left to right). The error variability shows the opposite pattern from temporal error.

Analysis

We calculated participants' temporal error (in ms) on each trial, i.e., the temporal difference between the time participants pressed the button and the target arrival at the interception zone. Negative values indicate early responses, whereas positive values indicate delayed responses. In addition to the temporal error, we also estimated participants' error variability by calculating the standard deviation of participants' temporal error on every combination of previous and current trial speeds for each ITI.

For each ITI and participant, we first removed outlier trials by calculating the median absolute deviation (MAD) and excluding all trials where the deviation was greater than 3 (Leys et. al, 2013; Rousseeuw et. al, 1993). We removed one participant from the 8 s ITI condition because more than 15% of trials were removed. Overall, 2.81% of trials were removed with this procedure.

We evaluated serial dependence by separately modeling participants' temporal error for each ITI using multiple linear regression. We modeled the current trial temporal error as a function of target speed on the current and previous trials. We used the coefficient of the previous trial regressor as our main indication of serial dependence for both temporal error and error variability. For temporal errors, negative slopes indicate an attractive serial bias toward the speed of the previous trial. In other words, the greater the speed of the previous trial, the more participants tend to anticipate their responses. In order to evaluate the presence of serial dependence for each ITI, we performed one-tailed one-sample t-tests on regression slopes. To evaluate the serial dependence dynamics over ITIs, previous trial slopes for each ITI were entered in a one-way repeated-measures ANOVA. A Helmert contrast was used to follow up on significant ANOVA results (alpha = 0.05). The analysis was performed on Python and statistical analysis was performed on JASP.

To check whether the serial dependence effect was not due to spurious correlations, we regressed the current trial error with the next trial speed expecting the absence of an effect (Fischer and Whitney, 2014).

Results

As expected, participants' temporal errors were centered around zero. Because merely reactive responses would lag behind, this suggests participants were indeed using the time to contact to perform the task rather than waiting for the target to hit the vertical bar (Figure 1B). As shown in Figures 1C-D, there was a general tendency of temporal errors to increase as a function of current speed and decrease as a function of the previous speed.

Error variability decreased as a function of the current target speed and increased as a function of the previous target speed. Our control analysis regressing current trial response errors by the next trial speed showed absence of evidence for this effect (ITI = 0.1s: t(19) = 0.357 , p = 0.712; ITI = 1 s: t(19) = 1.569, p = 0.133; ITI = 2 s: t(19) = 1.345, p = 0.195; ITI = 4 s: t(19) = 1.435, p = 0.168 ; ITI = 8 s: t(19) = 0.527, p = 0.605).

To verify the serial dependence effect across ITIs, we plotted the average temporal error across participants as a function of current (vertical axis) and previous (horizontal axis) target speeds (Figure 2A). As shown in previous studies, there is a clear attractive pattern of serial dependence. The faster the speed of the previous trial, the more participants presented anticipated responses. This pattern can be seen by the change from a darker red to lighter red color for the fastest current trial speed (28 dva/s) or by the change from red to blue for the intermediate current trial speed (24 dva/s) (Figure 2A). This pattern is present for all current trial speeds, but at different ranges of temporal error, as faster current trial speeds induced later responses. Importantly, we found a consistent serial dependence effect for all ITI conditions.

To assess these effects in more detail, we first analyzed the previous speed slopes obtained by regressing temporal error by previous and current trial speeds using multiple linear regressions performed for each participant at each ITI condition. Again, serial dependence was observed for all tested ITIs, as evidenced by negative slopes for previous trial speed [ITI = 0.1 s: t(19) = –9.632, p < 0.001 , Cohen' *d* = –2.154; ITI = 1 s: t(19) = –9.229, p < 0.001 , Cohen' *d* = –2.064; ITI = 2 s: t(19) = –8.472, p < 0.001 , Cohen' *d* = –1.894; ITI = 4 s: t(19) = –7.362, p < 0.001 , Cohen' *d* = –1.646; ITI = 8 s: t(19) = –8.593, p < 0.001 , Cohen' *d* = –1.971] (Figure 2B left). In addition, we found that serial dependence decreased with increasing time between trials [F(4,72) = 3.963, p = 0.006, ω^2 = 0.074], but that there was an abrupt decrease from 0.1 to 1 s [Helmert contrast between 0.1 s against all other ITIs: t(18) = –3.571, p <0.001], which then did not decrease between 1 s to 8 s ITIs [all p>0.05 for subsequent contrasts] (Figure 2B left).

