Implicit Learning of Temporal Behavior in Complex Dynamic Environments

Josh M. Salet

Department of Experimental Psychology

University of Groningen

Wouter Kruijne

Department of Experimental Psychology

University of Groningen

Hedderik van Rijn

Department of Experimental Psychology

University of Groningen

Author Note

This research was supported by a NWO 2017 awarded to Hedderik van Rijn.

Correspondence concerning this article should be addressed to Josh M. Salet, Department of Experimental Psychology, University of Groningen, Grote Kruisstraat 2/1, 9712 TS Groningen, the Netherlands. E-mail: j.m.salet@rug.nl

Abstract

To anticipate future events humans exploit predictive patterns in the environment. Such statistical learning is often outside of awareness. Timing research suggests that humans also adapt to temporal regularities. However, in these experiments, the intervals to be timed are isolated and explicitly cued, contrasting with everyday life where intervals are often unnoticed. In the present study, we found implicit adaptation to temporal regularities in an ecological setting. Ninety-eight participants played a game in which they responded to sudden-onset targets. While two targets appeared at random times, one target appeared every three seconds. In two experiments, we found adaptation to the regularity: Response times were lower, hit rates higher, and mouse cursor trajectories revealed anticipatory movements. Crucially, this was observed when participants were informed about the regularity, but also when they were unaware of it. Here, we for the first time, show implicit learning of temporal behavior in a complex environment.

Keywords: Implicit Memory, Prediction, Rhythm, Time Perception, Reaction Time

Implicit Learning of Temporal Behavior in Complex Dynamic Environments

While the brain is continuously confronted with a highly dynamic stream of sensory information, this sensory flood does contain regularities which the brain capitalizes upon to reduce processing load. For example, regularities of a stimulus sequence are learned (Brady & Oliva, 2008; Fiser & Aslin, 2002; Turk-Browne, Jungé, & Scholl, 2005), guide attention (Baker, Olson, & Behrmann, 2004; Olson & Chun, 2001; Yu & Zhao, 2015; Zhao, Al-Aidroos, & Turk-Browne, 2013), and bias which objects are encoded in working memory (Umemoto, Scolari, Vogel, & Awh, 2010). Oftentimes, this learning is completely implicit. Automatically extracting relevant information from the environment, often referred to as statistical learning, has been suggested to serve as a predictive coding strategy by means of which the brain generates an internal model of expected upcoming information to minimize the prediction error (Clark, 2013; Friston, 2005; Mehta, 2001).

It has been suggested that the brain, in a similar manner to ordinal regularities, may detect and utilize *temporal* regularities in the environment. A long tradition of classical conditioning research has shown that a consistent duration between events can be learned and guide behavior: Pavlov's dogs salivated in response to the bell, but only did so around the point in time of expected food delivery (Pavlov, 1927). Infants automatically track the interval of unattended auditory stimuli (Brannon, Roussel, Meck, & Woldorff, 2004), and adults optimize preparation at time points where an event is likely to occur (Coull & Nobre, 1998; Niemi & Näätänen, 1981; Woodrow, 1914). Los, Kruijne, and Meeter (2014, 2017) argued that such adaptation to regular intervals is driven by an implicit, automatic process. Recently, we have implemented this proposal in a computational framework and showed that such adaptive behavior can be accounted for by automatic Hebbian learning on the basis of past trials, without the use of explicit information about the duration of the intervals (Salet, Kruijne, Van Rijn, Los, and Meeter, manuscript submitted).

However, in typical timing experiments, the intervals to be timed are isolated and their on- and offsets are explicitly cued. This contrasts starkly with real world

environments in which regular intervals are embedded in a constant stream of both regular and irregular events. As a result it might be much more difficult to identify and utilize such regularities. While many statistical learning experiments involve regularities embedded in a stream of irregular events (e.g., Brady & Oliva, 2008; Fiser & Aslin, 2002; Olson & Chun, 2001; Turk-Browne et al., 2005), regularities in these experiments are formed by their order (e.g., stimulus 'B' always follows stimulus 'A'). It remains unclear whether temporal regularities (e.g., 'B' reappears after a predictable interval, yet could followed by either 'A' or 'C') can also be implicitly monitored and used to steer behavior.

In the present study, we investigated whether participants automatically extract and adapt to temporal regularities when they are 'hidden' in a continuous stream of stimuli and actions, as in everyday behavior. To this end, we developed a novel paradigm inspired by the arcade game 'Whac-A-Mole'. In our study, participants scored points by 'hitting' sudden-onset, briefly displayed targets (the 'moles') using the mouse. At two out of the three target positions, targets appeared at random times, but at one of the locations the target appeared consistently every three seconds. Assuming that participants adapt to this regularity, we expect better performance for the regular target than for the irregular targets. We analyzed performance expressed by response times, hit rates, and participants' cursor trajectories in epochs around target onsets. We split the experiment into an explicit and an implicit phase, where participants respectively were and were not informed about the regularity. This allowed us to investigate separately whether the temporal regularity can be used intentionally to optimize performance, and whether the temporal regularity might also be detected and utilized to drive behavior unintentionally.

