Title:

Cognitive resources predict value learning for outcome-irrelevant features

Authors:

Ido Ben Artzi¹, Roy Luria^{1,2}, Nitzan Shahar^{1,2}

Affiliations:

¹School of Psychological Sciences, Tel Aviv University, Tel Aviv, Israel

²Sagol School of Neuroscience, Tel Aviv University, Tel Aviv, Israel

Correspondence concerning this article should be addressed to:

Ido Ben Artzi

School of Psychological Sciences

Tel Aviv University

Tel Aviv, Israel

Email: idobenartzi@mail.tau.ac.il

Conflict of interest. The authors declare no conflicts of interest.

Abstract

To establish accurate action-outcome associations in the environment, individuals must refrain from assigning value to outcome-irrelevant features. However, reinforcement learning studies have largely ignored the role of attentional control processes on credit assignment (the process of assigning value to one's actions). In the current study, we examined the extent to which working memory – a system that can filter and block the processing of irrelevant information in one's mind – predicted credit assignment to outcome-irrelevant task features. One hundred and seventy-four individuals completed working memory capacity and outcome-irrelevant learning estimates. Outcome-irrelevant learning was estimated in a reinforcement learning task where only stimulus' visual features predicted reward, but not the response keys used to indicate one's selection. As expected, we found a consistent tendency to assign value to the tasks' response keys, reflecting outcome-irrelevant learning at the group level. However, we also found substantial individual differences, such that only 55% of participants demonstrated this effect. Importantly, working memory capacity significantly moderated individual differences in outcome-irrelevant learning; individuals with higher capacity were less likely to assign credit to the outcome-irrelevant feature (i.e., response key). We discuss the influence of working memory on outcome-irrelevant learning through the perspective of cognitive control failure.

Introduction

Throughout their lives, individuals infer causal associations between actions and outcomes as they experience and navigate dynamic and complex environments. Importantly, forming these associations requires cognitive flexibility and cognitive resources. Most environments are feature-rich, thus requiring the human agent to filter out irrelevant information when forming action-outcome associations in their minds ([Collins & Shenhav, 2022](#); [Niv, 2019](#); [Wilson & Niv, 2012](#)). For example, think of a child considering whether to eat an apple or a pear. The visual and tactile cues of each fruit should predict its taste, whereas its position on a table and the hand with which it is taken should be deemed irrelevant and thus not assigned value nor be considered in the child's choice of fruit. Recent studies have indicated that individuals attribute credit to outcome-irrelevant features of an action, even when they hold certain and explicit knowledge that those features have no predictive value for the outcome, a phenomenon we coined "outcome-irrelevant learning" ([Rmus, Zou, et al., 2021](#); [Shahar et al., 2019, 2021](#)). Studying outcome-irrelevant learning should shed light on the manner in which people assign credit to various features during reinforcement learning, yet the underlying processes that moderate outcome-irrelevant information during credit assignment are still largely unknown.

Most studies that have examined credit assignment processes in the reinforcement learning literature have disregarded the specific process of updating in regard to outcome-irrelevant information ([Daw et al., 2011](#); [Dolan & Dayan, 2013](#); [Niv et al., 2015](#)). However, outside the realm of reinforcement learning, attentional control studies have extensively examined the influence of irrelevant information on individuals' choices ([Folk et al., 1994](#); [Theeuwes & Burger, 1998](#); [Zanto & Gazzaley, 2009](#)). A well-known attentional control system is working memory. Working memory processing is known to reduce attention to irrelevant features of a stimulus including, for example, shape, color, or location of target stimuli ([Engle & Kane, 2004](#); [Kane & Engle, 2003](#); [Luria & Vogel, 2011](#)). Studies have systematically shown that even when observers have explicit and certain knowledge regarding the irrelevance of particular characteristics of a stimulus, attention regulation processes remain imperfect, resulting in a consistent influence of irrelevant information on decision-making. ([Engle & Kane, 2004](#); [Kane & Engle, 2003](#); [Luria & Vogel, 2011](#)). A distinct feature of working memory is its limited capacity ([Conway et al., 2003](#); [Cowan, 2010](#); [Luck & Vogel, 1997](#)). Previous studies have highlighted individual differences in working memory capacity, demonstrating that lower capacity is associated with a reduced ability to filter out irrelevant information ([Allon & Luria, 2017](#); [Conway et al., 2001](#); [Engle, 2002](#); [Kane et al., 2001](#); [Vogel et al., 2005](#)). Despite the vast knowledge regarding the limitations of working memory, surprisingly little is known regarding the role of attentional control processes in filtering and blocking outcome-irrelevant information during credit assignment and value-based decision-making.