The effect of current speed was also present for all ITIs [ITI = 0.1 s: t(19) = 5.831, p < 0.001 , Cohen' *d* = 1.304; ITI = 1 s: t(19) = 6.154, p < 0.001 , Cohen' *d* = 1.376; ITI = 2 s: t(19) = 5.029, p < 0.001 , Cohen' *d* = 1.125; ITI = 4 s: t(19) = 5.695, p < 0.001 , Cohen' *d* = 1.273; ITI = 8 s: t(19) = 5.587, p < 0.001 , Cohen' *d* = 1.282], but was not significantly modulated by ITIs [F(4,72) = 0.469, p = 0.758, ω^2 = 0.000](Figure 2A and 2B [right panel]).

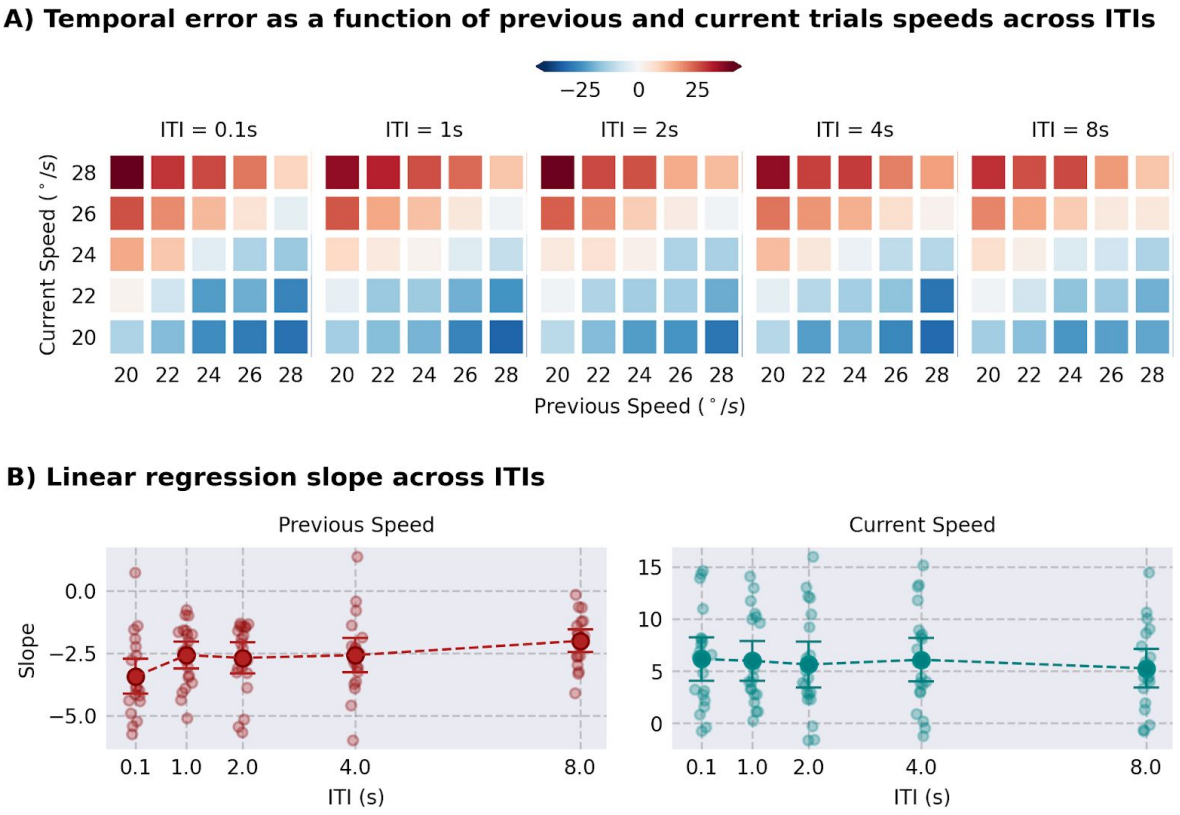

Figure 2. Serial dependence abruptly decreases after 0.1 s ITI but stabilizes afterward. (A) Attractive serial bias pattern of the temporal error as a function of current (vertical axis) and previous (horizontal axis) trial speed for all ITIs. The attractive serial bias is evidenced by temporal errors becoming more negative as previous trial speed increases. (B) Previous (left) and current (right) speed slopes of multiple linear regressions performed on individual participant data for all ITIs. The closer a value is to zero, the lower the regressor's contribution for explaining the temporal error. Large opaque dots represent average parameter estimates across participants, small and transparent dots represent individual data, and error bars represent 95% confidence interval over parameter estimates.

