Title
Learning traps lead to change blindness in dynamic environments

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors
Lee, Won Jae
Li, Amy X.
Lee, Jaimie E
et al.

Publication Date
2023

Peer reviewed

Learning traps lead to change blindness in dynamic environments

Won Jae Lee[1] (w.j.lee@student.unsw.edu.au), Amy X. Li [1,2] (amy.li@psy.ox.ac.uk), Jaimie E. Lee [1] (jaimieelizabethlee@gmail.com), Brett K. Hayes[1] (b.hayes@unsw.edu.au)

[1]School of Psychology, UNSW Sydney, Sydney, NSW 2052, Australia
[2]Department of Experimental Psychology, University of Oxford, Oxford, OX2 6GG

Abstract

The ability to selectively attend to stimuli increases the efficiency of learning. However, learning traps can develop when attention prematurely narrows to a subset of the features that predict outcomes, resulting in suboptimal decisions. The current work investigated the potential for learning traps to be particularly damaging in dynamic environments, where the features that predict rewards and losses change during learning. Two experiments (N=316) found that when learners received choice-contingent feedback, they frequently fell into a learning trap, using a suboptimal categorisation rule. Critically, these learners were unlikely to detect a subsequent rule change nor learn the new optimal rule. This change blindness was not attenuated by priming participants to expect change. These results show that the pernicious effects of learning traps are amplified in dynamic environments.

Keywords: exploration; decision-making; category learning; selective attention

Introduction

Learning from experience is critical for us to make informed decisions that shape our everyday lives. However, in certain situations, the mechanisms that facilitate experiential learning can lead to suboptimal decision-making. For example, imagine that you have moved into a new neighborhood and are exploring local cafés to discover who serves the best coffee. Early in your search, you visit a café where the coffee is not to your liking. Based on that experience you may form the belief that the café generally serves poor coffee, and avoid it in the future. This behavior would be adaptive if your belief was true. However, if the belief was false (e.g., you were served by a trainee so your coffee was below the café's usual standard), then avoidance means you will miss out on potentially rewarding experiences. Moreover, because we typically only receive feedback about choice options that we approach, the negative first impression about the café will not be corrected through further learning.

This example illustrates a *learning trap* (Erev, 2014; Rich & Gureckis, 2018) – a pattern of under-exploration of choice options based on false beliefs formed early in learning, which can lead one to miss available rewards. Learning traps can have far more serious consequences than missing out on good coffee. They can lead to persistent false impressions about others in social encounters and the formation of group stereotypes (Denrell, 2005; Le Mens & Denrell, 2011). They can also lead to under-exploration of alternatives, which may result in poor management and financial decisions (Denrell & March, 2001; Teodorescu & Erev, 2014a). Learning traps may also contribute to the under-exploration of rewarding prospects in clinical depression (Teodorescu & Erev, 2014b).

Learning traps may take a number of forms. One type of trap may arise when payoffs from choice prospects vary over time (sometimes negative, sometimes positive). Experience of a negative outcome following choice of a given prospect can lead to a "hot stove" effect such that the prospect is subsequently avoided, with a loss of potential rewards (Denrell, 2007; Denrell & March, 2001).

In the hot-stove case, previously experienced outcomes are the only guide for future choices. In many environments, however, learners interact with stimuli made up of multiple feature dimensions (Murphy, 2005) and use these features to predict choice outcomes (Schulz, Konstantinidis & Speekenbrink, 2018; Sher et al., 2022). In such situations, early experience of a negative outcome can lead people to form overly simplistic beliefs about the predictive relationship between stimulus features and choice outcomes. This means that people may fail to approach future choice options that actually yield a reward. Returning to the café example, you may note the name of the café company where you received the bad cup of coffee. Generalizing from your initial negative experience, you may subsequently avoid other café franchises that share that company logo, even though some actually serve good coffee.

Rich and Gureckis (2018) developed an experimental model of such traps in tasks where different categories (e.g., types of cartoon bees; job applicants with different profiles) were associated with either rewards or losses. A conjunctive rule involving two feature dimensions was a perfect predictor of the category bound and associated outcomes (e.g., approaching bees with two legs and single wings led to a loss of points; approaching bees with other feature combinations led to gains). On each learning trial, participants could choose to approach an exemplar and receive the consequent gain or loss, or avoid the exemplar (with no gain/loss). When outcome feedback was contingent on a decision to approach, a large proportion of participants fell into a learning trap, relying on a simple "one-dimensional" categorization rule (e.g., "avoid bees with two legs"). As a result, these participants missed out on additional rewards. In contrast, most in a baseline "full feedback" condition, who were provided with outcome feedback for both approached and forgone instances, learned the correct two-dimensional category rule.