The importance of working memory resources for reducing the likelihood of outcome-irrelevant learning is apparent when considering that outcome-irrelevant features may also be task-relevant. In an example of a child who is asked to choose one of two fruits – each of which is positioned on opposite sides of the table – the hand-related motor plans for executing the child's choice are task-relevant (as, presumably, the child will use one hand to reach the fruit

and eat it). Yet, the motor plan of using one hand or the other can be regarded as largely outcome-irrelevant, as the hand used has no (or very little) causal influence on the outcome (e.g., the fruit's taste). Therefore, we hypothesized that working memory resources are required to actively hold task-relevant information (e.g., task-related motor plans) in mind, while at the same time, refraining from assigning credit to that information if it is known to be outcome-irrelevant. Considering findings from the attentional control literature, we further predict that individuals with low, versus high, working memory capacity will be less capable to withhold credit from being assigned to outcome-irrelevant information, especially when it is relevant to an ongoing task ([Engle & Kane, 2004](#); [Fockert et al., 2001](#); [Kane & Engle, 2003](#)). It should be noted that some studies addressed working memory in the context of reinforcement learning as a one-shot learning system that perfectly retains recent action-outcome associations ([Collins, 2018](#); [Collins, Albrecht, et al., 2017](#); [Collins & Frank, 2012](#); [Yoo & Collins, 2021](#)). However, here we mostly follow the definition of working memory as an attentional system that enables the filtering of irrelevant information ([Conway et al., 2001](#); [Engle, 2002](#); [Kane et al., 2001](#); [Vogel et al., 2005](#); [Zanto & Gazzaley, 2009](#)). In the current study, we integrate reinforcement learning and attentional control literature by addressing a fundamental question of whether working memory capacity predicts individuals' ability to disregard outcome-irrelevant information during value-based learning.

In this pre-registered study, individuals performed a card game in which they were asked to choose cards to gain a monetary reward (i.e., multiple armed bandit task). Our task design allowed us to disentangle and estimate value updating to the outcome-irrelevant task features, since only cards, but not their randomly assigned response keys, predicted the monetary outcome. Furthermore, we included visual working memory encoding and retrieval phases with varying degrees of load in between trials of the reinforcement learning card task in order to examine the influence of visual working memory load on outcome-irrelevant learning. Finally, we estimated individuals' working memory capacity using a well-known change-detection task ([Balaban et al., 2019](#); [Luck & Vogel, 1997](#)). Our results replicated previous studies by demonstrating a robust group-level effect of outcome-irrelevant learning. Moreover, we found significant individual differences such that only 55% of individuals demonstrated outcome-irrelevant learning. Importantly, working memory capacity showed a substantial moderation effect, so that reward had a lower influence on the selection of outcome-irrelevant response keys for high vs. low-capacity individuals. The within-task working memory load manipulation did not influence outcome-irrelevant learning. We discuss these results by addressing the strengths and limitations of findings from reinforcement learning and attentional control literature, as well as the recent theoretical integration of the two perspectives.

Methods

Participants. 174 Prolific workers (age mean = 27.1, range 18 to 49; 80 males, 93 females, 1 other) completed three online sessions across three consecutive days in return for monetary compensation. All participants reported normal or corrected vision, no current or past psychiatric or neurological diagnosis ([see SI](#)). The study protocol was approved by the Research Ethics Council of Tel-Aviv University and all participants signed informed consent before participating in the study.

Procedure. In the first session, 200 participants performed a working memory capacity measurement. In the second and third sessions 178/174 participants respectively, completed a reinforcement learning task under three working memory load conditions (i.e., no-load, low-load, and high-load).

Reinforcement learning task. Participants completed a reinforcement learning bandit task interleaved between the memory array and test array stages of a working memory task ([Figure 1](#)). This design allowed us to examine outcome-irrelevant learning processes under different working memory load conditions. Trials started with a *memory array stage* where participants had to memorize a visual array (i.e., colored squares). Next, participants made two decisions on two sequential offers of a *reinforcement learning bandit task* where they had to choose one of two cards to gain monetary rewards. Finally, the trial ended with a *test array stage*, where participants were asked to report whether a newly presented colored square was part of the initial visual array or not. We will now describe the memory array, multi-armed bandit trials, and test array stages that were included on each trial with more details:

The memory array stage of the working memory task included the brief presentation (200 ms) of a visual array (i.e., colored squares) which the participants had to remember until the test array stage. To manipulate working memory load, we included three types of visual arrays ([see Figure S2](#)) ; (a) no-load (a fixed color for the entire block of trials), (b) low load (one colored square, with its color being randomly selected each trial by the computer), and (c) high load (four colored squares, with all colors being randomly selected by the computer each trial without replacement out of 9 possible colors). The three conditions (i.e., no load, low and high load) were manipulated between blocks. The location of the square/s was further randomized each trial by the computer ([see SI for further details](#)).

Reinforcement learning bandit task included two offers for each trial (interleaved between memory array and test array stages), which were designed to allow us to estimate credit assignment to the outcome-irrelevant response keys. In each trial, the computer allocated four cards (without return) to the two offers (i.e., first or second). Therefore, each offer allowed individuals to choose one of two cards, and these two cards were further randomly allocated to the right or left sides of the screen (see [Figure 1](#)). Offers in the reinforcement learning bandit task started with a fixation (900ms), followed by the presentation of two cards in the right/left location. Participants then choose a card freely using a right/left corresponding response-key press ('s' or 'k' keys in a QWERTY keyboard; until response with 6sec deadline). After making a choice the unchosen card disappeared and the chosen card remained to allow choice feedback

(500ms). Cards led probabilistically to reward (£0 or £1, play-pound coins; outcome presented for 750ms) according to a true expected value (which slowly drifted across trials according to a predefined random walk; [see SI](#)). Participants were asked to do their best to make choices that will maximize monetary return. To further ensure participants' motivation, we rewarded them with a monetary payment bonus at the end of the study according to their wins in the task. Importantly, only the cards predicted reward, but not locations or response-keys used to report cards' selection. This fact is important since it renders any credit assignment to the location/response key as outcome-irrelevant learning ([Shahar et al., 2019, 2021](#)). Participants were told explicitly during the instruction phase that only cards predicted reward and not the response-keys used for their selection ([see SI](#)). Before starting the task, participants were asked to complete a short quiz (9 questions long) to ensure they read and understood the instructions ([see SI](#)), which also included a specific question regarding which task features predict a monetary outcome. Participants had to show 100% accuracy in the quiz, and if they scored lower, they were prompted to the beginning of the instructions phase and were asked to retake the quiz.