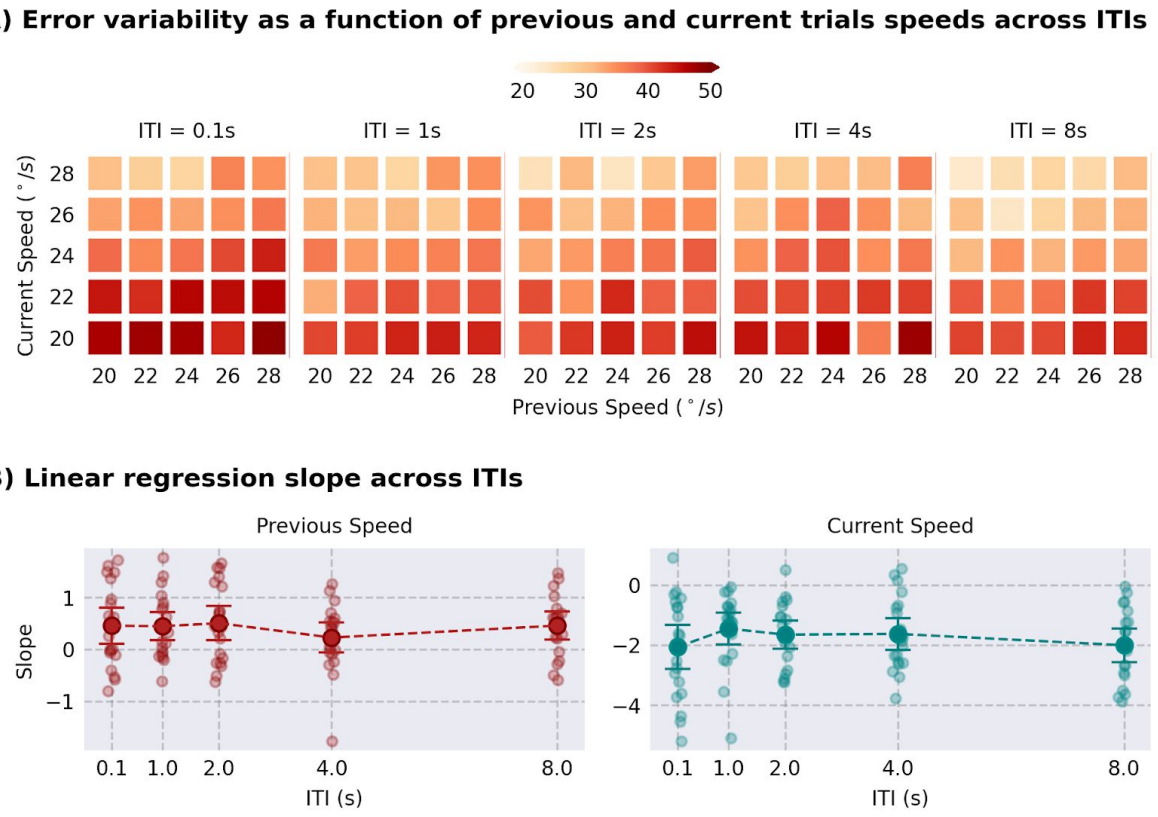

Figure 3. Error variability decreases with increased current target speed and is stable across ITIs. (A) Error variability as a function of current (vertical axis) and previous (horizontal axis) trial speed. There is a tendency for error variability to decrease as current trial speed increases (lighter colors from bottom to top) and increase as previous trial speed increases (darker colors from left to right). (B) Previous (left) and current (right) slopes of multiple linear regressions performed on individual participant error variability across ITIs. The closer a value is to zero, the lower the regressor's contribution for explaining the error variability. Large opaque dots represent average parameter estimates across participants, small and transparent dots represent individual data, and error bars represent 95% confidence interval over parameter estimates.

Given that the serial dependence effect was pronounced even after 8 s, we next asked how stable the memory trace that leads to serial dependence might be. According to theoretical models of short-term memory, long delays lead to higher response variability in short-term memory tasks due to memory trace diffusion (e.g., Schneegans and Bays 2017; Wimmer et al. 2014). If this is the case, we should see an increase in error variability as a function of ITI, allowing for the memory trace to diffuse more. To characterize how and whether variability increased as a function of the ITI, in a first step, we averaged across current and previous speeds and compared error variability over ITIs. Although we find evidence for a change in error variability over ITIs [$F(4, 72) = 5.902$, $p < 0.001$, $2= 0.029$], this was due to a decrease of error variability from 0.1 s compared to the rest of ITIs [Helmert contrast between 0.1 s against all other ITIs: $t(18) = 4$, $p <0.001$] and from 4 s compared to 8 s [Helmert contrast between 4 s and 8 s ITIs: $t(18) = 2.727$, $p = 0.008$] (Figure 1B [width of response distribution] and 3B).