This type of trap has been robustly replicated in a variety of paradigms where learners are required to make decisions about stimuli based on multiple feature dimensions (Blanco,

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.

Turner & Sloutsky, 2023; Lee, Li & Hayes, 2022; Li, Gureckis & Hayes, 2021; Liquin & Gopnik, 2022). These studies have shown that such traps emerge relatively early in the learning process and persist over extended periods of subsequent category learning.

The Current Studies: Learning Traps in Dynamic Environments

An important limitation of most previous studies of learning traps is that they involved the learning of static categorization rules: the feature configurations that allowed one to predict which instances were associated with positive or negative outcomes remained unchanged throughout learning (but see Blanco et al., 2023 for an exception). However, many if not most learning environments outside the laboratory are *dynamic*, where relationships between stimulus features and outcomes can change (e.g., Brown & Steyvers, 2009; Navarro, Newell & Schulze, 2016; Schulz, et al., 2018). In our café example, staff and coffee-making practices are likely to change over time leading to changes in the quality of their product. Likewise, the features that predict which companies listed on the stock market are likely to yield the best returns are likely to change in response to fluctuations in the broader economic environment.

However, no studies to date have examined how learning traps function in dynamic learning environments. This is surprising because, if anything, learning traps have the potential to be even more pernicious in dynamic than in static learning environments. If you are avoiding all outlets of a particular coffee company, then you will never discover that they have improved their barista training and consequently, the quality of their coffee. The learning trap means that not only are you missing out on currently available rewards, but you will be blind to changes in the features that predict rewards and losses.

The aim of the current studies, therefore, was to investigate the consequences of learning trap formation in a dynamic environment. The first stage of each study was similar to Rich and Gureckis (2018), where participants were tasked with learning which categories of cartoon bees were friendly (i.e., with a gain in points if approached) and which bees were dangerous (i.e., with a loss in points if approached). A conjunctive category rule involving two feature dimensions was a perfect predictor of category membership (e.g., instances with feature values "0" on both dimensions A and B, were dangerous; those with other feature combinations were friendly – see Figure 1 for an example).

Participants in a *contingent feedback* condition only received outcome feedback about exemplars they approached but not those they avoided. Based on previous results (e.g., Lee, et al., 2022; Li et al., 2021; Rich & Gureckis, 2018), we expected that a large proportion of learners in this condition would fall into a one-dimensional learning trap (e.g., approach bees with six legs, avoid bees with two legs) in the first few learning blocks.

Crucially, as illustrated in Figure 1, later in learning *there was a change in the conjunctive category rule* (e.g., the feature combination predicting dangerous bees was changed from "A0B0" to "A0B1"). In Experiment 1, this change was not signaled. In Experiment 2, participants were primed to expect change.

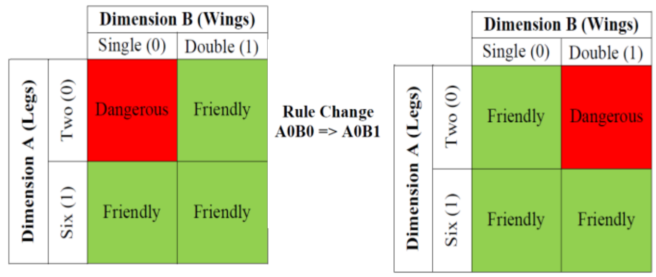

Figure 1: Illustration of an initial and changed two-dimensional category rule.

Our key prediction was that how people respond to this dynamic change in the environment depends on whether they are already in a one-dimensional learning trap before the change. If they are, and if the rule change involves a feature that they are currently avoiding, then they should be blind to the change. For example, if they had learned a one-dimensional rule based on Dimension A (legs), such that they were approaching stimuli with six legs but avoiding stimuli with two legs, then they would be unlikely to notice a rule change like that shown in Figure 1. Consequently, their decisions about which instances to approach or avoid should remain the same; they will continue to earn suboptimal levels of reward and will not learn the new categorization rule. Note that this prediction does not apply to people who initially learned a one-dimensional rule involving the alternate feature dimension (e.g., approach items with double wings, avoid items with single wings).