Test array stage included a target screen in which one colored square was randomly selected by the computer in a color that was either the 'same' or 'different' color as the square that appeared in the same location in the memory array stage. Participants were asked to respond 'same' or 'different' by pressing 's' or 'k' response-keys (keyboard mapping was counterbalanced between participants; until response with a 6sec deadline). Following the target screen. Participants were presented with feedback indicating whether their response was correct or incorrect, which was then followed by a fixation screen (inter-trial interval, 500ms fixation).

Across two sessions participants completed a total of 6 blocks of the reinforcement learning task. Each block included a different set of cards and had 50 trials. Participants received a monetary bonus at the end of the task based on their performance.

Figure 1. Trial sequence in the reinforcement learning task, which was performed under varying working memory load. Participants were asked to first memorize a visual array

(memory array stage), then make two choices across two card offers (reinforcement learning bandit task stage), and finally report whether a target was the same or different compared to the visual array that was memorized at the trial initiation (test array stage). Working memory load was manipulated between blocks by including in the memory stage either four random colored squares (high load), one random colored square (low load), or a fixed color squared (no-load).

Working memory capacity task. To measure working memory capacity participants were asked to maintain and retrieve visual information ([Balaban et al., 2019](#); [Luck & Vogel, 1997](#); [Pashler, 1988](#)). On each trial, a visual array of colored squares appeared (set size of 4 or 8 squares; [see SI](#)). Squares in each array had distinct colors and were evenly spread across the screen. Each trial started with the presentation of the squares array (i.e., memory array phase, 200 ms), followed by a fixation cross (i.e., retention phase, 900ms) and then a target (i.e., test array phase, until response with a 6sec deadline). The target screen included one colored square randomly selected by the computer in a color that was either the ‘same’ or ‘different’ color as the square that appeared in the same location in the memorized array. Participants were asked to respond ‘same’ or ‘different’ by pressing ‘s’ or ‘k’ response-keys (keyboard mapping is counterbalanced between participants). Following the target screen participants saw feedback indicating whether their response was correct or incorrect, followed by a fixation screen (inter-trial interval, 500ms fixation). Each participant completed 120 trials aimed to identify individual differences in working memory capacity.

Estimating outcome-irrelevant learning. To estimate outcome-irrelevant learning, we examined whether the outcome in the first offer (i.e. £0 vs £1) affected response key selection in the second offer ([see Figure 2A](#)). Specifically, we reasoned that credit will be assigned not only to the chosen card but also to a response key used for its selection ([Shahar et al., 2019, 2021](#)). This would mean that after reward (i.e., £1) was obtained in the first offer, participants should be more likely to stay and choose the second offer with the same response key that was selected in the first one. However, if the choice in the first offer was unrewarded, participants should be more likely to switch their response key selection in the second offer. For example, assume that in the first offer the participant selected card A with a left response key, and observed a reward (£1). The left response key is now assumed to be more valuable in the individual’s mind, making it more likely that the left response key will be used in the second offer, where a different set of cards was offered. Therefore, outcome-irrelevant learning was assessed by estimating a regression parameter coefficient in a hierarchical Bayesian logistic regression, where repetition of response key selection from the first to the second offer (i.e., 0 vs. 1; Stay_{response key}) was predicted using the outcome of the first offer (i.e., previous outcome, £0 vs £1; [Figure 2B](#)). The positive influence of previous-outcome on Stay_{response key} was considered as evidence for value assignment to the response key (i.e., outcome-irrelevant learning).

Estimating working memory capacity. We calculated the detection accuracy in the visual working memory task (where reinforcement learning trials were not included) for each

individual in each set-size condition (8 or 4 squares) to estimate working memory capacity. In each condition, capacity was calculated according to $K = \text{set_size} \cdot (2 \cdot \text{accuracy} - 1)$, and the average between the two conditions was used as a measure of overall capacity ([Cowan, 2001](#)).

Regression analysis. Hierarchical Bayesian regression analysis was performed using the ‘brms’ and ‘rstan’ R packages ([Bürkner, 2018](#); [Carpenter et al., 2017](#)). To examine posterior distributions we report 95% high-density interval (HDI) and probability of direction (pd) For all predictors, we used a pre-registered weakly informative prior, performed prior and posterior predictive checks, and verified that the conclusion did not change when using either wider or narrower priors ([see SI](#)). Bayes factors were calculated using the ‘bayestestR’ R package ([Makowski et al., 2019](#)) for both the null point estimate and the region of practical equivalence (ROPE). We used a ROPE interval in line with recent guidelines in bayesian statistics recommending to supplement a null-point Bayes factor estimate, with a more conservative Bayes factor calculating the portion of the posterior distribution falling within a region where the parameter estimates are theoretically negligible ([Makowski et al., 2019](#)). We set the ROPE interval to be -0.013 to 0.013 based on the rationale that a previous-outcome effect of 0.33% or lower on the probability to stay with the same response key should be considered negligible.