To explore further how different ITIs modulate response variability, we performed a similar multiple regression on error variability as before, checking the dynamics of error variability across ITIs as a function of previous and current trial speed for all participants and ITIs. Our results show that error variability decreased as target speed increased in the current trial, as evidenced by a positive slope in the multiple regression for all ITIs [ITI = 0.1 s: $t(19) = -5.454$, $p < 0.001$, Cohen' $d = -1.220$; ITI = 1 s: $t(19) = -5.309$, $p < 0.001$, Cohen' $d = -1.187$; ITI = 2 s: $t(19) = -6.911$, $p < 0.001$, Cohen' $d = -1.545$; ITI = 4 s: $t(19) = -5.980$, $p < 0.001$, Cohen' $d = -1.337$; ITI = 8 s: $t(18) = -7.008$, $p < 0.001$, Cohen' $d = -1.608$]. Interestingly, there was also evidence for the opposite pattern of increased variability with faster targets on the previous trial for all, except 4 s, ITIs [ITI = 0.1 s: $t(19) = 2.557$, $p = 0.019$, Cohen' $d = 0.572$; ITI = 1 s: $t(19) = 3.188$, $p = 0.005$, Cohen' $d = 0.713$; ITI = 2 s: $t(19) = 2.997$, $p = 0.007$, Cohen' $d = 0.670$; ITI = 4 s: $t(19) = 1.524$, $p = 0.144$, Cohen' $d = 0.341$; ITI = 8 s: $t(18) = 3.328$, $p = 0.004$, Cohen' $d = 0.763$] (Figure 3A). In addition, we found no evidence for a change in slopes across ITIs [current, $F(4,72) = 0.885$, $p = 0.478$; previous, $F(4,72) = 0.274$, $p=0.894$] (Figure 3B).

Discussion

How long does the serial dependence effect persist between trials? Previous studies have shown that the impact of a trial on future behavior decreases with time (Fischer and Whitney, 2014; Kalm and Norris 2018; Fritsche et al. 2020). However, this decrease in serial dependence might be explained by an overwriting of the representation of previously experienced stimuli or responses (Matthey, Bays and Dayan 2015) or by decreasing their representation precision due to interference (Kalm and Norris 2018), not temporal decay alone. To understand the sole effect of the passage of time into the memory trace that underlie serial dependence, we had participants perform a coincident timing task at different inter-trial intervals in separate blocks. Our results indicate that the serial dependence effect fades abruptly after 1 s, but from there on remains pronounced for up to 8 s. In addition, participants' response variability seemed to be stable even after long intervals between trials, as evidenced by decreased response variability for longer ITIs. These results indicate that the memory trace that causes serial dependence is stable over at least 8 s.

