



## OPEN Memory reactivation during rest forms shortcuts in a cognitive map

Cal M. Shearer<sup>1,2</sup>, Annalise B. Rawson<sup>2,3</sup>, Helen C. Barron<sup>2,3,4</sup>✉ & Jill X. O'Reilly<sup>1,4</sup>✉

Efficient and flexible cognition relies upon cognitive maps—representations of concepts and the relations between them. Cognitive maps integrate relations that were learned separately into a cohesive whole. Memory reactivation during rest and sleep may contribute to cognitive map formation in two ways: by simply strengthening memories for directly experienced relations, or by reorganising concepts and creating new relations that capture the underlying structure. We designed a multi-stage learning task to test whether reactivation during rest is involved in restructuring memories as opposed to simply consolidating what was experienced. We causally manipulated memory reactivation during rest using awake, contextual targeted memory reactivation. We found that promoting memory reactivation during rest qualitatively reorganises the cognitive map by *forming 'shortcuts' between events which have not been experienced together*. These shortcuts in memory extend beyond direct experience to facilitate our ability to make novel inferences. Using a series of control tests we show that inference performance cannot be explained by quantitative strengthening of the experienced component links. Interestingly, we show that representing a shortcut may come with limitations, as shortcuts cannot be readily updated in response to rapid changes in the environment. Together, these findings reveal how memories are reorganised during awake rest to construct a cognitive map of our environment, while highlighting the constraints set by a trade-off between efficient and flexible behaviour.