Data Treatment. The first trial on each block and trials with implausibly quick reaction times (<200ms), or exceptionally slow reaction times (>4000ms) were omitted (5.52% of all trials). We omitted participants for the following reasons: (a) Participants with a lower than chance accuracy in the working memory capacity task with set size four (1 participant), (b) participants with more than 30% excluded trials in the reinforcement learning task (3 participants), and (c) participants who repeated the same response-key more than 80% of the trials (1 participant), in total 5 participants (age mean = 32, range 21 to 33; 3 males, 2 females) were excluded altogether.

Results

Our main aim was to examine whether outcome-irrelevant learning changed as a function of individuals’ working memory capacity, and load manipulation (i.e., no-load, low, and high load). For this aim, we fitted a hierarchical Bayesian logistic regression. The dependent variable was the individuals’ tendency to repeat response key selection from the first to the second offer of the reinforcement learning bandit task (i.e., 0 vs. 1 for *different* or *same*, respectively; Stay_{response key}). We examined three predictors including previous-outcome (i.e. £0 vs £1, play pounds), working memory load (i.e., no-load, low load, and high load), individual’s working memory capacity, and their paired and triple interactions ([see SI](#)). Following recent guidelines on Bayesian regression analysis ([Bergh et al., 2021](#)), we start by performing nested model comparisons, allowing us to drop predictors that have a negligible contribution to the prediction of response key selection, and continue by examination of the posteriors of the best performing model. We assembled four nested models that included the following fixed effects: (a) Model 1 (full) previous-outcome (unrewarded vs. rewarded; describing outcome-irrelevant learning), set-size manipulation, working memory capacity, and all paired and triple interaction as predictors. (b) Model 2 (set-size) excluded working memory capacity parameters, (c) Model

3 (capacity) included working memory capacity parameters but excluded set-size parameters, and finally, (d) Model 4 (null) included only previous-outcome as a predictor. Since we were interested in the population level (fixed effects) parameters, all models included a random effect of subjects on intercept, previous-outcome, set-size, and previous-outcome x set-size parameters. Leave-one-out cross-validation model comparison ([Vehtari et al., 2017](#)) indicated that the best explanatory model was Model 3 (with previous-outcome, capacity, and their paired interaction). Specifically, stacking weights with uniform model priors showed 71% support for Model 3, 26% for Model 1, 3% for Model 2, and 0% for Model 4. We then continued to examine the posterior parameter's distribution for the best fitting model, Model 3.

We found strong evidence for outcome-irrelevant learning, such that participants were more likely to stay with their response key selection after the first offer was rewarded (48%) vs. unrewarded (45%; posterior median =0.09, HDI_{95%} between 0.03 and 0.16; probability of direction (pd) 99.75%; 0% in ROPE (-0.013 - 0.013) and Bayes Factor (BF) of 7.66 against the null and of 7.29 against the modified ROPE; [Figure 2C](#)). This replicates previous findings of outcome-irrelevant learning in human individuals ([Rmus, Zou, et al., 2021](#); [Shahar et al., 2019, 2021](#)). Importantly, we found support for our pre-registered hypothesis that individuals with low working memory capacity will demonstrate increased outcome-irrelevant learning compared to high capacity individuals such that the interaction between previous-outcome and capacity was negative (posterior median =-0.14, HDI_{95%} between -0.22 and -0.07; probability of direction (pd) 100%; 0% in ROPE (-0.013 - 0.013) and Bayes Factor (BF) of 280 against the null and of 289 against the modified ROPE; [Figure 2E](#)). [Figure 2D](#) illustrated a point-estimate of individual coefficients for previous-outcome (posterior mean) as a function of their estimated capacity. To provide the reader with a more intuitive effect-size of the association between working memory capacity and outcome-irrelevant learning, we calculated and found a correlation of $r=-0.31$ (CI_{95%} -0.44 and -0.17) between these two estimates.

To assert participants complied with our instructions and performed both the working memory and reinforcement learning components above chance, we examined accuracy rates. For the reinforcement learning task component, we found that participants were able to choose the card with the higher true expected value in 57% of trials. For the working memory task component, we found 94%, 88%, and 67% accuracy rates for no-load, low-load, and high-load, respectively. Thus, we conclude that performance was well beyond random in all parts of the task ([see SI](#) for Bayesian analysis of accuracy rates).

To sum up, results demonstrated support for outcome-irrelevant learning at the group level and supported the mediation of working memory capacity on outcome-irrelevant learning estimates. However, our results did not support our hypothesis for interaction between working memory load and previous-outcome ([see SI](#) and Table S1 for further information regarding evidence in favor of the null hypothesis for working memory load manipulation effect).