Previous studies have also investigated the temporal properties of serial dependence, although with substantial differences in results. In Bliss et al. (2017), serial dependence was tested in a delayed reproduction task. The authors found that serial dependence sharply dropped from 3 to 6 s, and for intervals of 10 s, it turned into a repulsive bias. Conversely, and consistent with the results of the present work, Papadimitrou et al. (2015) showed that for 2 out of 3 monkeys, serial dependence was still high after 6 s of inter-trial interval, and that, similar to Bliss et al. (2017), one of the monkeys had serial dependence drop to zero after 6 s. A possible explanation for these contradictory results could be found in task differences. For instance, Bliss et al. (2017) used a delayed visual reproduction task, which might rely more heavily on visual memory than our coincident timing task and Papadimitrou et al. (2015) delayed saccade task. Another critical difference is that, in our task, participants were not required to keep target information actively in memory. In contrast, in Bliss et al. (2017) and Papadimitrou et al. (2015), keeping current trial information in working memory was required to perform the task.

It might be the case that implicitly and explicitly encoded memory traces have different temporal dynamics. There is evidence that working memory representation loses precision over longer delay intervals (Rademaker et al. 2018; Shin et al. 2017). These results have been interpreted as an increase in the uncertainty of memory representation over time (c.f. Pouget et al. 2013). Assuming memory information is combined with current sensory information weighted by their uncertainty to produce serial dependence (van Bergen and Jehee, 2019; Fritsche et al. 2020), a steady decrease of the memory trace precision should have led to a steady reduction of serial dependence in our experiment. However, we observed an abrupt decrease in response bias for short ITIs (1 s), from where it is still present for up to 8 s. In contrast to previous working memory studies (Rademaker et al. 2018; Shin et al. 2017), these results suggest that the memory trace's precision implicitly acquired in a coincident timing task does not vanish quickly in the absence of interference.

It has been suggested that serial dependence could result from a reminiscent and persistent activity encoding the previous target until the next trial (Papadimitrou et al. 2015). Assuming that such a proposal is true, current attractor network models used to explain short-term memory storage would predict that the longer the delay, the greater the variability in participants' responses (e.g., Wimmer et al. 2014). Given that we observed a decrease in

response variability for longer intervals between trials in our data, we speculate that this is due to a different information storage mode. Our findings are more consistent with recent models that have proposed that information could be stored in short-term memory through changes in synaptic weight (Mongilo et al. 2008). This mechanism would lead to a reduced diffusion of the memory content stored in the network (Itskov et al. 2011; Seeholzer et al. 2019), resulting in a temporally stable representation of memory contents. Our results also suggest that such an "activity-silent" mechanism should decrease the rate of decay of the memory trace. Although the idea of memory being stored in silent, hidden states has received electrophysiological evidence (Barbosa et al. 2020; Stokes 2015; Wolff et al. 2017; Rose et al. 2016), these results are far from unanimous (Papadimitrou et al. 2017), and it might be the case that the exact mechanism is task-dependent. However, it is essential to note that such models were exclusively created based on prefrontal neurons' activity, and the memory trace that leads to serial dependence might be stored elsewhere (c.f. Christophel et al. 2018).

An alternative explanation of why we do not find evidence for an increase of error variability over time is a possible limitation of our task when compared to purely perceptual tasks: the motor component of the coincident timing task might overshadow the noise due to diffusion. It could be argued that noise present in generating the motor command and processing sensory information would be unsusceptible to the effect of time. However, if these two sources of noise exist and are larger than the impact of noise due to diffusion of the memory trace, then the effects of diffusion would be difficult to show in our data. Future studies using tasks that control other sources of noise sources might help to resolve this issue.