Experiment 1

Method

Participants. Forty-eight first-year psychology students of the University of Groningen were recruited and received course credit for their participation. Sample size

was motivated on the basis of previous findings using Reaction Time (RT) as a measure of implicit learning (Zhao et al., 2013): A power analysis suggested a sample size of 39 participants to detect an effect of 'regularity' (Cohens's d = 0.46, power = 80%, alpha level = 0.05). The higher sample size was chosen to fully counterbalanced the location of the regular target across experimental conditions. The experimental procedure was approved by the Ethical Committee Psychology of the University of Groningen with ID 18018-P, and participants gave written informed consent prior to testing.

Stimuli and Task. The experiment was implemented using OpenSesame (Mathôt, Schreij, & Theeuwes, 2012) version 3.2.8 with the PsychoPy back-end (Peirce et al., 2019). Participants were seated in a dimly lit room at approximately 60 cm viewing distance from a 22" CRT monitor (100 Hz, Iiyama MA203DT) and were instructed to score as many points as possible by moving the mouse cursor towards sudden-onset targets. Three blue, unfilled circles (r = 1.5 cm) placed at the three corners of an equilateral triangle (sides: 24 cm) marked the positions of possible targets (Figure 1). A white fixation dot (r = 0.24 cm) was presented centrally. Target onsets were presented by filling the circles for 0.5 s. In each block, targets were presented 20 times at each of the three locations. One of the locations was the 'regular target location' at which a target appeared every three seconds. At the other two irregular locations, targets appeared at pseudorandomly generated intervals, while ensuring that the interval between two subsequent presentations in the same location was minimally 0.75 s (0.25 s between offset and new onset) and that there was always at least one irregular target between two regular target presentations.

Whenever participants moved the cursor within 3 cm of one of the target locations a white circle appeared around it indicating a hit at that location (Figure 1). Participants scored points if they could hit a target within 0.5 s after its onset. Each hit was rewarded with 5 points and accompanied by a positively valued sound of a dropping coin. If participants were too late to hit the target (> 0.25 s after its offset) or 'hit' a location that did not contain a target, three points were deducted as a penalty.

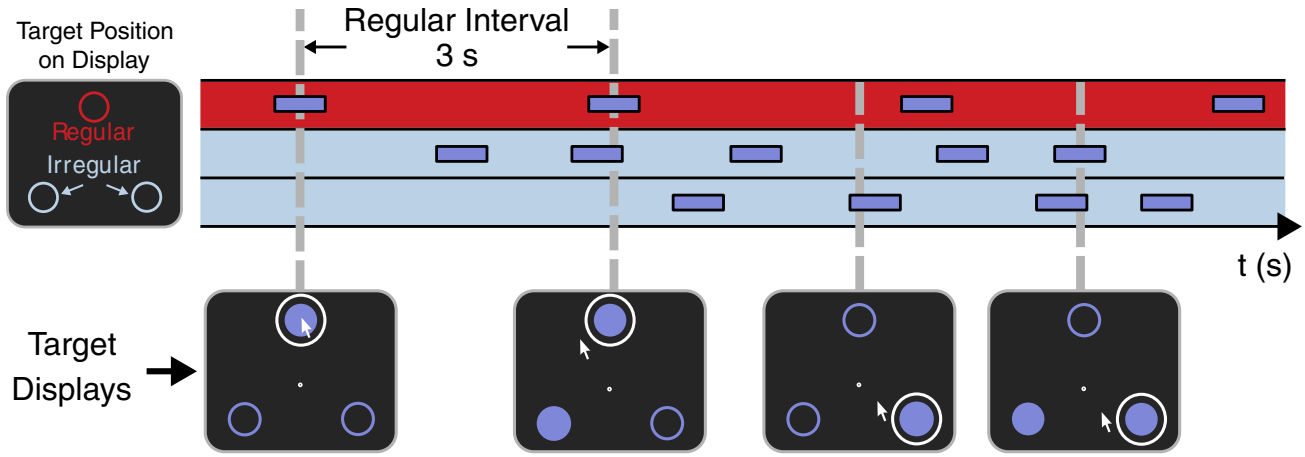

Figure 1. Segment of an exemplary stimulus sequence. The three colored rows in the chart (top figure) represent the three possible target locations: At one location (red row) the regular target appears every three seconds and in the other two locations (blue rows) targets appear at pseudorandomly determined moments. Example 'Target Displays' are depicted at four different time points. Note that colors and sizes for illustrative purposes do not match the displays of the experiment exactly. An example video of the experiment is available on OSF (https://osf.io/9fp43/).

Procedure. The experiment consisted of two phases of eight blocks, with 60 target presentations in total. A block lasted approximately one minute after which feedback was provided on the total score, the number of hits, misses per block, and the overall high score across all blocks completed so far. After the first phase, the temporal regularity was assigned to a new location. In each of the two phases, the regularity was either implicit or explicit, and the order was counterbalanced across participants. Participants who started with the implicit phase were initially not informed about any regularity, but were, before the second phase started, fully informed about the nature and location of the upcoming regular target. They were encouraged to make use of this information, but to not ignore the targets at the other two locations. For participants starting with the explicit phase these instructions were provided at the start of the experiment. Before the second phase

(the implicit phase) these participants were informed that the former regularity would no longer hold, without mentioning the upcoming new regular target location. The location of the temporal regularity in the first and second phase were counterbalanced between participants. Both phases started with two practice blocks. In the explicit phase, regular targets were colored red during practice, and at the start of each block the regular location was briefly colored red, up to the moment of first target appearance.