In contrast, if people initially learned the correct two-dimensional category rule, then they should notice the rule change (regardless of which features are involved). After the change occurs, these participants should approach stimuli that were previously deemed safe but which now result in a loss. This should lead to some confusion and a reduction in rewards earnings, but this is likely to be temporary as people gradually learn the new two-dimensional rule.

Experiment 1

This study tested our predictions about the consequences of learning trap formation for learners' sensitivity to dynamic change in the learning environment, when this change was not signaled. Our key predictions apply to the *contingent feedback* condition, where outcome feedback was contingent on a decision to approach a stimulus. In Experiment 1, we also examined how people responded to a dynamic change in the relevant categorization rule when they received *full feedback* about outcomes regardless of whether an instance was approached. Previous work suggests that people receiving such feedback are unlikely to fall into an early learning trap; most should learn the correct two-dimensional rule prior to the rule change. Hence, they should also notice the later rule change and are likely to eventually learn the new rule. This condition served as a baseline, such that we could observe how people responded to dynamic change under optimal feedback conditions.

Method

Participants

We recruited 214 adults (M_{age} = 37.25 years, 108 males, 103 females, 3 other) through the Prolific online platform. Equal numbers were randomly assigned to contingent or full feedback conditions. Participants were paid a base reward of £2.00 for task completion, and could earn a bonus of up to £1.70 based on points accrued during the task. Each point earned corresponded to a bonus of £0.01.

Materials and Procedure

We adapted the paradigm used by Rich and Gureckis (2018) for the present study. Training stimuli were images of cartoon bees constructed from a combination of three binary-valued visual feature dimensions. Feature dimensions were legs (two or six), body (spotted or striped) and wings (single or double), generating 8 unique stimuli (see Figure 2 for an example). Two dimensions were relevant to categorization. A conjunctive rule involving values on these two dimensions predicted which bees were friendly or dangerous (see Figure 1). One feature dimension was irrelevant; values on this dimension were equally likely to be associated with friendly or dangerous bees. Assignment of specific dimensions as relevant/irrelevant and the feature combination that predicted dangerous bees at the start of learning, were randomly determined by the experimental program for each participant.

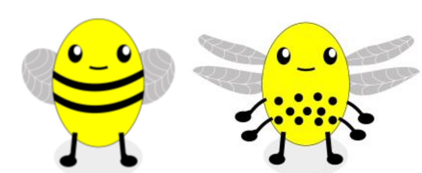

Figure 2: Examples of bee exemplars with differing values on the three feature dimensions *legs* (two vs. six), *body* (striped vs. spotted), and *wings* (single vs. double).

Participants were told that they were beekeepers tasked with collecting as much honey as possible (and associated reward points) from different beehives. Some bees could be friendly and give honey (+1 point when approached); other bees would be dangerous and sting (-3 points).

Learning Phase 1 (Pre-Change). On successful completion of a short test of comprehension of the experimental instructions, participants progressed to learning phase 1. This consisted of 4 blocks each consisting of 16 trials. In each block, the 8 unique bee stimuli were presented twice in random order. Hence, within each block there were 4 dangerous and 12 friendly stimuli. Block transitions were not indicated to participants.

On each trial, participants were presented with an image of a bee stimulus and made an approach or avoid decision by clicking the appropriate on-screen button. Approaching a bee led to points gain or loss depending on whether the item was friendly or dangerous. If the bee was avoided, no points were gained or lost. Those in the contingent feedback group received outcome feedback only when they decided to approach. Those in the full feedback group were given outcome feedback on every trial, including those where the bee was avoided. Feedback about the outcome ("harvested honey" or "stung") and points gain/loss was presented on a separate screen that appeared after a response was made. A cumulative point tracker was visible on screen throughout learning.

Learning Phase 2 (Post-change). After block 4, the categorization rule changed. From this point, a new combination of relevant dimension features predicted dangerous bee stimuli. The rule change involved switching the binary feature value on one of the relevant dimensions (as illustrated in Figure 1). As such, some bees that were friendly in learning phase 1 now led to a points loss upon approach, and vice versa. Which of the relevant dimensions was subject to this change was determined randomly for each participant. Participants completed a further 6 learning blocks in this phase. The rule change was not signalled to participants. At the end of the learning, participants were asked if they were aware of the rule change (yes vs. no/unsure).