Our results provide empirical evidence that serial dependence is sustained over time and the memory trace underlying it is stable, suggesting that future studies and models should look for explanations that accommodate such slow decay and stability. We speculate that the memory trace that leads to serial dependence is different from the one proposed to keep information through a persistent activity during the delay period in working memory tasks and favor the hypothesis that implicit short-term memory is stored through short-term synaptic plasticity mechanisms.

References

Barbosa, J., Stein, H., Martinez, R. L., Galan-Gadea, A., Li, S., Dalmau, J., … Compte, A. (2020). Interplay between persistent activity and activity-silent dynamics in the prefrontal cortex underlies serial biases in working memory. Nature Neuroscience, 23(8), 1016–1024. https://doi.org/10.1038/s41593-020-0644-4

Blake, R., Cepeda, N. J., & Hiris, E. (1997). Memory for Visual Motion. Journal of Experimental Psychology: Human Perception and Performance, 23(2), 353–369. https://doi.org/10.1037/0096-1523.23.2.353

Bliss, D. P., Sun, J. J., & D'Esposito, M. (2017). Serial dependence is absent at the time of perception but increases in visual working memory. Scientific Reports, 7(1), 1–13. https://doi.org/10.1038/s41598-017-15199-7

Brainard, D. H. (1997). The Psychophysics Toolbox. Spatial Vision, 10(4), 433–436. https://doi.org/10.1163/156856897X00357

Brooks, J. L. (2012). Counterbalancing for serial order carryover effects in experimental condition orders. Psychological Methods, 17(4), 600–614. https://doi.org/10.1037/a0029310

Christophel, T. B., Allefeld, C., Endisch, C., & Haynes, J. D. (2018). View-Independent Working Memory Representations of Artificial Shapes in Prefrontal and Posterior Regions of the Human Brain. Cerebral Cortex (New York, N.Y. : 1991), 28(6), 2146–2161. https://doi.org/10.1093/cercor/bhx119

Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. Nature Neuroscience. https://doi.org/10.1038/nn.3689

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, 1–32. https://doi.org/10.7554/eLife.55389

Itskov, V., Hansel, D., & Tsodyks, M. (2011). Short-term facilitation may stabilize parametric working memory trace. Frontiers in Computational Neuroscience, 5(October), 1–19. https://doi.org/10.3389/fncom.2011.00040

Kalm, K., & Norris, D. (2018). Visual recency bias is explained by a mixture model of internal representations. Journal of Vision, 18(7), 1–15. https://doi.org/10.1167/18.7.1

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in psychtoolbox-3. Perception, 36(14), 1-16.

Kwon, O.-S., & Knill, D. C. (2013). The brain uses adaptive internal models of scene statistics for sensorimotor estimation and planning. Proceedings of the National Academy of Sciences, 110(11), E1064–E1073. https://doi.org/10.1073/pnas.1214869110

Lewandowsky, S., Oberauer, K., & Brown, G. D. A. (2009). No temporal decay in verbal short-term memory. Trends in Cognitive Sciences, 13(3), 120–126. https://doi.org/10.1016/j.tics.2008.12.003

Leys, C., Ley, C., Klein, O., Bernard, P., & Licata, L. (2013). Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. Journal of Experimental Social Psychology, 49(4), 764–766. https://doi.org/10.1016/j.jesp.2013.03.013

Liberman, A., Fischer, J., & Whitney, D. (2014). Serial dependence in the perception of faces. Current Biology, 24(21), 2569–2574. https://doi.org/10.1016/j.cub.2014.09.025

Liberman, A., Zhang, K., & Whitney, D. (2016). Serial dependence promotes object stability during occlusion. Journal of Vision, 16(15), 1–10. https://doi.org/10.1167/16.15.16

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. Nature Neuroscience, 17(3), 347–356. https://doi.org/10.1038/nn.3655

Magnussen, S., & Greenlee, M. W. (1992). Retention and Disruption of Motion Information in Visual Short-Term Memory. Journal of Experimental Psychology: Learning, Memory, and Cognition, 18(1), 151–156. https://doi.org/10.1037/0278-7393.18.1.151