After the implicit phase, participants answered three questions to assess their awareness of the temporal regularity: (1) An open question whether they had noticed anything remarkable during the experiment, and (2) after being informed that there had been a temporal regularity, a binary question whether they had been aware of the regularity. Finally, (3) participants indicated (or guessed) the location of the temporal regularity.

Inclusion Criteria. Two participants indicated that they had been aware of the temporal regularity in the implicit phase, and could correctly identify its location. Their data were excluded from further analyses. We additionally excluded data from blocks where the hit rate of one of the three targets was less than 25%, as this suggested that participant might not have divided their attention across all targets. For six participants, more than three out of eight blocks were discarded in either phase, which led us to exclude their data entirely. Responses faster than 0.15 s (1.2% of cases) were excluded, as well as responses to targets that appeared consecutively at the same location (9.6% of cases), as due to the design such repetitions could only occur for irregular targets.[1]

Analyses. Analyses were conducted in R 3.6.1 (R Core Team, 2019). Analyses scripts as R-notebook are available on OSF (https://osf.io/9fp43/). The analyses aimed to determine whether participants adapted their behavior to the regular targets, and whether this was modulated by awareness of the regularity (explicit versus implicit). To this end, we analyzed performance expressed by RT and hit rates, and in addition we analyzed

[1] Analyses without excluding direct repetitions yielded the same conclusions.

participants' cursor trajectories in epochs around target onsets.

We analyzed RT by means of Linear Mixed Models (LMMs; Baayen, Davidson, & Bates, 2008; Bates, Mächler, Bolker, & Walker, 2015). We log-transformed RT to overcome the typical skew in RT distributions. Hit rates were analyzed with general (G)LMMs with a logistic link function. For both RT and hit rate models always included 'participant' as random intercept, and we relied on model comparisons by means of Bayes Factors estimated from BIC values (Wagenmakers, 2007) to assess whether the addition of predictors as fixed effects or random slopes was warranted.

The main interest of the study was whether regularity of targets (regular versus irregular), phase (explicit versus implicit), and their possible interaction affected performance. In search of the model that best captured the variation of the data we additionally assessed the effect of phase order (explicit - implicit versus implicit - explicit), response type on the preceding target (hit versus miss), and practice effects (progression in the experiment indexed by trial number). Here, we only report statistics on the effects of our main interest: 'regularity', 'phase', and their interaction. For completeness, we also report the outcome of Likelihood Ratio Tests for these comparisons. The Supplemental materials offers a complete report of the statistical procedure, including traditional repeated measures analyses of variance (ANOVA; including the standardized effect sizes) to facilitate comparison with existing studies.

To analyze cursor trajectories we selected data epochs in a time window from 0.5 s before to 1.0 s after each target onset, and calculated the distance of the cursor to the target every 0.01 s. We assessed the effects of 'regularity' and 'phase' at each timepoint by means of two-way repeated measures ANOVAs, and identify significant differences using nonparametric permutation tests (5000 iterations, alpha level 0.05; Maris & Oostenveld, 2007).

Results

Questionnaire. None of the participants showed any indication of awareness of the regularity when asked whether they noticed anything 'remarkable'. In the yes/no forced choice question, six out of 48 participants (12.5%) indicated to be aware of the temporal regularity. In the identification task, 18 participants (37.5%) correctly chose the location of the regular target out of the three options. Of the six participants that indicated being aware of the temporal regularity, only two participants (33%) chose the correct location. Data from these two participants were excluded from further analyses.

Reaction Times. Figure 2a depicts the mean RT split on 'regularity' and 'phase'. Visual inspection suggests that RT to the regular target was lower, and that this effect was more pronounced in the explicit phase than the implicit phase. Indeed, model comparisons revealed support for the inclusion of the factors 'regularity' (ΔBIC = 77.40, BF > 1000; $\chi^2(3) = 108.25$, $p < 0.001$), 'phase' (ΔBIC = 78.51, BF > 1000; $\chi^2(2) = 99.08$, $p < 0.001$), and their interaction (ΔBIC = 47.22, BF > 1000; $\chi^2(1) = 57.50$, $p < 0.001$). The interaction was driven by the finding that regular and irregular targets differed more in the explicit than the implicit phase. Critically, post hoc Tukey's HSD test revealed that RT to the regular target were still significantly lower in both the explicit ($p < 0.001$) and implicit phase ($p < 0.001$).