Figure 2. (A) Outcome-irrelevant learning was estimated as a tendency to stay with the previous response key selection as a function of previous outcome (i.e., £0 vs £1), despite explicit knowledge that response keys did not predict outcomes in the task. (B) marginal means estimates for response key stay probability show a greater stay tendency after reward vs. unrewarded trial, reflecting group level outcome-irrelevant learning (error bars represent HDI_{95%}) (C) posterior distribution for the regression coefficient of previous-outcome effect on response key stay probability describing group level outcome-irrelevant learning (red/blue lines indicate median/HDI_{95%}, respectively). (D) Outcome-irrelevant learning differed substantially between individuals and was mediated by working memory capacity estimates (E) posterior distribution for the regression coefficient indicating the mediation effect (i.e., paired interaction) between previous-outcome and working memory capacity.

Discussion

To engage in goal-directed behavior, human agents must form action-outcome associations within the environment ([Gershman et al., 2015](#)). Reinforcement learning studies have described, in depth, the mechanisms involved in assigning credit to preceding actions when inferring causal associations ([Sutton & Barto, 1998](#)). However, little is known regarding how cognitive control processes might prevent individuals from assigning credit to aspects of the environment that are known – with high certainty – to be irrelevant to the observed outcome ([Rmus, McDougle, et al., 2021](#)). To the best of our knowledge, this preregistered study, was the first to demonstrate a moderating role for a well-known cognitive control system, namely working memory capacity ([Cowan, 2017](#)), on the tendency to assign value to irrelevant features of the environment ([Rmus, Zou, et al., 2021](#); [Shahar et al., 2019, 2021](#)).

Our results replicated, at the group level, a newly described phenomenon coined “outcome-irrelevant learning,” which refers to a tendency to assign credit to outcome-irrelevant elements of a task ([Rmus, Zou, et al., 2021](#); [Shahar et al., 2019, 2021](#)). Specifically, in the current study, participants completed a reinforcement learning task in which they were asked to select cards to gain rewards. Card selection was indicated by a keypress; importantly, only the cards’ visual features, but not the response key used for selection, predicted monetary outcome. Outcome-irrelevant learning was illustrated by an increased probability to stay with a response key selection when it was followed by a reward compared to an unrewarded outcome. Note that this result does not reflect a general tendency to repeat one’s previous actions, and thus should not be considered as an example of a perseveration tendency ([Miller et al., 2019](#); [Wood & Rünger, 2016](#)). Finally, an important aspect of outcome-irrelevant learning is that the effect was demonstrated despite both explicit instructions and practice trials indicating that the response keys were irrelevant to the monetary outcome. Therefore, unlike previous reinforcement learning studies that focused on an ability to learn which features are relevant to an outcome in an unknown environment ([Leong et al., 2017](#); [Niv, 2019](#); [Niv et al., 2015](#); [Wilson & Niv, 2012](#)), outcome-irrelevant learning reflects a fundamental human difficulty to avoid assigning credit to features that are known, with high certainty, as not been associated with a certain outcome ([Rmus, Zou, et al., 2021](#); [Shahar et al., 2019, 2021](#)).

An important, and largely unanswered, question regards the underlying cognitive mechanisms of outcome-irrelevant learning. We argue that outcome-irrelevant learning reflects a control failure, whereby human agents find it difficult to avoid assigning credit to representations that are actively held in working memory. ([Dreisbach & Fröber, 2019](#); [Dreisbach & Haider, 2008, 2009](#)). In the current study, individuals were faced with a duality, such that response keys are both task-relevant (thus, must be actively held in mind) and outcome-irrelevant (thus, should be disregarded and not be assigned value regardless of the outcome). One hypothesis for the mechanism allowing individuals to modulate outcome-irrelevant learning might be that control is exerted during value updating, thereby resulting in task-relevant, yet outcome-irrelevant, information to maintain a neutral value. Another possibility might be that response keys are assigned a value, but that cognitive control processes regulate their involvement during the decision-making process. Both explanations lead to the similar conclusion that cognitive resources are required for the mental shielding of outcome-relevant from irrelevant information during value updating. If outcome-irrelevant learning does indeed reflect a cognitive control failure, it should increase when attentional control resources are low.

Our main finding indicated that individuals with low, as compared to high, working memory processing capacity showed increased outcome-irrelevant learning. Working memory capacity is a well-studied phenomenon in the cognitive control literature, and it has been shown to predict individuals' ability to avoid distraction from irrelevant information ([Allon & Luria, 2017](#); [Balaban et al., 2019](#); [Engle, 2001](#); [Luck & Vogel, 1997](#)). Individual differences in working memory capacity were previously found to predict interference from irrelevant information across multiple tasks such as the Stroop, dichotic listening, and anti-saccade tasks ([Conway et al., 2001](#); [Kane et al., 2001](#); [Kane & Engle, 2003](#)). Our main finding extends this literature by demonstrating that working memory capacity also predicts value-based attentional control. This finding, although correlative in nature, is in line with the speculation that lower working memory capacity leads to difficulty in holding distinct representations of the task set, which consequently leads value to be assigned to representations that are activated at the time of selection. It should be noted that, to date, most reinforcement learning studies have not directly addressed the influence of working memory capacity on credit assignment. Studies that have explored the role of working memory in credit assignment mostly manipulated set-size (e.g., number of arms in a multi-armed bandit task), thus primarily demonstrating an effect of decay and noise on the representations of relevant information ([Collins, Albrecht, et al., 2017](#); [Collins, Ciullo, et al., 2017](#); [Collins & Frank, 2012](#); [Yoo & Collins, 2021](#)). Our findings suggest that working memory capacity, as measured by a visual attention task, can predict the extent of inaccurate credit assignment.