Makin, A. D. J., Poliakoff, E., Chen, J., & Stewart, A. J. (2008). The effect of previously viewed velocities on motion extrapolation. Vision Research, 48(18), 1884–1893. https://doi.org/10.1016/j.visres.2008.05.023

Matthey, L., Bays, P. M., & Dayan, P. (2015). A Probabilistic Palimpsest Model of Visual Short-term Memory. PLoS Computational Biology, 11(1), 1–33. https://doi.org/10.1371/journal.pcbi.1004003

Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic Theory of Working Memory. Science, 319(5869), 1543–1546. https://doi.org/10.1126/science.1150769

Papadimitriou, C., Ferdoash, A., & Snyder, L. H. (2015). Ghosts in the machine: Memory interference from the previous trial. Journal of Neurophysiology, 113(2), 567–577. https://doi.org/10.1152/jn.00402.2014

Pelli, D. G. (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies, Spatial Vision 10:437-442.

Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: Knowns and unknowns. Nature Neuroscience, 16(9), 1170–1178. https://doi.org/10.1038/nn.3495

Rademaker, R. L., Park, Y. E., Sack, A. T., & Tong, F. (2018). Evidence of gradual loss of precision for simple features and complex objects in visual working memory. Journal of Experimental Psychology: Human Perception and Performance, 44(6), 925–940. https://doi.org/10.1037/xhp0000491

Rose, N. S., LaRocque, J. J., Riggall, A. C., Gosseries, O., Starrett, M. J., Meyering, E. E., & Postle, B. R. (2016). Reactivation of latent working memories with transcranial magnetic stimulation. Science, 354(6316), 1136–1139. https://doi.org/10.1126/science.aah7011

Rousseeuw, P. J., & Croux, C. (1993). Alternatives to the Median Absolute Deviation. Journal of the American Statistical Association, 88(424), 1273–1283.

Schneegans, S., & Bays, P. M. (2017). Neural architecture for feature binding in visual working memory. Journal of Neuroscience, 37(14), 3913–3925. https://doi.org/10.1523/JNEUROSCI.3493-16.2017

Schneegans, S., & Bays, P. M. (2018). Drift in neural population activity causes working memory to deteriorate over time. Journal of Neuroscience, 38(21), 4859–4869. https://doi.org/10.1523/JNEUROSCI.3440-17.2018

Seeholzer, A., Deger, M., & Gerstner, W. (2018). Stability of working memory in continuous attractor networks under the control of short-term plasticity. In Stability of working memory in continuous attractor networks under the control of short-term plasticity. https://doi.org/10.1101/424515

Shin, H., Zou, Q., & Ma, W. J. (2017). The effects of delay duration on visual working memory for orientation. Journal of Vision, 17(14), 1–24. https://doi.org/10.1167/17.14.10

Stokes, M. G. (2015). “Activity-silent” working memory in prefrontal cortex: A dynamic coding framework. Trends in Cognitive Sciences, 19(7), 394–405. https://doi.org/10.1016/j.tics.2015.05.004

van Bergen, R. S., & Jehee, J. F. M. (2019). Probabilistic Representation in Human Visual Cortex Reflects Uncertainty in Serial Decisions. The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 39(41), 8164–8176. https://doi.org/10.1523/JNEUROSCI.3212-18.2019

Wimmer, K., Nykamp, D. Q., Constantinidis, C., & Compte, A. (2014). Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. Nature Neuroscience, 17(3), 431–439. https://doi.org/10.1038/nn.3645

Wolff, M. J., Jochim, J., Akyürek, E. G., Buschman, T. J., & Stokes, M. G. (2020). Drifting codes within a stable coding scheme for working memory. PLoS Biology, 18(3), 1–19. https://doi.org/10.1371/journal.pbio.3000625

Wolff, M. J., Jochim, J., Akyürek, E. G., & Stokes, M. G. (2017). Dynamic hidden states underlying working-memory-guided behavior. Nature Neuroscience, 20(6), 864–871. https://doi.org/10.1038/nn.4546