Dimensional Rule Scoring. Participants' approach/avoid decisions in each block of 16 trials were used to identify the type of categorization rule that they were using. Perfect conformity to the optimal two-dimensional (2D) rule would result in approaching all 12 friendly bees and avoiding all 4 dangerous bees (a net gain of 12 points for that block). Perfect conformity to a one-dimensional (1D) rule would result in approaching 8 friendly bees but avoiding 4 dangerous and 4 friendly bees (i.e., a net gain of 8 points). In a given block, participants were classified as "2D rule users" or "1D rule users" if their approach choices were consistent with the relevant rule on at least 15 out of 16 trials. If the choice pattern did not meet either criterion, the participant was said to be using an "unclassified rule".

The same general process was used to identify participants' rule use in learning phases 1 and 2. However, approach/avoid patterns in each phase were compared against different categorization rules (i.e., relevant pre-change and post-change rules as illustrated in Figure 1).

Results and Discussion

Figure 3 (upper panels) shows the proportion of participants using 1D, 2D, or unclassified rules in each learning block. In both feedback conditions, 2D rule use increased as participants advanced through learning phase 1, with unclassified rule use becoming less prevalent.

Prior to the rule change (block 4), the relative proportions of participants using the optimal 2D rule or a 1D rule depended on the type of feedback provided. In the full feedback group, a substantial proportion of participants learned the 2D rule, while use of the 1D rule was rare. In the contingent feedback condition, however, use of a 1D rule (signifying the formation of a learning trap) was common.

Learning of the post-change rule was also affected by feedback type. When full feedback was provided, over 60% of participants eventually learned the new 2D categorization rule. In the contingent feedback condition, learning of the new rule was less common, and persistent use of a 1D rule was more common.

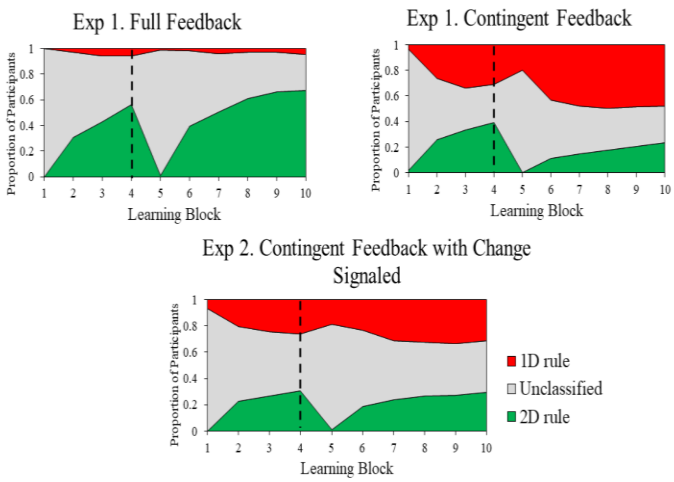

Figure 3: Proportion of participants using categorization rules over learning blocks.

These results were confirmed in multinomial logistic regression analyses which compared the ability of models that included feedback and block number to predict rule use to a null model without these predictors. Likelihood ratio tests showed that adding each predictor significantly improved model predictions of rule use in both learning phase 1 (feedback: $\chi^2(4) = 81.52$, $p<.001$; block: $\chi^2(6) = 200.43$, $p<.001$) and learning phase 2 (feedback: $\chi^2(4) = 390.51$, $p<.001$; block: $\chi^2(6) = 194.54$, $p<.001$).

Pre-change rule use and post-change rule use

Our key prediction was that the categorization rule that an individual used *prior to the rule change would influence rule learning after change*. To examine this, we defined four sub-groups within the contingent feedback condition[1], according to participants' rule use in the last block of learning phase 1: 2D, 1D attended, 1D unattended and unclassified. Those in the 1D attended and 1D unattended groups both showed a pattern of early approach/avoidance that was consistent with a suboptimal one-dimensional category rule. Those in the 1D unattended group learned a one-dimensional rule that led to avoidance of instances whose feature status changed after block 4 (e.g., in reference to Figure 1, "approach instances with six legs, avoid instances with two legs"). This group was predicted to be the least likely to detect the rule change and learn the new rule. In contrast, those in the 1D attended group learned a one-dimensional rule involving the alternate dimension (e.g., approach instances with double wings, avoid instances with single wings), and so were more likely to notice the rule change. Figure 4 shows the proportion of participants in these sub-groups using various rules at the end of learning.

