Causal Learning With Delays Up to 21 Hours

Yiwen Zhang (yiwenzhang@pitt.edu)

Benjamin M. Rottman (rottman@pitt.edu)
Department of Psychology, University of Pittsburgh
Pittsburgh, PA 15260 USA

Abstract

Delays between causes and effects are commonly found in cause-effect relationships in real life. However, previous studies have only investigated delays on the order of seconds. In the current study we tested whether people can learn a cause-effect relation with hour long delays. The delays between the cause and effect were either 0, 3, 9, or 21 hours, and the study lasted 16 days. Surprisingly, we found that participants were able to learn the causal relation about equally as well in all four conditions. These findings demonstrate a remarkable ability to accurately learn causal relations in a realistic timeframe that has never been tested before.

Keywords: causal learning; delay; preparedness of learning

Introduction

An important question in the fields of causal, reinforcement, and associative learning, is the impact of the delay between a cause and effect (or cue and outcome) on learning.

Temporal delays have been studied in animal conditioning and reinforcement learning for decades. Initially it was believed that learning is worse with longer intervals between the conditioned stimulus and unconditioned stimulus (Renner,1964). Subsequent research on taste aversion has found that learning with delays ranging from one hour to 24 hours is possible (Logue, 1979), but this phenomenon of learning over a long delay, often called 'preparedness of learning' is believed to be an evolutionarily adapted exception for food-related conditioned stimuli and certain fearful stimuli and phobias (Dunlap & Stephens, 2014). Computational models of associative learning that have attempted to model the influence of delay on learning ('trace learning' paradigms) have primarily attempted to capture how longer delays, even on the order of seconds, produce considerably less learning. For example, associative learning theories that are more biologically inspired assume a short window of associability (e.g., Gluck & Thompson, 1987; Grossberg & Schmajuk, 1989).

Whereas trial-by-trial models of learning (e.g., Rescorla & Wagner, 1972) are not sensitive to intra-trial temporal dynamics, subsequent reinforcement learning models such as Sutton and Barto's Temporal Difference model (1990) was specifically designed to capture this finding by using eligibility traces. However, there are still open debates in the theorizing of the role of delay. For example, Gallistel and Gibbon's (2000) timing model proposes that the effect of delay is proportional to the intertrial interval; if both are increased in proportion then there is no impact of delay. But most importantly, delay has played a central role in the field of animal learning and Reinforcement Learning.

Within the field of human causal learning, there have also been debates about the role of delay. Initially it was believed that humans have difficulty learning cause-effect relations with longer delays. Shanks, Pearson and Dickinson (1989) investigated the role of temporal contiguity between a cause and effect with delays from 0 to 16 seconds and found that participants' judgements of causal efficacy were significantly reduced with delays longer than 4 seconds.

However, subsequent studies showed that causal learning is mediated by temporal assumptions (Buehner & McGregor, 2006; Buehner, 2005; Hagmayer & Waldmann, 2002). Buehner and colleagues argued that Shanks et al.'s (1989) results were due to learners having an expectation of an immediate succession of causes and effects. Buehner and McGregor (2006) had people learn the relation between releasing a ball into a chute and a light coming on triggered by the ball reaching the end of a chute. They had two conditions, one which had a steeper chute, and another with a more gradual chute, so participants had different expectations about how long it would take for the ball to reach the end of the chute. When the chute was gradual, and the delay between the insertion of a ball and the light was longer (their expectations matched the delay), participants gave stronger causal ratings that inserting the ball made the light come on than if the delay was so short that it violated their expectations. Other research investigated the roles of the variability of delay and number of intervening events as opposed to delay per se (Lagnado & Speekenbrink, 2010), and how people use distributions of delays for inferring causal structures among multiple variables (Bramley, Gerstenberg, Mayrhofer, & Lagnado, 2018).

Longer Delays and Current Study

Despite all this important empirical and theoretical work on delay, an important open question is how delays impact human learning in real-life situations. Almost all the prior research, with the exception of the preparadness of learning research with animals, has focused on delays on the order of seconds. However, many real-life causal events occur with delays of several minutes, hours, or days. The primary goal of this research is to investigate how well people are able to learn cause-effect relations with delays on the order of hours.