Hit Rate. Figure 2b depicts the average hit rate split on 'regularity' and 'phase'. The hit rates suggests a higher hit rate to the regular target, and that this effect was slightly more pronounced in the explicit phase. Despite this apparent trend, the GLMM revealed only support for the inclusion of the factor 'regularity' (ΔBIC = 5.30, BF = 14.17; $\chi^2(1) = 15.67$, $p < 0.001$), but not for 'phase' (ΔBIC = $-$ 10.35, 1/BF = 177.18; $\chi^2(1) = 0.02$, $p = 0.90$), and evidence for their interaction was inconclusive (ΔBIC = $-$ 0.02; 1/BF = 1.01), though note that likelihood ratio tests suggested that this interaction was significant ($\chi^2(2) = 20.72$, $p < 0.001$). Post hoc comparisons using Tukey's HSD test revealed that hit rate to the regular target was significantly higher in the explicit phase (p

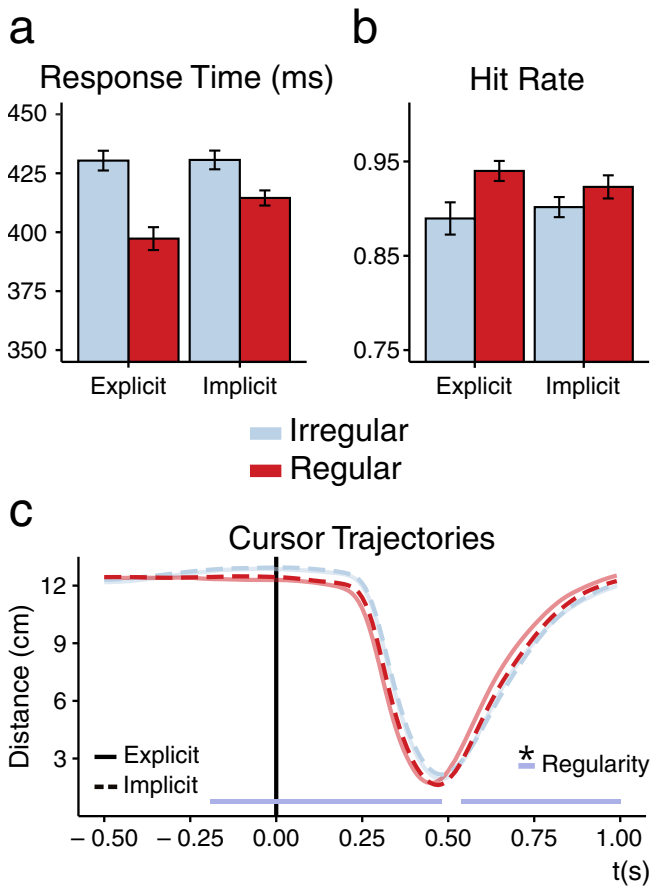

Figure 2. Results of Experiment 1. a) Average response times, separately depicted for explicit and implicit phases, split for regular versus irregular targets. b) Similar, but for hit rates. Error bars represent 95% Cousineau-Morey confidence intervals (Cousineau, 2005; Morey, 2008). c) Average mouse cursor distance to target in a − 0.5 to + 1.0 s window around target onset, for regular and irregular targets in the explicit and implicit phase. The horizontal bars at the bottom mark significant clusters for the main effect 'regularity'.

< 0.001) and also significant in the implicit phase, albeit with $p = 0.0496$.

Mouse cursor trajectories. Figure 2c displays the average cursor trajectories in a window of 0.5 s before to 1.0 s after target onset, split on 'regularity' and 'phase'. Two significant clusters were found indicating a main effect of 'regularity': One from 0.2 s before to 0.48 s after target onset ($p < 0.001$) and one from 0.54 s to 1.0 s after target onset ($p < 0.001$). Figure 2c illustrates that these clusters are indicative of earlier initiation

of movement towards regular targets compared to irregular targets, which effectively yields a leftwards shift of the distance curve.

Discussion

Only two out of 48 participants could accurately report on the temporal manipulation in the implicit phase. In the remaining participants, we did not find significant support for a higher hit rate to the regularity in the implicit phase. We nevertheless found that RT was lower and cursor movements were initiated earlier to the regular target, as if participants anticipated the upcoming event. Together, these results indicate that participants not only adapted to the regular target intentionally (explicit phase), but also implicitly without detecting its presence (implicit phase).

Experiment 2

The results from Experiment 1 revealed implicit adaptation to the 'hidden' temporal regularity. However, in-depth analyses of the pseudorandomly generated stimulus sequences revealed that the average inter-stimulus interval (ISI) differed between the regular and irregular targets. To rule out that any of the observed effects were driven by this higher-order dependency, we replicated Experiment 1 with a more stringent randomization procedure that took ISI into account.[2]

Method

Power analyses based on Experiment 1 indicated that a sample size of 32 participants would suffice to detect RT differences between regular and irregular targets (Cohens's d = 0.51, power = 80%, alpha level = 0.05). However, as Experiment 2 served as a stringent replication of Experiment 1, we continued data collection until data from 48 participants

[2] For Experiment 2 the hypotheses and analysis plan were pre-registered: https://osf.io/hdycr/. In the Appendix, we discuss this issue in more detail as well as the measures taken against it.

were included following the same 'Inclusion Criteria' as described for Experiment 1, maintaining full counterbalancing of experimental conditions. In total, 54 participants participated for course credits (i.e., six participants did not meet the inclusion criteria). The experimental procedure was approved by the Ethical Committee Psychology of the University of Groningen with ID 18233-S, and participants gave written informed consent prior to testing.