It is evident that pre-change rule use had a profound effect on subsequent learning. A multinomial logistic regression that included pre-change rule use as a predictor of final rule use provided a better fit to the data than a null model, ($\chi^2(10) = 68.649$, $p<.001$). To better understand the magnitude and direction of the effect of pre-change rule use, we examined the relevant Wald statistics and Odds Ratios (OR). Participants who were using a 2D rule before the change were 8 times more likely to learn the new optimal categorization rule by the end of learning, compared to participants who were using a 1D rule before the change that involved the unattended dimension (OR = 8.00, *95% CI* for OR [1.61, 39.69], Wald's $\chi^2 = 6.48$, $p=.01$). Most participants using a 1D unattended rule in learning phase 1 (19 out of 21) failed to learn the new rule and continued to use the same 1D rule through subsequent blocks. By comparison, most of those in the 1D-attended subgroup did not continue to use that rule after the change. However, most in this sub-group also failed to learn the new 2D rule (i.e., at the end of learning most were using an unclassified rule).

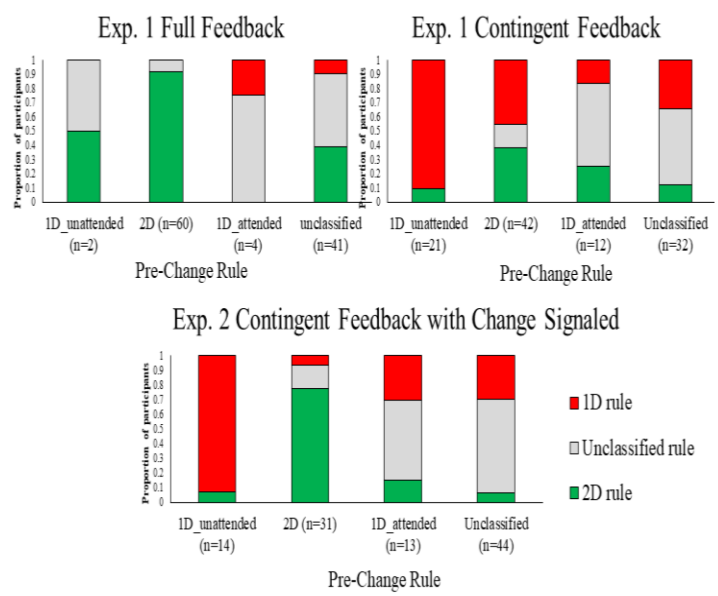

Figure 4: Proportion of participants using categorization rules at the end of learning (block 10) as a function of pre-change rule use.

Responses to the post-test question about noticing a rule change were in line with the rule-use results. A majority in the contingent feedback condition using a 2D rule (69%) or a 1D attended rule (67%) before the rule change, noticed the change (the corresponding proportions for the full feedback

[1] The pattern of results did not change substantially if those in the full-feedback condition were included, except that the proportion learning the post-change 2D rule increased.

condition were 88% and 100%). In contrast, no one in the 1D unattended sub-group noticed the rule change.

When feedback was contingent on approaching an exemplar, a considerable proportion of participants fell into a one-dimensional learning trap within the first few learning blocks. This replicates a key finding of previous studies (e.g., Li et al., 2021; Rich & Gureckis, 2018).

Those who acquired an optimal 2D categorization rule early in learning generally adapted to an unsignaled change in the rule. Most eventually learned the new 2D rule. In contrast, those who had already fallen into a learning trap such that they were not attending to the dimension where the change occurred (1D unattended group), were blind to this change. They did not report noticing a rule change and most never learned the new 2D rule.

We did not have strong predictions about the final rule learning in the 1D-attended sub-group. Most of these participants noticed the rule change and abandoned their 1D rule by end of learning. However most never learned the new 2D rule. At a minimum, this suggests that awareness of rule change is necessary but not sufficient for learning the new rule. We return to this issue in the General Discussion.

Experiment 2

The previous study found that falling into a learning trap that involves use of an overly simplistic category rule means that people will often be blind to unsignaled changes to the rule. This was despite learners having several additional blocks of training after the rule change. Experiment 2 examined whether this change blindness is reduced by signaling the possibility of a rule change.