The finding that working memory capacity, but not load manipulation, was associated with outcome-irrelevant learning should be discussed. It leads us to assume that storage of values (reinforcement learning task components) and visual stimuli (working memory task component) did not tap into similar processes, since increasing demands on one did not affect the other. However, our results clearly showed that working memory capacity limitations (as measured in a separate task, with only a visual working memory component) predicted

outcome-irrelevant learning. It is therefore plausible that working memory capacity estimates reflect a general cognitive ability to filter out irrelevant information. Alternatively, more specific explanations may apply. Studies have repeatedly shown that loading working memory beyond its capacity leads to control failures and mental interference from task-irrelevant information ([Desimone & Duncan, 1995](#); [Woodman & Luck, 2004](#)). One reason that we did not find an effect of working memory load manipulation on outcome-irrelevant learning might be that we did not include a strong enough load manipulation; however, the no-load (zero squares to remember) and the high-load (four squares to remember) conditions that we used should reflect the extreme ends of the load manipulation ([Woodman & Luck, 2004](#)). Another plausible explanation might be that participants sacrificed effort on one task for the other, however, overall accuracy rates in both tasks were above chance, suggesting that participants did in fact devote attention and effort to both tasks. We, therefore, suggest a theoretical explanation for our finding, such that the visual stimuli (working memory task component) and the cards' values (reinforcement learning task component) did not tap into the same working memory buffer, thus leading to the observed null effect. In a somewhat similar manner, and in line with the activity-silent memory theory ([Lewis-Peacock et al., 2012](#); [Rose et al., 2016](#); [Stokes, 2015](#)), information that is not immediately relevant for a present task is maintained with minimal neural activity ([Stokes et al., 2013](#)). Therefore, it could be that visual information from the working memory task component was maintained in activity-silent memory during the performance of the reinforcement learning task component ([Kreithner et al., 2022](#)). Another theoretical possibility might be that credit assignment to response keys is a relatively automatic process that happens regardless of the task demand and the current working memory load.

Finally, we note that the outcome-irrelevant learning demonstrated in this study was primarily discussed in regard to the assignment of credit to response keys, but we note that it can also be interpreted in terms of value assignment to spatial location (since the card in the right/left location was always selected using a right/left response key, respectively). The present study was designed to determine the effects of working memory capacity and load manipulation on credit assignment to outcome-irrelevant features of the environment. The results of this investigation showed that individual differences in working memory capacity predict outcome-irrelevant learning. We conclude that resource-demanding cognitive control is exerted during credit assignment, in such a way that credit can be accurately assigned to outcome-relevant actions. These findings contribute to our understanding of how credit assignment is regulated and provide a basis for future cognitive and clinical research.

References

- [Allon, A. S., & Luria, R. \(2017\). Compensation mechanisms that improve distractor filtering are short-lived. *Cognition*, 164, 74–86.](#)
<https://doi.org/10.1016/j.cognition.2017.03.020>
- [Balaban, H., Fukuda, K., & Luria, R. \(2019\). What can half a million change detection trials tell us about visual working memory? *Cognition*, 191, 103984.](#)
<https://doi.org/10.1016/j.cognition.2019.05.021>
- [Bergh, D. van den, Clyde, M. A., Gupta, A. R. K. N., de Jong, T., Gronau, Q. F., Marsman, M., Ly, A., & Wagenmakers, E.-J. \(2021\). A tutorial on Bayesian multi-model linear regression with BAS and JASP. *Behavior Research Methods*, 53\(6\), 2351–2371.](#)
<https://doi.org/10.3758/s13428-021-01552-2>
- [Bürkner, P.-C. \(2018\). Advanced Bayesian multilevel modeling with the R package brms. *The R Journal*, 10\(1\), 395.](#)
<https://doi.org/10.32614/RJ-2018-017>
- [Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. \(2017\). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76, 1–32.](#)
<https://doi.org/10.18637/jss.v076.i01>
- [Collins, A. G. \(2018\). The tortoise and the hare: Interactions between reinforcement learning and working memory. *Journal of Cognitive Neuroscience*, 30\(10\), 1422–1432.](#)
https://doi.org/10.1162/jocn_a_01238
- [Collins, A. G. E., Albrecht, M. A., Waltz, J. A., Gold, J. M., & Frank, M. J. \(2017\).](#)

[Interactions among working memory, reinforcement learning, and effort in value-based choice: A new paradigm and selective deficits in schizophrenia. *Biological Psychiatry*, 82\(6\), 431–439. <https://doi.org/10.1016/j.biopsych.2017.05.017>](#)

[Collins, A. G. E., Ciullo, B., Frank, M. J., & Badre, D. \(2017\). Working memory load strengthens reward prediction errors. *Journal of Neuroscience*, 37\(16\), 4332–4342. <https://doi.org/10.1523/JNEUROSCI.2700-16.2017>](#)

[Collins, A. G. E., & Frank, M. J. \(2012\). How much of reinforcement learning is working memory, not reinforcement learning? A behavioral, computational, and neurogenetic analysis. *European Journal of Neuroscience*, 35\(7\), 1024–1035. <https://doi.org/10.1111/j.1460-9568.2011.07980.x>](#)