Recently we have begun to study how well people can learn cause-effect relations from data presented one trial per day for

a series of days. The reason is that one way in which standard causal-learning paradigms are artificial is that all the trials are presented in quick succession, whereas in the real world (e.g., learning if a medicine is working, or what factors influence sleep), the experiences are spaced out over much longer periods of time. We have found that people can learn the relation between a single cause and a single effect about as well when spaced out one trial per day as when presented rapidly within a few minutes, and the extent of illusory correlation is also about the same across short and long timeframe presentations (Willett & Rottman, in press).

In a subsequent study testing how well people can learn about two causes and one effect, one cause had a strong unambiguous influence on the effect, but the other cause had a weaker influence, and assessing the influence required 'controlling' for the influence of the first cause. We found that people could learn about the stronger cause, but in the long timeframe people were not able to control for one cause when assessing the influence of the other cause (Willett & Rottman, 2020). Though people were able to control for second causes in the rapid presentation format, there also was not a significant difference between the two conditions.

Critically, in both of these experiments, the cause and effect were presented at the same time, so even though they are more realistic in one sense that the trials are more spaced out, they are still artificial in that there is zero delay between the cause and effect. The findings that participants are able to learn single cause-effect relations quite well, and are able to learn about two causes to some degree, possibly represent an overly optimistic picture of real-world causal learning.

This study aimed to assess the effect of temporal delays on causal learning over 16 days. We adopted a trial-by-trial learning paradigm and spaced it out to one trial per day. We manipulated the intervals between the cause and the effect within a trial, ranging from a few seconds to roughly 21 hours, to investigate whether long term causal learning is impeded with delay and the extent to which people can accurately learn cause-effect relations with long delays.

Methods

Participants

202 participants completed the study. 76 participants were recruited within the Pittsburgh community (mainly undergraduate students) and attended an in-person lab session on the first day of study. Due to the COVID-19 pandemic, the rest of participants were recruited through social media (e.g. Facebook) and attended a remote lab session over Zoom on the first day of study. Participants who successfully completed the entire study were paid $40. The final analyses included 200 participants, excluding 1 participant who explicitly reported they wrote down data during the study and 1 participant due to a programming error.

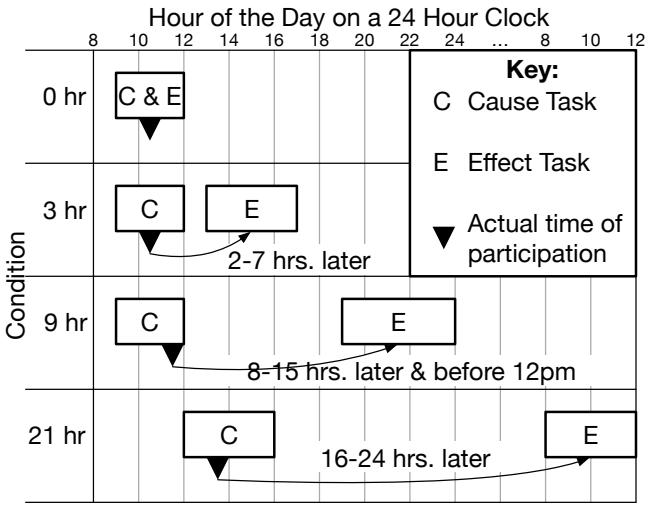

Figure 1: The time windows for participation.

Design

The study had employed a 2x4 between-subject design. There were two types of learning datasets (positive correlation vs. negative correlation) and four temporal delay conditions of roughly 0, 3, 9, or 21 hours between the cause and effect.

Dataset Type The positive correlation dataset used the following data: the cause and effect were both present 6 times (A cell), both absent 6 times (D cell), the cause was present and the effect was absent 2 times (B cell), and the cause was absent and the effect present 2 times (C cell). For the negative dataset, the cell frequencies were reversed [A=2, B=6, C=6, D=2]. According to the ΔP rule (Allan, 1980), the causal strengths were .5 and -.5 for the two datasets respectively. In order to ensure that the contingency kept the same for the first 8 days and the latter 8 days, we divided the whole dataset into two identical sets ([A=3, B=1, C=1, D=3] for the positive; [A=1, B=3, C=3, D=1] for the negative) and randomly ordered 8 trials across each half.