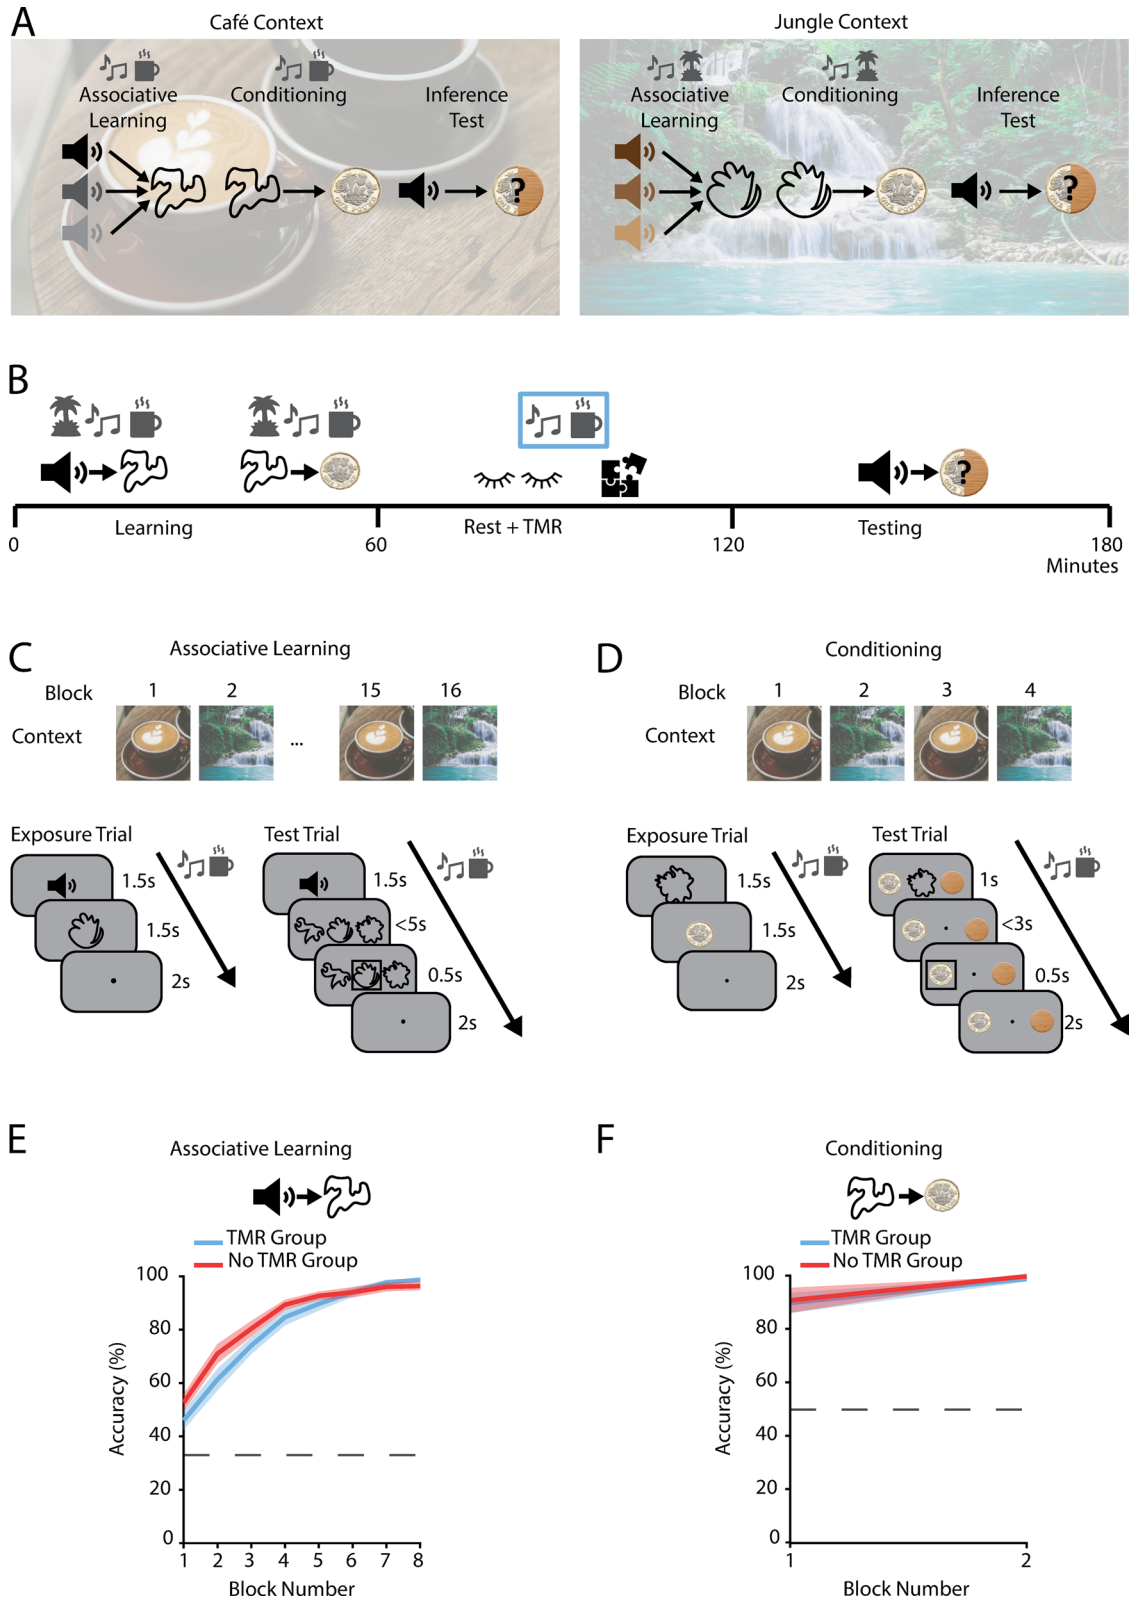
It is proposed that we understand our environment by constructing a cognitive map<sup>1</sup>. Within a cognitive map we can represent events in the external world and the relations between them. Importantly, the cognitive map is thought to go beyond the sum of its parts, by 'knitting together' events and relations that have been experienced separately and by building entirely new relations that have not been directly experienced. This allows a cognitive map to not simply aid navigation along familiar paths, but can also support navigation into unknown territories, whether spatial or more abstract.