Results

Questionnaire. None of the participants showed any indication of awareness of the regularity when asked whether they noticed anything 'remarkable'. In the yes/no forced choice question, five out of 54 participants (10.4%) indicated to have been aware of the temporal regularity. In the identification task, 23 participants (47.9%) indicated the correct location of the regular target. The two participants that indicated being aware of the temporal regularity and correctly identified its location were excluded from further analyses.

Reaction Times. Figure 3a shows average RT pattern which revealed a similar pattern as for Experiment 1. Indeed, statistical analyses again supported the inclusion of the factors 'regularity' (ΔBIC = 281.79, BF > 1000; $\chi^2(3) = 312.24$, $p < 0.001$), 'phase' (ΔBIC = 242.47, BF > 1000; $\chi^2(2) = 263.44$, $p < 0.001$), and their interaction (ΔBIC = 222.02, BF > 1000; $\chi^2(1) = 232.50$, $p < 0.001$). This interaction was driven by the finding that regular and irregular targets differed more in the explicit than the implicit phase. Post hoc Tukey's HSD test revealed that RT to the regular target were still significantly lower in both the explicit ($p < 0.001$) and implicit phase ($p < 0.001$).

Hit Rate. Results regarding hit rate (Figure 3b) were again similar to Experiment 1. Indeed, the GLMM revealed support for the inclusion of the factor 'regularity' (ΔBIC = 58.39, BF > 1000; $\chi^2(3) = 90.12$, $p < 0.001$). Whereas in Experiment 1 we did not find statistical support for the apparent interaction between 'regularity' and 'phase', in

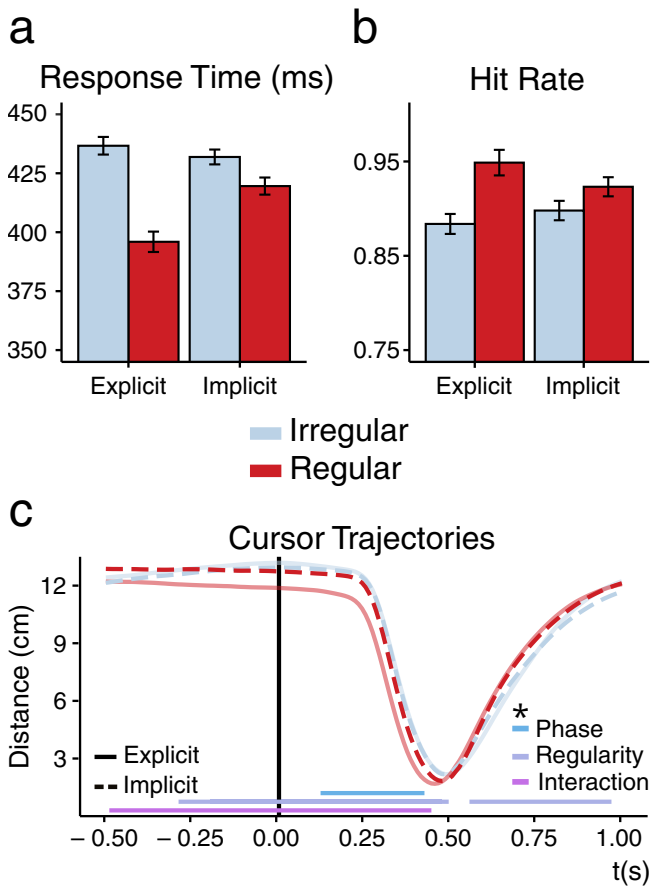

Figure 3. Results of Experiment 2. Reaction time (a), hit rate (b), and cursor trajectories (c) split by 'regularity' and 'phase'. Plotted as in Figure 2. Error bars represent 95% Cousineau-Morey confidence intervals (Cousineau, 2005; Morey, 2008).

Experiment 2 we find support for the additional inclusion of the factor 'phase' (ΔBIC = 31.90, BF > 1000; $\chi^2(2) = 52.05$, $p < 0.001$) and an interaction between these two effects (ΔBIC = 31.90, BF > 1000; $\chi^2(2) = 52.05$, $p < 0.001$). The interaction shows that the difference between the regular and irregular targets was more pronounced in the explicit phase. Critically, post hoc Tukey's HSD test revealed that hit rate was nevertheless significantly higher to the regular target in both the explicit ($p < 0.001$) and implicit phase ($p < 0.001$).

Mouse cursor trajectories. In Experiment 2, cursor trajectories were significantly modulated by 'regularity', 'phase' and their interaction (Figure 3c). Two

significant clusters for 'regularity' were found in the time window 0.29 s before to 0.49 s after target onset ($p < 0.001$) and 0.55 s to 0.96 s after target onset ($p < 0.001$), which is similar to our findings in Experiment 1. A main effect of 'phase' was marked by a significant cluster from 0.12 s to 0.42 s after target onset ($p = 0.02$), and a cluster for their interaction was found from 0.49 s before to 0.44 s after target onset ($p < 0.001$). This interaction was indicative of a stronger bias towards regular targets in the explicit phase, compared to the implicit phase (Figure 3c).