Temporal Delays We manipulated the temporal delays within each trial. Participants observed 16 trials of events and each trial contained a cause task in which participants learned whether the cause is present or absent, and an effect task in which participants learned whether the effect is present or not. In the 0-delay condition, participants did the cause and the effect task back to back each day. In the 3-hour delay condition, participants did the cause task in the morning and the effect task in the afternoon around 3 hours (min = 2, max = 7) later than the morning task. In the 9-hour delay condition, participants did the cause task in the morning and the effect task in the evening around 9 hours (min = 8, max = 15) later. In the 21-hour delay condition, participants did the cause task in the afternoon and the effect task next morning roughly 21 hours (min = 16, max = 24) later. See Figure 1 for a visualization of these windows of time to participate.

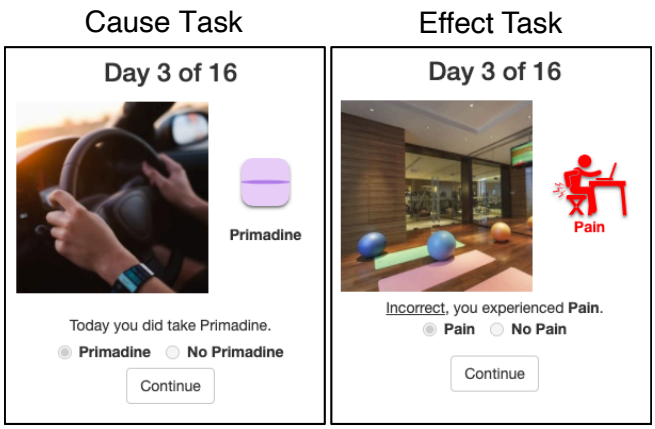

Figure 2: Screenshots of the cause task and the effect task. In the cause task, the participant has already verified whether the cause is present or not based on the icon. In the effect task, the participant predicted that the effect would be absent (no pain), and is now receiving feedback (that they experienced pain) and are verifying the pain.

The study was run automatically through a custom built website using the psychcloud.org framework. This website sent automated text message reminders, and allowed participants to login only at the allocated times. When participants were supposed to do the task, they were sent a text message, and if they did not they received hourly reminders.

If a participant did not do one of the tasks, either cause or effect task within the window of time that they were allotted in a given day, they were not allowed to participate for the the rest of the day, and they received the same trial the subsequent day. If participants missed more than 4 days, they could not continue to participate in the study. In total 6 participants were dropped from the study due to missing more than 4 days.

Procedures

The entire study which was conducted on participants' mobile phones, contained one practice task which happened in the lab (or over Zoom) on Day 0, one 16-day learning task and one final judgement task which happened on Day 17.

On the first day (Day 0), participants did a practice task in the lab (or over Zoom) which took 20 minutes. The task contained a four-trials learning session and a testing session afterwards. In the learning session, the trials were shown back-to-back. Participants finished one trial by pressing a button and then they moved to the next one. The goal for the practice task is to become familiar with the experiment procedure.

The long-term task began on Day 1. At the beginning, the participants read about a cover story that they are taking a new medication during the next 16 days and they need to figure out whether the medicine improves or worsens or has no influence on their back pain. The entire learning session contained 16 trials, one trial per day. Each day participants conducted two tasks, a cause task and an effect task (see Figure 2).

In the cause task, participants first saw an image of a scene.[1] After they clicked the 'Continue' button, they were shown an icon and a text of whether the cause is present or absent that day. Participants then verified the status of the cause by clicking a button. Only after they responded correctly would a 'Continue' button appear allowing them to continue. Finally, they were asked to "tell a story that links both pictures together" by typing down their story.[2]

The effect task followed a similar procedure, except that before seeing whether the effect was present or absent, participants were asked to predict the status of the effect (whether or not they have back pain). After they submitted their prediction, they received text feedback of their prediction and an icon showing whether they have back pain or not.

On Day 17, the day after the 16-day learning task, they did a 15-minute final judgment task. The task consisted of two parts. In the causal strength part, we measured causal strength, future prediction strength and frequency strength. In the memory part, we measured participants' recognition memory of the contextual images and episodic memory of causes and effects within each day, but we are not going to analyze the data from the memory part in this article.

Measures

We used five different measures of participants beliefs about the strength of the relation between the cause and the effect. All the measures were scaled in a range of [-1,1] for analysis.