As a simple example, if an observer has seen the pairwise associations  $A \rightarrow B$  and  $B \rightarrow C$ , they can infer the *unobserved* association  $A \rightarrow C$ <sup>2</sup>. There are two ways in which the inferred  $A \rightarrow C$  relationship could be represented: as a chain of learned associations ( $A \rightarrow B \rightarrow C$ ), or by formation of a new 'shortcut' or direct relation ( $A \rightarrow C$ ). The formation of a new shortcut ( $A \rightarrow C$ ) would represent a *qualitative* restructuring of memories, to build an entirely new relation that has not been directly experienced. The ability to make such qualitative changes in the organisation of knowledge is considered an essential feature for a cognitive map that supports flexible behaviour<sup>3–6</sup>. Here we investigate the hypothesis that offline periods contribute to the qualitative restructuring of memories. Specifically, we predict that while associative memories are formed during active learning, memory reactivation of these associations during periods of rest and sleep may introduce 'shortcuts' into the cognitive map, to support flexible behaviour.

The process of constructing a cognitive map is thought to be facilitated by offline periods of rest/sleep<sup>7–14</sup>. During these periods, memories for previous experiences are reactivated or 'replayed'<sup>15–17</sup>. Replay involves temporally structured spiking activity that recapitulates previous waking experience<sup>18–20</sup> to facilitate memory consolidation<sup>15–17,21</sup> and subsequent decision-making<sup>22</sup>. Importantly, growing evidence suggests replay may play a particular role in extending the cognitive map beyond previous waking experience to anticipate upcoming events<sup>23,24</sup>, restructure knowledge<sup>25</sup>, or even "join-the-dots" between spatial trajectories or events that were not directly experienced together<sup>26–28</sup>. Here, we explicitly test the hypothesis that offline periods of rest *qualitatively* reorganise memories, rather than *quantitatively* changing the strength of associations that were actually experienced. We define qualitative re-organisation as the creation of new inferred links (or shortcuts) that are independent of experienced associations, thus forming a cognitive map that is more than the sum of its parts.

<sup>1</sup>Department of Experimental Psychology, University of Oxford, Oxford, UK. <sup>2</sup>Medical Research Council Brain Network Dynamics Unit, Nuffield Department for Clinical Neurosciences, University of Oxford, Oxford, UK.

<sup>3</sup>The Oxford Centre for Integrative Neuroimaging, FMRIB, John Radcliffe Hospital, University of Oxford, Oxford, UK. <sup>4</sup>Helen C. Barron and Jill X. O'Reilly contributed equally to this work. ✉email: helen.barron@ndcn.ox.ac.uk; jill.oreilly@psy.ox.ac.uk



To test our hypothesis, we designed a multi-stage inference task adapted from a protocol previously implemented in humans and mice<sup>26</sup>. We causally manipulated memory reactivation during awake rest using contextual auditory TMR<sup>29,30</sup>. Consistent with our predictions, we demonstrate that biasing memory reactivation during awake rest improves the ability to infer novel relationships between sensory cues. Furthermore, using behavioural and physiological (eye-tracking) data, we show that this improvement in inference is driven by the formation of novel shortcuts between indirectly linked cues. As a result of these shortcuts, participants are less reliant on memory of intermediary associations (A → B and B → C) since they can instead use the direct, yet unobserved, shortcut formed during awake rest (A → C) to inform their decisions. However, although beneficial