Conflict Situations. For Experiment 2 we pre-registered an additional hypothesis regarding hit rates (https://osf.io/hdycr/). In 'conflict situations' where two targets are simultaneously on screen, we expected that participants would be more inclined to hit the regular target. Our analyses, outlined in the Supplemental material, did not support this hypothesis, neither in the explicit or implicit phase. This might have been the result of a lack of power, as conflict situations selected for analysis were very rare (~ 5% of trials).

Discussion

Experiment 2 served as a stringent replication of Experiment 1, to ascertain that the observed implicit adaptation was exclusively caused by the temporal regularity. Indeed, we found this to be the case. All crucial effects of Experiment 1 were replicated: RT was found to be lower and the cursor movements to be initiated earlier to the regular targets compared to irregular targets. Whereas in Experiment 1 the support for higher hit rates to the regular targets was inconclusive, we found support for this hypothesis in Experiment 2.

Compared to Experiment 1 we additionally found an interaction between 'regularity' and 'phase' in the cursor trajectory analyses. This interaction cluster indicated a bias towards regular targets well before target onset in the explicit, but not in the implicit phase (Figure 3c). Despite the lack of such apparent anticipation in the implicit phase, the faster RT and higher hit rate for the regular target still implies future anticipation of target onset.

General Discussion

Do we make use of temporal regularities in our environment even if we are oblivious to them? To answer this question, we developed a paradigm in a game-like environment where participants scored points by 'hitting' sudden-onset targets. One of these targets appeared every three seconds and the others with random timing, which made the temporal regularity difficult to detect. Indeed, only ~ 4% of the participants across two experiments reported to have detected this regularity in blocks without prior instruction. Nevertheless, we found that unaware participants still displayed better performance to regular targets compared to irregular targets, as measured by RT, hit rates, and cursor movements. Thus, participants adapted to the regularity without detecting its presence, and thus without intentionally utilizing timing information.

Previous research on temporal preparation has similarly suggested that behavior adapts to temporal regularities. However, these experiments constitute warned RT- or detection tasks with a variable interval between the cue and the target. In these studies, timing is often characterized as 'implicit' since participants are not required to make an overt time estimation (see Coull & Nobre, 2008, for a taxonomy). However, as each trial is isolated and its onset is explicitly cued, the temporal information and its importance for performance are transparent. This contrasts with everyday behavior in which we base our timing judgments on complex, uncued dynamically unfolding events and are hardly, if ever, aware of the temporality of our behavior. The paradigm of the present study is in line with such an environment.

Interestingly, a recent study from our group (Damsma, Taatgen, de Jong, & van Rijn, 2019) has used a comparable design to study adaptation to implicit temporal regularities. In their experiments, a visual search display was preceded by sequences of stimuli, and one location or feature in this sequence was presented at regular intervals. It was hypothesized that attention would be attracted to the temporal regularity, but no evidence for such adaptation was found. A critical difference that could account for these discrepant results

is that in our paradigm the regularity was part of the task's goal, thereby immediately pairing the regular target with a motor response, and rewarding feedback upon a successful hit. Previous research has demonstrated that motor learning is facilitated when predictive cues are directly associated with a motor plan and immediate reward (Howard, Wolpert, & Franklin, 2015; Manohar et al., 2015; Wolpert & Flanagan, 2016). It is likely that in our design participants therefore attended the regularity more than in Damsma et al. (2019), which has previously been argued to be crucial for statistical learning (Baker et al., 2004; Richter & de Lange, 2019; Turk-Browne et al., 2005).

A potential neural correlate of adaptation to temporal regularities is rhythmic neural entrainment, a process where neural oscillations align to the temporal structure of a stimulus sequence (Large & Jones, 1999; Schroeder & Lakatos, 2009). Indeed, Olson and Chun (2001) have shown implicit adaptation to rhythmic stimulus presentations, with durations ranging between ~ 0.2 - 1.5 s. In our design, the 'rhythm' was formed by a much longer interval of three seconds. Research on motor control typically distinguishes between 'fast' rhythmic and 'slow' discrete periodic stimuli, separated by a boundary at ~ 1.5 s (Huys, Studenka, Rheaume, Zelaznik, & Jirsa, 2008; Miall & Ivry, 2004). The interval in our design clearly falls in the latter category. Although neural entrainment is typically studied using 'fast' rhythmic stimuli, Herbst and Obleser (2017, 2019) investigated such entrainment in a 'slow' discrete timing task. In these experiments, they found indications that slow neural oscillations implement temporal predictions by a 'phase-locking' mechanism that aligns the more excitable phases of neural oscillations with target's onset. Our paradigm would enable one to study such neural entrainment in cases where there is no explicitly cued interval to be timed.