Causal Strength Participants made a standard "causal strength" judgment by answering whether the medicine "worsens, improves, or has no influence on pain" (on a scale of -10 to $+10$, -10 = strongly worsens, 0 = no influence, and $+10$ = strongly improves). This question was asked both in the middle of the learning session (after Trial 8) and in the testing session (after Trial 16).

Future Prediction Strength Participants were asked about the probability of having pain given that they did or did not take the medicine ("Imagine that 'tomorrow' (Day 17) you take/do not take" medicine). On a scale of 0 to 100%, what do you think is the likelihood that you would experience pain?" The future prediction strength was derived by subtracting participants' responses of when they do not take the medicine from when they do take the medicine - similar to the ΔP rule.

[1] These scenes were included to help provide some context and for other analyses involving episodic learning that are not presented in the current paper.

[2] We wanted to ensure that participants were paying attention and encoding the stimuli, not just clicking through the task, given that this task was embedded in their daily lives and could happen while they were doing other things. We were also concerned that if learning in all conditions was at floor it could be explained merely due to a lack of processing. One potential concern is that this task may have led to increased salience, perhaps leading to an overly optimistic picture of learning with delays. Though possible, these stories were still quite short and likely took 10-20 seconds to write. In comparison, many real-world events that people care to learn about are likely to be much more salient and important in one's life (e.g., pain, sleep) leading to deeper processing than in the current task.

Future Use Participants answered whether they believe they should continue to use the medicine on a scale of -10 to $+10$, -10 = definitely no, 0 = unsure, $+10$ = definitely yes.

Frequency Strength We asked about their memories of the frequencies of A, B, C, and D cells (e.g., "Of the 16 days in the study, how many days did you see a picture in which you did take the medicine and did experience pain"). We calculated frequency strength by calculating $p(\text{effect} \mid \text{cause}) - p(\text{effect} \mid \neg\text{cause})$ from participants memories of A, B, C, and D cells. We excluded one participant from data analysis due to this participant's frequency strength being incalculable due to a division by zero problem, which can happen if some pairs of cells are judged as zero.

Trial by Trial Prediction Strength We computed "prediction strength" from participants' predictions about the presence or absence of the effect from Trial 9 to Trial 16 using a similar equation as calculating frequency strength. We only use the last 8 trials because we assume that participants have experienced some data so that they could give a relatively reliable prediction.

Results

The analysis follows our pre-registered plan (https://osf.io/8tcvq). For ease of interpretation, we inverse coded the judgments for the negative datasets so that they are positive.

Figure 3 shows summaries of the six measures. As can easily be seen in the figure, participants' judgments were above zero for all six measures and for all 4 conditions, which provides evidence that participants were able to learn the contingency between the cause and effect in every condition. All 24 t-tests against zero were significant (see Figure 3).

We conducted two analyses and provided both p-values and Bayes Factors (BFs, Kruschke, 2014). First, we conducted an ANOVA for each dependent measure. If learning becomes weaker with longer delays, there should be a main effect of delay.[3] These ANOVAs also included a main effect of dataset (positive vs. recoded negative) and an interaction.

Table 1 presents the ANOVA results for the six measures. We first discuss the main effect of delay. In four of measures (causal strength after Trial 8, and after Trial 16, future use strength, and future prediction strength), there was no significant effect of delay, and the BFs were in the range of .03-.06, meaning that the evidence is roughly 20 to 1 in favor of the null hypothesis of no influence of delay. In the frequency strength measure, the BF was weaker, about 2 to 1 in favor of a null effect. The only measure that found an effect of delay was the trial-by-trial prediction strength measure. As can be seen in Figure 3, participants' predictions were stronger for the 0 and 3 hour conditions than for the 9 and 21 hour conditions, suggesting that learning was somewhat better with the shorter delays. The p-value was .006, though the BF was not especially strong; 5.126 to 1 in favor of the alternate. We compared the trial by trial prediction strength measure across the 0 vs. 3, 3 vs. 9, and 9 vs. 21-hour delay conditions. Out of these three comparisons, the only significant difference was 3 vs. 9 hours, $F(1,95) = 7.048$, $p = 0.009$.