◀ **Fig. 1.** Inference task design and learning performance. **(A)** Three-stage inference task adapted from a previous protocol<sup>22</sup>. In learning phase 1, participants learnt to associate auditory cues with visual cues ('Associative Learning'), with a many-to-one mapping. In learning phase 2, participants then learnt to associate visual cues with an outcome (rewarded—pound coin; neutral—wood coin) ('Conditioning'). During the 'Inference Test', the auditory cues were presented in isolation, and we assessed participants' ability to infer the value of any given auditory cue. Cues were divided into two sets, each grouped according to the contextual background music that played throughout learning (set 1: café context—left; set 2: jungle context—right). **(B)** Timeline of the task design in minutes. The task was split into three phases—learning, awake rest, and testing. During each block of the learning phase, the contextual background music for the relevant context was played. During the awake rest period, the contextual background music for only one set of cues was played to bias memory reactivation towards cues in this set (TMR group). After the awake rest period, participants were tested on their ability to infer associations between the auditory cues and the outcome cues. **(C, D)** Structure of the associative learning phase **(C)** and conditioning phase **(D)**. Top: learning trials were blocked by context and alternated between the two contexts using the contextual background music. Bottom: Example trials (schematic). Blocks of trials consisted of both exposure trials (left) and test trials (right). **(E, F)** Behavioural performance during the associative learning **(E)** and conditioning **(F)** phases (mean accuracy  $\pm$  SEM). Participants only proceeded to the inference test if they reached criterion ( $>85\%$  performance accuracy for each group in both associative learning and conditioning). There was no significant difference in the maximum accuracy reached between sensory cues in the two contexts (café vs jungle;  $p=0.385$ , paired bootstrapped test, two-tailed) or between the group of cues that were subsequently subject to TMR and no TMR (TMR group vs no-TMR group;  $p=0.185$ , paired bootstrapped test, two-tailed). Background images by Rifqi Ramadhan on Pexels.com (jungle) and Photo by [shche\\_team](#) on [Unsplash](#) (coffee), with permission.

for inference, we show that building cognitive shortcuts may come with limitations: when the component associations are modified (e.g.  $B \rightarrow C$  changes to  $B \rightarrow C^*$ ), this change is not immediately transferred to the shortcut, leaving the old shortcut ( $A \rightarrow C$ ) to compete with the updated chain of associations ( $A \rightarrow B, B \rightarrow C^*$ ).

Together our findings show that memory reactivation enhances inferential choice through the formation of shortcuts between indirect associations. However, these shortcuts cannot be rapidly updated. We thus demonstrate that memory reactivation during awake rest contributes to constructing a cognitive map that is more than the sum of its parts, allowing for efficient behaviour that is limited in flexibility.

## Results

### Task design and learning performance

To investigate the formation of shortcuts in a cognitive map, we designed an inference task (Fig. 1A,B). The inference task was adapted from a previous inference protocol implemented in humans and mice<sup>26</sup> that leveraged a sensory preconditioning paradigm<sup>2</sup>. The inference task included three stages.

#### Learning phase

In the first learning stage ( $A \rightarrow B$ ), participants learnt to associate pairs of auditory and visual cues ('associative learning'), with a many-to-one mapping from auditory to visual cues. The presentation of task cues was blocked into two groups, distinguished by different contextual background music that played throughout learning (either café or jungle music). One of these contextual background tracks was later played during rest to bias memory reactivation to this group of cues (TMR manipulation). After training, participants' mean accuracy on the associative learning trials was 96.67% (SD = 7.41%), with no significant difference in the maximum accuracy reached between sensory cues in the two contexts (café vs jungle:  $p=0.385$ , paired bootstrapped test, two-tailed) or between the group of cues that were subsequently subject to TMR and the group of cues that were not (TMR group vs no-TMR group:  $p=0.185$ , paired bootstrapped test, two-tailed); Fig. 1E).

In the second learning stage ( $B \rightarrow C$ ) of the inference task, participants learnt to associate the visual cues with either a rewarding outcome (represented by a pound coin) or a neutral outcome (represented by a wooden coin) ('conditioning'). As in the associative learning phase, contextual background music was played throughout learning. After training, participants achieved a mean accuracy of 99.50% on the conditioning trials (SD = 2.50%) (Fig. 1F).