Compared to traditional experiments on temporal preparation (e.g., Coull & Nobre, 1998; Los et al., 2017; Niemi & Näätänen, 1981; Woodrow, 1914), our paradigm is much more in line with a complex, dynamic real-world environment. However, temporal regularities in the real world are typically much more imprecise. For example, temporally

predictable events are not only formed by continuous rhythms and the timing of predicted events is often more imprecise than the exact three second interval used in our paradigm. An interesting open question is therefore whether the implicit adaptation we found is persistent and flexible enough to guide behavior in more imprecise scenarios.

Conclusion

The brain is thought to automatically capitalize upon temporal regularities in a complex, dynamic environment. Here, we developed a paradigm that adheres to an ecological setting in which temporal regularities were embedded in a stream of events. We found that participants, indeed, automatically exploited these regularities, without being able to report them afterwards. In this way, we provide, for the first time, direct evidence for the long-held belief that regularities in the time between the events we encounter in our daily life, despite being undetected, can be used to anticipate events.

References

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412. doi:10.1016/j.jml.2007.12.005

Baker, C. I., Olson, C. R., & Behrmann, M. (2004). Role of Attention and Perceptual Grouping in Visual Statistical Learning. *Psychological Science*, *15*(7), 460–466. doi:10.1111/j.0956-7976.2004.00702.x

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, *67*(1), 1–48. doi:10.18637/jss.v067.i01

Brady, T. F. & Oliva, A. (2008). Statistical Learning Using Real-World Scenes: Extracting Categorical Regularities Without Conscious Intent. *Psychological Science*, *19*(7), 678–685. doi:10.1111/j.1467-9280.2008.02142.x

Brannon, E. M., Roussel, L. W., Meck, W. H., & Woldorff, M. (2004). Timing in the baby brain. *Cognitive Brain Research*. Neuroimaging of Interval Timing, *21*(2), 227–233. doi:10.1016/j.cogbrainres.2004.04.007

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(3), 181–204. doi:10.1017/S0140525X12000477

Coull, J. T. & Nobre, A. C. (1998). Where and When to Pay Attention: The Neural Systems for Directing Attention to Spatial Locations and to Time Intervals as Revealed by Both PET and fMRI. *Journal of Neuroscience*, *18*(18), 7426–7435. doi:10.1523/JNEUROSCI.18-18-07426.1998

Coull, J. & Nobre, A. (2008). Dissociating explicit timing from temporal expectation with fMRI. *Current Opinion in Neurobiology*. Cognitive neuroscience, *18*(2), 137–144. doi:10.1016/j.conb.2008.07.011

Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorials in Quantitative Methods for Psychology*, *1*. doi:10.20982/tqmp.01.1.p042

Damsma, A., Taatgen, N., de Jong, R., & van Rijn, H. (2019). No evidence for an attentional bias towards implicit temporal regularities. *Attention, Perception, & Psychophysics*. doi:10.3758/s13414-019-01851-z

Fiser, J. & Aslin, R. N. (2002). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(3), 458–467. doi:10.1037/0278-7393.28.3.458

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *360*(1456), 815–836. doi:10.1098/rstb.2005.1622

Herbst, S. K. & Obleser, J. (2017). Implicit variations of temporal predictability: Shaping the neural oscillatory and behavioural response. *Neuropsychologia*, *101*, 141–152. doi:10.1016/j.neuropsychologia.2017.05.019

Herbst, S. K. & Obleser, J. (2019). Implicit temporal predictability enhances pitch discrimination sensitivity and biases the phase of delta oscillations in auditory cortex. *NeuroImage*, *203*, 116198. doi:10.1016/j.neuroimage.2019.116198

Howard, I. S., Wolpert, D. M., & Franklin, D. W. (2015). The Value of the Follow-Through Derives from Motor Learning Depending on Future Actions. *Current Biology*, *25*(3), 397–401. doi:10.1016/j.cub.2014.12.037

Huys, R., Studenka, B. E., Rheaume, N. L., Zelaznik, H. N., & Jirsa, V. K. (2008). Distinct Timing Mechanisms Produce Discrete and Continuous Movements. *PLOS Computational Biology*, *4*(4), e1000061. doi:10.1371/journal.pcbi.1000061

Large, E. W. & Jones, M. R. (1999). The dynamics of attending: How people track time-varying events. *Psychological Review*, *106*(1), 119–159. doi:10.1037/0033-295X.106.1.119

Los, S. A., Kruijne, W., & Meeter, M. (2014). Outlines of a multiple trace theory of temporal preparation. *Frontiers in Psychology*, *5*. doi:10.3389/fpsyg.2014.01058

Los, S. A., Kruijne, W., & Meeter, M. (2017). Hazard versus history: Temporal preparation is driven by past experience. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(1), 78–88. doi:10.1037/xhp0000279

Manohar, S. G., Chong, T. T.-J., Apps, M. A., Batla, A., Stamelou, M., Jarman, P. R., . . . Husain, M. (2015). Reward Pays the Cost of Noise Reduction in Motor and Cognitive Control. *Current Biology*, *25*(13), 1707–1716. doi:10.1016/j.cub.2015.05.038

Maris, E. & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 14.