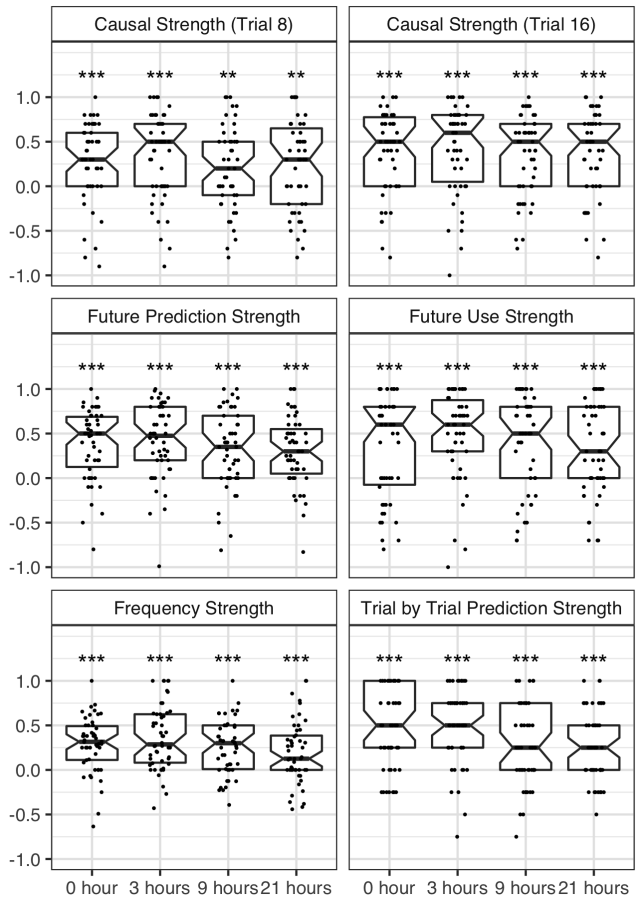

Figure 3: Judgments in four delay conditions and t-test results. *p-values against 0*: *<.05,**<.01, ***<.001

Four out of the six ANOVAs also found very strong main effects of dataset. This is because participants tended to give somewhat stronger judgments in the negative condition than the positive condition. This could have been due to a bias to think that the medicine is effective - that the presence of the medicine would help prevent the back pain. Though the cover story explicitly said that the medicine could improve or worsen the back pain, this bias is understandable, and is not of primary importance to the study. There were no significant interactions between delay and dataset, with most of the BFs roughly in the range of 10 to 1 in favor of the null.

[3]There are two highly related ways to conduct this analysis. One way involves testing for an interaction between dataset and delay; if participants have more difficulty learning the cause-effect relations then their judgments for the positive and negative datasets would get closer together over longer delays. Here as preregistered, we took a simpler approach of inverse coding the judgments for the negative datasets so that they are positive and then testing for a main effect of delay. These two approaches are very similar mathematically, only here we are primarily interested in a main effect of delay whereas in the other version we would primarily be interested in the interaction.

Table 1: ANOVA results for the six measures.

	F	*p*	BF
Causal Strength (After Trial 8)			
Delay	0.62	0.603	0.05
Dataset	3.98	0.048	1.03
Delay:Dataset	1.13	0.340	0.19
Causal Strength (After Trial 16)			
Delay	0.18	0.908	0.03
Dataset	14.17	<0.001	128.6
Delay:Dataset	0.27	0.846	0.08
Future Prediction Strength			
Delay	0.76	0.518	0.06
Dataset	14.75	<0.001	147.3
Delay:Dataset	0.57	0.635	0.10
Future Use Strength			
Delay	0.85	0.467	0.06
Dataset	12.57	0.001	55.71
Delay:Dataset	0.36	0.783	0.08
Frequency Strength			
Delay	2.24	0.085	0.41
Dataset	1.65	0.200	0.34
Delay:Dataset	0.46	0.709	0.09
Trial by Trial Prediction Strength			
Delay	4.28	0.006	5.13
Dataset	12.12	0.001	45.9
Delay:Dataset	0.78	0.510	0.13

Table 2: Regression results for the six measures.