[Collins, A. G. E., & Shenhav, A. \(2022\). Advances in modeling learning and decision-making in neuroscience. *Neuropsychopharmacology*, 47\(1\), 104–118. <https://doi.org/10.1038/s41386-021-01126-y>](#)

[Conway, A. R. A., Cowan, N., & Bunting, M. F. \(2001\). The cocktail party phenomenon revisited: The importance of working memory capacity. *Psychonomic Bulletin & Review*, 8\(2\), 331–335. <https://doi.org/10.3758/BF03196169>](#)

[Conway, A. R. A., Kane, M. J., & Engle, R. W. \(2003\). Working memory capacity and its relation to general intelligence. *Trends in Cognitive Sciences*, 7\(12\), 547–552. <https://doi.org/10.1016/j.tics.2003.10.005>](#)

[Cowan, N. \(2001\). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24\(1\), 87–114. <https://doi.org/10.1017/S0140525X01003922>](#)

Cowan, N. (2010). The magical mystery four: How is working memory capacity limited, and why? *Current Directions in Psychological Science*, 19(1), 51–57.

<https://doi.org/10.1177/0963721409359277>

Cowan, N. (2017). The many faces of working memory and short-term storage. *Psychonomic Bulletin & Review*, 24(4), 1158–1170. <https://doi.org/10.3758/s13423-016-1191-6>

Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6), 1204–1215. <https://doi.org/10.1016/j.neuron.2011.02.027>

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193–222.
<https://doi.org/10.1146/annurev.ne.18.030195.001205>

Dolan, R. J., & Dayan, P. (2013). Goals and habits in the brain. *Neuron*, 80(2), 312–325.
<https://doi.org/10.1016/j.neuron.2013.09.007>

Dreisbach, G., & Fröber, K. (2019). On how to be flexible (or not): Modulation of the stability-flexibility balance. *Current Directions in Psychological Science*, 28(1), 3–9.
<https://doi.org/10.1177/0963721418800030>

Dreisbach, G., & Haider, H. (2008). That's what task sets are for: Shielding against irrelevant information. *Psychological Research*, 72(4), 355–361.
<https://doi.org/10.1007/s00426-007-0131-5>

Dreisbach, G., & Haider, H. (2009). How task representations guide attention: Further evidence for the shielding function of task sets. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(2), 477–486. <https://doi.org/10.1037/a0014647>

Engle, R. W. (2001). What is working memory capacity? In *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 297–314). American Psychological Association. <https://doi.org/10.1037/10394-016>

Engle, R. W. (2002). Working memory capacity as executive attention. *Current Directions in Psychological Science*, 11(1), 19–23. <https://doi.org/10.1111/1467-8721.00160>

Engle, R. W., & Kane, M. J. (2004). Executive attention, working memory capacity, and a two-factor theory of cognitive control. In *The psychology of learning and motivation: Advances in research and theory*, Vol. 44 (pp. 145–199). Elsevier Science.

Fockert, J. W. de, Rees, G., Frith, C. D., & Lavie, N. (2001). The role of working memory in visual selective attention. *Science*. <https://doi.org/10.1126/science.1056496>

Folk, C. L., Remington, R. W., & Wright, J. H. (1994). The structure of attentional control: Contingent attentional capture by apparent motion, abrupt onset, and color. *Journal of Experimental Psychology. Human Perception and Performance*, 20(2), 317–329. <https://doi.org/10.1037//0096-1523.20.2.317>

Gershman, S. J., Norman, K. A., & Niv, Y. (2015). Discovering latent causes in reinforcement learning. *Current Opinion in Behavioral Sciences*, 5, 43–50. <https://doi.org/10.1016/j.cobeha.2015.07.007>

Kane, M. J., Bleckley, M. K., Conway, A. R., & Engle, R. W. (2001). A controlled-attention view of working-memory capacity. *Journal of Experimental Psychology. General*, 130(2), 169–183. <https://doi.org/10.1037//0096-3445.130.2.169>

Kane, M. J., & Engle, R. W. (2003). Working-memory capacity and the control of attention: The contributions of goal neglect, response competition, and task set to Stroop

interference. *Journal of Experimental Psychology. General*, 132(1), 47–70.

<https://doi.org/10.1037/0096-3445.132.1.47>

Kreither, J., Papaioannou, O., & Luck, S. J. (2022). Active working memory and simple cognitive operations. *Journal of Cognitive Neuroscience*, 34(2), 313–331.

https://doi.org/10.1162/jocn_a_01791

Leong, Y. C., Radulescu, A., Daniel, R., DeWoskin, V., & Niv, Y. (2017). Dynamic interaction between reinforcement learning and attention in multidimensional environments. *Neuron*, 93(2), 451–463. <https://doi.org/10.1016/j.neuron.2016.12.040>

Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2012). Neural evidence for a distinction between short-term memory and the focus of attention. *Journal of Cognitive Neuroscience*, 24(1), 61–79. https://doi.org/10.1162/jocn_a_00140

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279–281. <https://doi.org/10.1038/36846>

Luria, R., & Vogel, E. K. (2011). Shape and color conjunction stimuli are represented as bound objects in visual working memory. *Neuropsychologia*, 49(6), 1632–1639. <https://doi.org/10.1016/j.neuropsychologia.2010.11.031>

Makowski, D., Ben-Shachar, M. S., Chen, S. H. A., & Lüdecke, D. (2019). Indices of Effect Existence and Significance in the Bayesian Framework. *Frontiers in Psychology*, 10, 2767. <https://doi.org/10.3389/fpsyg.2019.02767>

Makowski, D., Ben-Shachar, M. S., & Lüdecke, D. (2019). bayestestR: Describing effects and their uncertainty, existence and significance within the Bayesian framework. *Journal of Open Source Software*, 4(40), 1541.