Mathôt, S., Schreij, D., & Theeuwes, J. (2012). OpenSesame: An open-source, graphical experiment builder for the social sciences. *Behavior Research Methods*, *44*(2), 314–324. doi:10.3758/s13428-011-0168-7

Mehta, M. R. (2001). Neuronal Dynamics of Predictive Coding. *The Neuroscientist*, *7*(6), 490–495. doi:10.1177/107385840100700605

Miall, R. C. & Ivry, R. (2004). Moving to a different beat. *Nature Neuroscience*, *7*(10), 1025–1026. doi:10.1038/nn1004-1025

Morey, R. (2008). Confidence Intervals from Normalized Data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology*, *4*. doi:10.20982/tqmp.04.2.p061

Niemi, P. & Näätänen, R. (1981). Foreperiod and simple reaction time. *Psychological Bulletin*, *89*(1), 133–162. doi:10.1037/0033-2909.89.1.133

Olson, I. R. & Chun, M. M. (2001). Temporal contextual cuing of visual attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(5), 1299–1313. doi:10.1037/0278-7393.27.5.1299

Pavlov, I. P. (1927). *Conditioned Reflexes: An investigation of the Physiological Activity of the Cerebral Cortex.* London, United Kingdom: Oxford University Press.

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., ... Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203. doi:10.3758/s13428-018-01193-y

Richter, D. & de Lange, F. P. (2019). Statistical learning attenuates visual activity only for attended stimuli. *eLife*, *8*, e47869. doi:10.7554/eLife.47869

Salet, J., Kruijne, W., Van Rijn, H., Los, S. A., & Meeter, M. (2019). fMTP: A Unifying Computational Framework of Temporal Preparation Across Time Scales. *Manuscript submitted for publication*.

Schroeder, C. E. & Lakatos, P. (2009). Low-frequency neuronal oscillations as instruments of sensory selection. *Trends in Neurosciences*, *32*(1), 9–18. doi:10.1016/j.tins.2008.09.012

Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The Automaticity of Visual Statistical Learning. *Journal of Experimental Psychology: General*, *134*(4), 552–564. doi:10.1037/0096-3445.134.4.552

Umemoto, A., Scolari, M., Vogel, E. K., & Awh, E. (2010). Statistical Learning Induces Discrete Shifts in the Allocation of Working Memory Resources. *Journal of experimental psychology. Human perception and performance*, *36*(6), 1419–1429. doi:10.1037/a0019324

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, *14*(5), 779–804. doi:10.3758/BF03194105

Wolpert, D. M. & Flanagan, J. R. (2016). Computations underlying sensorimotor learning. *Current Opinion in Neurobiology*. Neurobiology of cognitive behavior, *37*, 7–11. doi:10.1016/j.conb.2015.12.003

Woodrow, H. (1914). The measurement of attention. *The Psychological Monographs*, *17*(5), 1–158. doi:10.1037/h0093087

Yu, R. Q. & Zhao, J. (2015). The persistence of the attentional bias to regularities in a changing environment. *Attention, Perception, & Psychophysics*, *77*(7), 2217–2228. doi:10.3758/s13414-015-0930-5

Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention Is Spontaneously Biased Toward Regularities. *Psychological science*, *24*(5), 667–677. doi:10.1177/0956797612460407

Appendix

Experiment 2 was a direct, more stringent replication of Experiment 1, and was conducted because we found that the inter-stimulus interval (ISI) in Experiment 1 differed for regular versus irregular targets. To rule out that any observed effects could be driven by this higher-order dependency, we replicated Experiment 1 with a more stringent randomization regime for the stimulus sequences.

In order to minimize ISI differences in conditions while otherwise maintaining the randomization procedure as unconstrained as in Experiment 1, we simulated experiments in a random fashion and assessed the ISI distributions for regular and irregular targets by means of the Kullback-Leibler (KL) divergence. That is, for each participant, we simulated a full experiment existing of eight unique blocks for both the explicit and implicit phase. For each simulated experiment, we retrieved the ISI's split on 'regularity'. Following, we estimated the ISI distributions for these values using a Gaussian kernel density estimation with a fixed bandwidth of 0.2. Finally, we compared the estimated ISI distributions by computing the KL-divergence as a measure of how one probability distribution (regular) diverges from a second probability distribution (irregular). A randomly generated sequence was only accepted as an experimental sequence if the KL-divergence between the two ISI

Table A1

Descriptive summary of the ISI distributions of Experiment 1 (Exp1) and Experiment 2 (Exp2): mean (μ), median ($\mu_{1/2}$), standard deviation (σ), and minimum and maximum ISI (min, max), all displayed in seconds.

	Exp1				Exp2			
	μ	$\mu_{1/2}$	σ	min, max	μ	$\mu_{1/2}$	σ	min, max
Regular	0.998	0.850	0.706	0.059, 2.75	0.983	0.870	0.673	0.095, 2.75
Irregular	0.964	0.870	0.648	0.071, 2.75	0.978	0.870	0.669	0.094, 2.75

distributions was lower than 0.001 (for comparison, in Experiment 1 the KL-divergence was ~ 0.1). Table A1 provides a descriptive summary of the ISI distribution used in both experiments. The code used to generate these sequences is included in the OSF-repository (https://osf.io/9fp43/).