	β	*p*	BF
Causal Strength (After Trial 8)			
Interval	-0.003	0.417	0.37
Dataset	-0.101	0.291	0.46
Interval:Dataset	-0.004	0.647	0.30
Causal Strength (After Trial 16)			
Interval	-0.002	0.616	0.28
Dataset	-0.241	0.010	5.52
Interval:Dataset	0.0002	0.981	0.25
Future Prediction Strength			
Interval	-0.004	0.237	0.46
Dataset	-0.261	0.001	27.3
Interval:Dataset	0.005	0.439	0.32
Future Use Strength			
Interval	-0.003	0.525	0.30
Dataset	-0.293	0.006	8.61
Interval:Dataset	0.004	0.661	0.27
Frequency Strength			
Interval	-0.006	0.035	2.07
Dataset	-0.117	0.087	1.04
Interval:Dataset	0.007	0.254	0.48
Trial by Trial Prediction Strength			
Interval	-0.011	0.003	13.6
Dataset	-0.277	0.001	30.2
Interval:Dataset	0.009	0.229	0.45

Second, we conducted linear regressions to test for an effect of the actual delay that a participant experienced given that participants in the same condition could have experienced somewhat longer or shorter delays based on exactly when they choose to complete the task. The regressions used the actual delay intervals, rather than the 4 delay conditions. The actual time interval for an individual participant was the average time interval between the cause task and the effect task over the 16 days. The delay is coded in terms of hours for interpreting the regression coefficient.

Table 2 presents the results of the regressions, which were similar to the ANOVA results. Again, we only found a main effect of delay in the trial by trial prediction strength measure, with a somewhat stronger BF than the ANOVA; 13.6 to 1. There was a marginal effect of delay for the frequency measure, but the BF was very weak. The main effects of dataset and the interaction mirror the ANOVA results.

Discussion

This is the first study to investigate human causal learning in which the trials were spaced out over time (once per day) and there were considerable delays between the cause and the effect ranging from 0 to 21 hours. These two features were designed to make the task mimic a real-world causal learning situation in which an individual is learning about causes and effects over longer periods of time (e.g., what factors improve or worsen one's sleep or one's mood). Surprisingly, this study showed that people are capable of learning cause-effect relationships with delays up to 21 hours, and for the most part causal learning was not affected by the length of delay, ranging from 0 to 21 hours.

Exceptions to the General Finding We did find some exceptions. Most notably, the length of delay affected the trial by trial prediction strength. This judgment is somewhat different from the rest. The trial by trial prediction strength was calculated from participants' predictions of the effect after knowing that the cause was present or absent on each day during the latter 8 days of learning. At a conceptual level, prediction-based reinforcement learning models (e.g., Rescorla & Wagner, 1972) assume that learners spontaneously make predictions of outcomes, and that these predictions are required for learning. Because predictions were made during learning in this study, they can be viewed as interim measures of learning, but they should be tied to the other measures. Though the predictions should be related to the final measures of learning, the predictions are the most implicit or at least farthest removed from all the other measures. For example, we did not remind participants of the cause when they predicted the effect, and retrieving the cause for making predictions may require more effort in the longer delay conditions so that participants may have been less likely

to make predictions of the effect based on the cause they saw earlier. All of these reasons could have contributed to effects of delay appearing primarily for this measure.

At the same time, there are some other findings that do not entirely fit with this interpretation. Most importantly, there was not an effect of delay in the causal strength measure after Trial 8, suggesting that learning was not considerably impaired by the delay even earlier on. Furthermore, the proposal that participants weren't thinking about the cause when predicting the effect in the longer delay conditions is problematic from a reinforcement learning perspective. Prediction is necessary in reinforcement learning, and if participants aren't predicting the effect from the cause reinforcement learning models would not be able to say how participants had accurate judgments by the end of learning. Furthermore, the future predictive strength measure also did not find an effect of delay. We asked the future predictive strength measure in order to have another measure similar to the trial by trial prediction strength measure, but assessed after learning. It is not entirely clear why these two measures would produce different outcomes except for the fact that the trial by trial measure is happening earlier during learning.

Another exception to the overall findings of no influence of delay is the frequency strength measure. Though the frequency strength judgment was asked after learning, it is also somewhat different from the other questions and can be viewed as a less explicit measure of causal strength. This measure had people to recall their memories for the 16 events rather than make a judgment about the cause and/or effect. It is possible that this measure taps into people's episodic memories, or that this judgment is recreated from participants' beliefs about the relationship between the cause and effect; this is a topic of current study. However, the effects of delay on frequency strength were only found in the regression analyses, and though significant, the BF was very weak, suggesting that any influence of delay was quite modest.