Previous work (e.g., Navarro et al., 2016) has shown that instructions which lead people to expect a change in their learning environment can increase exploration of different choice options. In our paradigm, having an expectation that the categorization rule may change, could lead to greater exploration of exemplars. In particular, after forming an initial belief about which exemplar features predict gains or losses, people may occasionally approach bees that they believed were dangerous to discover whether the categorization rule had changed. This could lead to higher levels of change detection and learning of the new 2D rule. However, it remains to be seen whether the signaling of possible change is sufficient to reduce change blindness in those who have fallen into an early 1D learning trap.

Method

Participants

102 participants (M_{age} = 37.15 years, 69 males, 30 females, 3 other) were recruited through the Prolific platform, with the same payment arrangements as Experiment 1.

Procedure

All participants were allocated to a contingent feedback group where the possibility of change in the learning rule was signaled before learning commenced. The data from this "change signaled" condition was compared to the contingent feedback condition in Experiment 1, where rule change was not signaled. There was no full feedback condition.

The procedure for the new change-signaled condition was identical to the contingent feedback group in the previous experiment except that the initial instructions contained the following: "*Bees can sometimes experience a seasonal change in behavior. As such, the features that predict which bees are friendly and which bees are dangerous can change. This seasonal change happens at most ONCE during the experiment*." The instructions did not indicate when this rule change would occur. For all participants, this change took place after learning block 4, as in the previous study. Before proceeding to the training phase, participants had to pass a brief comprehension test that demonstrated they understood the instructions, including the expectation of rule change.

Results and Discussion

Rule use across learning blocks is shown in Figure 3. Multinomial logistic regressions compared rule use in the change-signaled condition with the contingent feedback condition of Experiment 1. Separate analyses were carried out for pre-change and post-change learning phases. In the pre-change phase, signaling the possibility of change reduced the likelihood that participants would use a 2D rule over an unclassified rule (OR = 0.61, 95% CI [0.42, 0.88], Wald's χ^2 = 6.98, p=.008). In the post-change phase, however, receiving signaling instructions more than doubled the likelihood of participants using a 2D rule over a 1D rule (OR = 2.15, 95% CI [1.55, 2.98], Wald's χ^2 = 20.88, p<.001). That is, signaling participants to expect change reduced learning of defined categorisation rules during early learning, but increased eventual learning of the new 2D rule.

Pre-change rule use and post-change rule use

Participants were again divided into sub-groups based on the categorization rule they used in the last block of the pre-change stage. These subgroups and their respective rule use in the final learning block is shown in the lower panel of Figure 4.

To examine how change signaling affected final rule learning in the various sub-groups, we ran multinomial regressions on data from the new change-signaled group and the unsignaled contingent feedback group from Experiment 1. As shown in Figure 4, for the subgroup who learned a 2D rule in the pre-change phase, participants were much more likely to learn the new 2D rule when change was signaled than when it was unsignaled, (OR = 14.25, *95% CI* for OR [2.91, 69.77], Wald's χ^2 = 10.75, p=.001). However, signaling change had little effect on learning of the new rule in any of the other pre-change sub-groups (lowest p =.38). In particular, there was little effect of change signaling on learning of the new rule for those in the 1D unattended sub-group, (OR = 0.73, *95% CI* for OR [0.06, 8.92], p=.81). In the signaled condition, nearly all participants in this sub-group (93%) remained insensitive to the rule change, continuing to use a 1D rule at the end of learning. This was similar to the proportion in the corresponding sub-group in the contingent feedback condition in Experiment 1 (90%).

In sum, for those who avoided the one-dimensional trap early in learning, signaling the possibility of rule change enhanced learning of the new 2D rule. However, change-signaling provided no benefit for those who had fallen into a 1D trap early in learning – they continued to remain blind to the change in the categorization rule when the change involved features on an unattended dimension.

The results for the post-test probe of awareness of rule change followed a similar pattern. In the signaled condition, most of those in the 2D (94%) and 1D-attended subgroups (77%) noticed the change. These proportions were higher than for the corresponding subgroups in Experiment 1. Change signaling, however, had little impact on awareness of rule change in the 1D-unattended subgroup (only one participant noticed the change).