Miller, K. J., Shenhav, A., & Ludvig, E. A. (2019). Habits without values. *Psychological Review*, 126(2), 292–311. <https://doi.org/10.1037/rev0000120>

Niv, Y. (2019). Learning task-state representations. *Nature Neuroscience*, 22(10), 1544–1553. <https://doi.org/10.1038/s41593-019-0470-8>

Niv, Y., Daniel, R., Geana, A., Gershman, S. J., Leong, Y. C., Radulescu, A., & Wilson, R. C. (2015). Reinforcement learning in multidimensional environments relies on attention mechanisms. *Journal of Neuroscience*, 35(21), 8145–8157. <https://doi.org/10.1523/JNEUROSCI.2978-14.2015>

Pashler, H. (1988). Familiarity and visual change detection. *Perception & Psychophysics*, 44(4). <https://doi.org/10.3758/BF03210419>

Rmus, M., McDougale, S. D., & Collins, A. G. (2021). The role of executive function in shaping reinforcement learning. *Current Opinion in Behavioral Sciences*, 38, 66–73. <https://doi.org/10.1016/j.cobeha.2020.10.003>

Rmus, M., Zou, A., & Collins, A. G. E. (2021). Plucking a string or playing a G? Degree of choice abstraction impacts human reinforcement learning. *bioRxiv*. <https://doi.org/10.1101/2021.08.25.457707>

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

Shahar, N., Hauser, T. U., Moran, R., Moutoussis, M., Bullmore, E. T., & Dolan, R. J. (2021). Assigning the right credit to the wrong action: Compulsivity in the general population is associated with augmented outcome-irrelevant value-based learning.

[Translational Psychiatry, 11\(1\), 1–9. https://doi.org/10.1038/s41398-021-01642-x](https://doi.org/10.1038/s41398-021-01642-x)

[Shahar, N., Moran, R., Hauser, T. U., Kievit, R. A., McNamee, D., Moutoussis, M., Consortium, N., & Dolan, R. J. \(2019\). Credit assignment to state-independent task representations and its relationship with model-based decision making. *Proceedings of the National Academy of Sciences, 116*\(32\), 15871–15876. https://doi.org/10.1073/pnas.1821647116](https://doi.org/10.1073/pnas.1821647116)

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

[Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., & Duncan, J. \(2013\). Dynamic coding for cognitive control in prefrontal cortex. *Neuron, 78*\(2\), 364–375. https://doi.org/10.1016/j.neuron.2013.01.039](https://doi.org/10.1016/j.neuron.2013.01.039)

[Sutton, R. S., & Barto, A. G. \(1998\). *Introduction to reinforcement learning* \(Vol. 135\). MIT press Cambridge.](#)

[Theeuwes, J., & Burger, R. \(1998\). Attentional control during visual search: The effect of irrelevant singletons. *Journal of Experimental Psychology: Human Perception and Performance, 24*\(5\), 1342–1353. https://doi.org/10.1037/0096-1523.24.5.1342](https://doi.org/10.1037/0096-1523.24.5.1342)

[Vehtari, A., Gelman, A., & Gabry, J. \(2017\). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and computing, 27*\(5\), 1413–1432. https://doi.org/10.1007/s11222-016-9696-4](https://doi.org/10.1007/s11222-016-9696-4)

[Vogel, E. K., McCollough, A. W., & Machizawa, M. G. \(2005\). Neural measures reveal individual differences in controlling access to working memory. *Nature, 438*\(7067\),](#)

500–503. <https://doi.org/10.1038/nature04171>

Wilson, R. C., & Niv, Y. (2012). Inferring relevance in a changing world. *Frontiers in Human Neuroscience*, 5, 189. <https://doi.org/10.3389/fnhum.2011.00189>

Wood, W., & Rünger, D. (2016). Psychology of habit. *Annual Review of Psychology*, 67(1), 289–314. <https://doi.org/10.1146/annurev-psych-122414-033417>

Woodman, G. F., & Luck, S. J. (2004). Visual search is slowed when visuospatial working memory is occupied. *Psychonomic Bulletin & Review*, 11(2), 269–274. <https://doi.org/10.3758/BF03196569>

Yoo, A., & Collins, A. (2021). How working memory and reinforcement learning are intertwined: A cognitive, neural, and computational perspective. *Journal of Cognitive Neuroscience*, 1–17. https://doi.org/10.1162/jocn_a_01808

Zanto, T. P., & Gazzaley, A. (2009). Neural suppression of irrelevant information underlies optimal working memory performance. *Journal of Neuroscience*, 29(10), 3059–3066. <https://doi.org/10.1523/JNEUROSCI.4621-08.2009>