Open Questions Though this study has been designed to push the limits of external validity of the learning paradigm - to study causal learning as closely as it occurs over long timeframes in everyday life - there are still important future directions. One is that participants received the information of an absence of the cause or the effect. However, in many real life situations, nothing alerts the learner to an absence, which raises important challenges for research (Gallistel, Craig, & Shahan, 2019). For example if someone forgets to take medicine, or does not feel pain, these absences may go unnoticed. However, modifying the paradigm for real absences will be hard to implement because if a participant fails to do a task when they receive a text message it will turn what should be a presence into an absence. Another important future direction has to do with a different sort of delay. For example, antidepressant medications can take multiple weeks before starting to have an influence and in this case the medicine is taken each day, whereas in the current study the medicine was taken on 50% of days.

Conclusions This research makes an important empirical contribution to the field of human causal learning specifically, and learning more generally, showing that learning is not necessarily degraded even with considerable delays. Empirically, this raises the possibility that people can accurately learn about the contingencies between events in their daily lives, at least in simple cases with only one cause and effect. Theoretically, this research requires a reexamination of the mental processes that underlie human causal and statistical learning, which primarily assume that learning is degraded with increasing delay.

Acknowledgments

This work was supported by NSF BCS 1651330.

References

Allan, L. G. (1980). A note on measurement of contingency between two binary variables in judgment tasks. *Bulletin of the Psychonomic Society*, *15*(3), 147–149.

Bramley, N. R., Gerstenberg, T., Mayrhofer, R., & Lagnado, D. A. (2018). Time in causal structure learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *44*(12), 1880.

Buehner, M. J. (2005). Contiguity and covariation in human causal inference. *Learning & Behavior*, *33*(2), 230–238.

Buehner, M. J., & McGregor, S. (2006). Temporal delays can facilitate causal attribution: Towards a general timeframe bias in causal induction. *Thinking & Reasoning*, *12*(4), 353–378.

Dunlap, A. S., & Stephens, D. W. (2014). Experimental evolution of prepared learning. *Proceedings of the National Academy of Sciences*, *111*(32), 11750–11755.

Gallistel, C., Craig, A. R., & Shahan, T. A. (2019). Contingency, contiguity, and causality in conditioning: Applying information theory and weber's law to the assignment of credit problem. *Psychological Review*, *126*(5), 761.

Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, *107*(2), 289.

Gluck, M. A., & Thompson, R. F. (1987). Modeling the neural substrates of associative learning and memory: A computational approach. *Psychological Review*, *94*(2), 176.

Grossberg, S., & Schmajuk, N. A. (1989). Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Networks*, *2*(2), 79–102.

Hagmayer, Y., & Waldmann, M. R. (2002). How temporal assumptions influence causal judgments. *Memory & Cognition*, *30*(7), 1128–1137.

Kruschke, J. (2014). *Doing bayesian data analysis: A tutorial with r, jags, and stan*. Academic Press.

Lagnado, D. A., & Speekenbrink, M. (2010). The influence of delays in real-time causal learning. *The Open Psychology Journal*, *3*(1).

Logue, A. W. (1979). Taste aversion and the generality of the laws of learning. *Psychological Bulletin*, *86*(2), 276.

Renner, K. E. (1964). Delay of reinforcement: A historical review. *Psychological Bulletin*, *61*(5), 341.

Rescorla, R. A. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Current research and theory*, 64–99.

Shanks, D. R., Pearson, S. M., & Dickinson, A. (1989). Temporal contiguity and the judgement of causality by human subjects. *The Quarterly Journal of Experimental Psychology*, *41*(2), 139–159.

Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of pavlovian reinforcement. In M. G. . J. Moore (Ed.), *Learning and computational neuroscience: Foundations of adaptive networks* (p. 497–537). The MIT Press.

Willett, C. L., & Rottman, B. M. (2020). Causal learning with two causes over weeks. *Proceedings of the 42nd annual conference of the cognitive science society*, 2007–2013.

Willett, C. L., & Rottman, B. M. (in press). The accuracy of causal learning over long timeframes: An ecological momentary experiment approach. *Cognitive Science*.