General Discussion

Previous work has shown that learning with choice-contingent feedback can lead people to fall into a learning trap – an overly simplistic representation of the environment that reduces exploration of potentially rewarding options (Li et al., 2021; Rich & Gureckis, 2018). The current work extended the investigation of learning traps by examining the consequences of trap formation in a dynamic environment where learners had to detect and adapt to change.

In two studies, choice-contingent feedback led many people to fall into a learning trap within the first few learning blocks. This trap involved use of a one-dimensional rule, whereby participants made approach/avoid decisions based on the features of a single stimulus dimension.

The important novel finding was that this early trap formation had significant consequences for subsequent learning in the dynamic environment. Those in the one-dimensional learning trap were oblivious to a change in the categorization rule, when this change involved features in the unattended dimension. This sub-group continued to use the same one-dimensional rule throughout the course of learning. Signaling the possibility of change in Experiment 2 did little to alleviate this change blindness.

In contrast, a majority of those who avoided the trap and learned the correct two-dimensional rule early in learning, noticed the subsequent rule change and successfully learned the new 2D rule. For this group, signalling change enhanced learning of the new rule.

How might learning traps blind learners to changes in the stimuli? When learning to categorize multidimensional stimuli, the formation of attentional learning traps is hypothesized to be driven by selective attention to one stimulus dimension based on an incorrect understanding of which dimensions are relevant to categorisation (Rich & Gureckis, 2018). Once a learner falls into a learning trap, the blindness effect may be driven by *learned inattention*; over the course of learning people reduce attention to features that are seen to be irrelevant for accurate categorization (Hoffman & Rehder, 2010). Hence, those in a learning trap become less likely to attend to dimensions other than the one they are currently using to guide categorization decisions. Such overly selective attentional has been shown in other learning contexts to cause people to miss changes in the reward structure of the environment (e.g., Blanco, et al., 2023).

Consistent with this explanation is our finding that not all learning traps led to change blindness. This effect was limited to traps where change occurred in the unattended feature dimension (e.g., when the trap was based on exclusive attention to features of Dimension A but the rule change involved features from Dimension B). When the initial learning trap involved attention to the dimension that was subsequently involved in the rule change (i.e., the 1D-attended sub-group), participants often noticed the change. Somewhat surprisingly, however, unlike those who initially used a 2D rule, most participants in the 1D-attended subgroup failed to learn the new rule. The reason for this remains unclear. It could reflect individual differences in susceptibility to learning traps, arising from differences in learning efficiency (i.e., ability to learn the associations between features and outcomes) and/or the propensity to explore novel stimuli rather than exploit known options (e.g., Chang, Jara-Ettinger & Baskin-Sommers, 2022; Sang, Todd & Goldstone, 2011).

Our results show that the pernicious effects of learning traps extend beyond what has been recognised in previous research. In the sorts of dynamic environments that are common outside the laboratory, those who have fallen into an early learning trap are likely to be insensitive to environmental change. In economic contexts, this could lead to further shortfalls in the earning of potential rewards. In social contexts, it could cause one to miss important changes in relationships between individuals or social groups.

Given the potentially serious negative effects of learning traps in both static and dynamic environments, it is important to examine ways of reducing trap prevalence. To date, success in achieving this goal has been modest. Rich and Gureckis (2018) attempted to reduce trap formation by implementing changes designed to promote attention to multiple feature dimensions (e.g., individuating exemplars, adding stochasticity to outcomes, occluding some features). Unfortunately, none of these interventions led to trap reduction or increased learning of the optimal rule.

Lee et al. (2022) had more success by providing learners with “summative feedback”, which allowed them to compare current rewards earnings with optimal earnings. When such feedback was provided frequently during learning, fewer people fell into a learning trap and more learned the correct categorization rule. However, this result was obtained using a static version of the learning traps task, and whether summative feedback can reduce change-blindness in dynamic environments remains an open question. More generally, the future development of successful methods for ameliorating the negative effects of learning traps will require a better understanding of the fundamental learning and attentional mechanisms that give rise to traps.

Acknowledgements

This work was supported by an Australian Research Council Discovery Grant DP220101592 to BKH.

References

Blanco, N. J., Turner, B. M., & Sloutsky, V. M. (2023). The benefits of immature cognitive control: How distributed attention guards against learning traps. *Journal of Experimental Child Psychology, 226*, 105548. https://doi.org/10.1016/j.jecp.2022.105548

Brown, S. D., & Steyvers, M. (2009). Detecting and predicting changes. *Cognitive Psychology*, *58*(1), 49-67. https://doi.org/10.1016/j.cogpsych.2008.09.002

Chang, S. A. A., Jara-Ettinger, J., & Baskin-Sommers, A. (2022). Resource scarcity compromises explore-exploit decision-making. *Journal of Experimental Social Psychology, 98*, 104254. https://doi.org/10.1016/j.jesp.2021.104254

Denrell, J. (2005). Why most people disapprove of me: Experience sampling in impression formation. *Psychological Review, 112*(4), 951–978. https://doi.org/10.1037/0033-295X.112.4.951

Denrell, J., & Le Mens, G. (2007). Interdependent sampling and social influence. *Psychological Review, 114*(2), 398–422. https://doi.org/10.1037/0033-295X.114.2.398

Denrell, J., & March, J. G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science, 12*(5), 523–538. http://dx.doi.org/ 10.1287/ orsc.12.5.523.10092

Erev, I. (2014). Recommender systems and learning traps. In M. Ge & F. Ricci (Eds.), *Proceedings of the first international workshop on decision making and recommendersystems* (pp. 38–41). Free University of Bozen-Bolzano.

Hoffman, A. B., & Rehder, B. (2010). The costs of supervised classification: The effect of learning task on conceptual flexibility. *Journal of Experimental Psychology: General*, *139*(2), 319. https://doi.org/10.1037/a0019042

Lee, J. E., Li, A. X., & Hayes, B. K. (under review). Overcoming learning traps with summative feedback.

Le Mens, G., & Denrell, J. (2011). Rational learning and information sampling: On the "naivety" assumption in sampling explanations of judgment biases. *Psychological Review*, *118*(2), 379–392. http://dx.doi.org/10.1037/a0023010

Li, A. X., Gureckis, T. M., & Hayes, B. (2021). Can losses help attenuate learning traps? In T. Fitch, et al. (Ed.s) *Proceedings of the 43rd Annual Meeting of the Cognitive Science Society*, pp. 1201-1207.

Liquin, E. G., & Gopnik, A. (2022). Children are more exploratory and learn more than adults in an approach-avoid task. *Cognition, 218*, 104940. https://doi.org/https://doi.org/10.1016/j.cognition.2021.104940

Murphy, G. L. (2005). The study of concepts inside and outside the laboratory: Medin versus Medin. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. W. Wolff (Eds.), *Categorization inside and outside the laboratory* (pp. 179–195). Washington, DC: American Psychological Association.

Navarro, D. J., Newell, B. R., & Schulze, C. (2016). Learning and choosing in an uncertain world: An investigation of the explore–exploit dilemma in static and dynamic environments. *Cognitive Psychology, 85*, 43-77. https://doi.org/10.1016/j.cogpsych.2016.01.001

Rich, A. S., & Gureckis, T. M. (2018). The limits of learning: Exploration, generalization, and the development of learning traps. *Journal of Experimental Psychology: General, 147*(11), 1553–1570. https://doi.org/10.1037/xge0000466

Sang, K., Todd, P. M., & Goldstone, R. L. (2011). Learning near-optimal search in a minimal explore/exploit task. *Proceedings of the Thirty-Third Annual Conference of the Cognitive Science Society*. (pp. 2800-2805). Boston, Massachusetts: Cognitive Science Society.

Schulz, E., Konstantinidis, E., & Speekenbrink, M. (2018). Putting bandits into context: How function learning supports decision making. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *44*(6), 927–943. https://doi.org/10.1037/xlm0000463

Sher, S., McKenzie, C. R. M., Müller-Trede, J., & Leong, L. (2022). Rational choice in context. *Current Directions in Psychological Science*, *31*(6), 518–525. https://doi.org/10.1177/09637214221120387

Teodorescu, K., & Erev, I. (2014a). On the decision to explore new alternatives: The coexistence of under- and over-exploration. *Journal of Behavioral Decision Making, 27*, 109-123. https://doi.org/10.1002/bdm.1785

Teodorescu, K., & Erev, I. (2014b). Learned helplessness and learned prevalence: Exploring the causal relations among perceived controllability, reward prevalence, and exploration. *Psychological Science*, *25*(10), 1861–1869. https://doi.org/10.1177/0956797614543022