#### Rest and TMR

Following the second learning phase, we applied a TMR manipulation<sup>31–34</sup> to causally manipulate memory reactivation by biasing each participants resting memory reactivation towards a randomly allocated subset of cues. Participants underwent the TMR manipulation during a period of awake rest<sup>29</sup> (60 min), after learning stage 2 (i.e., after completing the associative learning and conditioning phases of the task), but before performing the inference test. Throughout this awake rest period, we played the contextual background track associated with cues in one of the two groups (i.e., either café music or jungle music, fully counterbalanced across participants) whilst participants completed a relaxing task unrelated to the learning task (namely, they completed a jigsaw puzzle). We predicted that this manipulation should bias the content of memory reactivation towards one group of cues (A, B, C) and any learned associations within that group (i.e. A-B, B-C). Moreover, we hypothesised that if memory reactivation driven by TMR also facilitates the construction of a cognitive map that is more than the sum of its parts, memory reactivation should build shortcuts between cues that have not been directly experienced together (e.g. A-C), to facilitate subsequent inference across cues in the TMR group.

### Test phase

Importantly, auditory cues were never paired with outcome cues, providing an opportunity to assess evidence for an inferred relationship between these indirectly related cues. Accordingly, participants performed an ‘inference test’. In the inference test, we presented auditory cues in isolation, without visual cues or outcome cues, and without the contextual background music. In response to these isolated auditory cues, we measured evidence for inference by asking participants which outcome was associated with each auditory cue (i.e., we tested the  $A \rightarrow C$  relationship), despite the fact that they had never directly experienced this pairing.

During the test phase, the inference task was performed twice, before and after a ‘value-flip’ manipulation. During the value-flip ( $B \rightarrow C^*$ ), a subset of the associations between visual cues and outcomes were ‘Flipped’, so that previously rewarded visual cues became neutral and vice versa. Participants learned these new associations, with the contextual background soundtrack playing, as in learning phase 2. However, this value-flip was not followed by a period of rest. Instead, participants immediately repeated the inference task.

The purpose of the value-flip manipulation was to investigate to what extent learning of the updated associations  $B \rightarrow C^*$  was transferred to the inferred shortcut  $A \rightarrow C$ . We reasoned that if participants relied upon a chain of directly observed associations  $A \rightarrow B \rightarrow C$  to support inference, learning of the new link  $B \rightarrow C^*$  would result in a new chain  $A \rightarrow B \rightarrow C^*$ . However, if participants use an inferred shortcut  $A \rightarrow C$ , formed during awake rest, this shortcut would not be immediately updated in response to learning the new  $B \rightarrow C^*$  associations. Instead,  $A \rightarrow C$  might only be updated to  $A \rightarrow C^*$  during another period of awake rest.

### TMR improves inference performance

After the period of awake rest, participants performed the inference test (Fig. 2A,B). Here, we presented the auditory cues in isolation and tested participants’ ability to infer a relationship with the corresponding outcome cue (rewarding or neutral). To assess the effect of TMR on inference, we compared participants’ accuracy in response to auditory cues in the TMR and no TMR groups.

Participants were better at inferring the correct outcome for auditory cues in the TMR group than the non-TMR group ( $p = 0.013$ , Fig. 2C—this and all subsequent  $p$  values, unless otherwise noted, from paired bootstrap tests, one-tailed, as described in *methods*). The inference task was carried out both in isolation, and in the presence of two distractor tasks (a one-back task and a semantic judgement task, see *Methods*). The distractor tasks were designed to increase the load on visual working memory or general cognitive resources respectively, to modulate task difficulty and avoid ceiling effects when inference was performed in isolation. Specifically, we predicted that including a visual working memory distractor should impair inference if participants rely on recall for the intermediate visual cue (Figure S1A). However, participants were better at inferring the correct outcome for auditory cues in the TMR group under all three variants of the inference task ( $p = 0.013$ , no distractor task;  $p = 0.031$ , visual working memory task;  $p = 0.031$ , non-memory task; Fig. S1), with no significant difference in the TMR effect between conditions ( $p = 0.237$ , no distractor task vs memory task;  $p = 0.741$ , no distractor task vs non-memory task; Fig. S1). Hence, hereafter results are pooled across the three distractor conditions. Overall, participants were better at inferring the correct outcome for auditory cues in the TMR group ( $p = 0.013$ , Fig. 2C). This result demonstrates that biasing memory reactivation during periods of rest with awake, contextual TMR facilitates inferential choice.

### TMR creates a shortcut in memory between indirectly linked cues

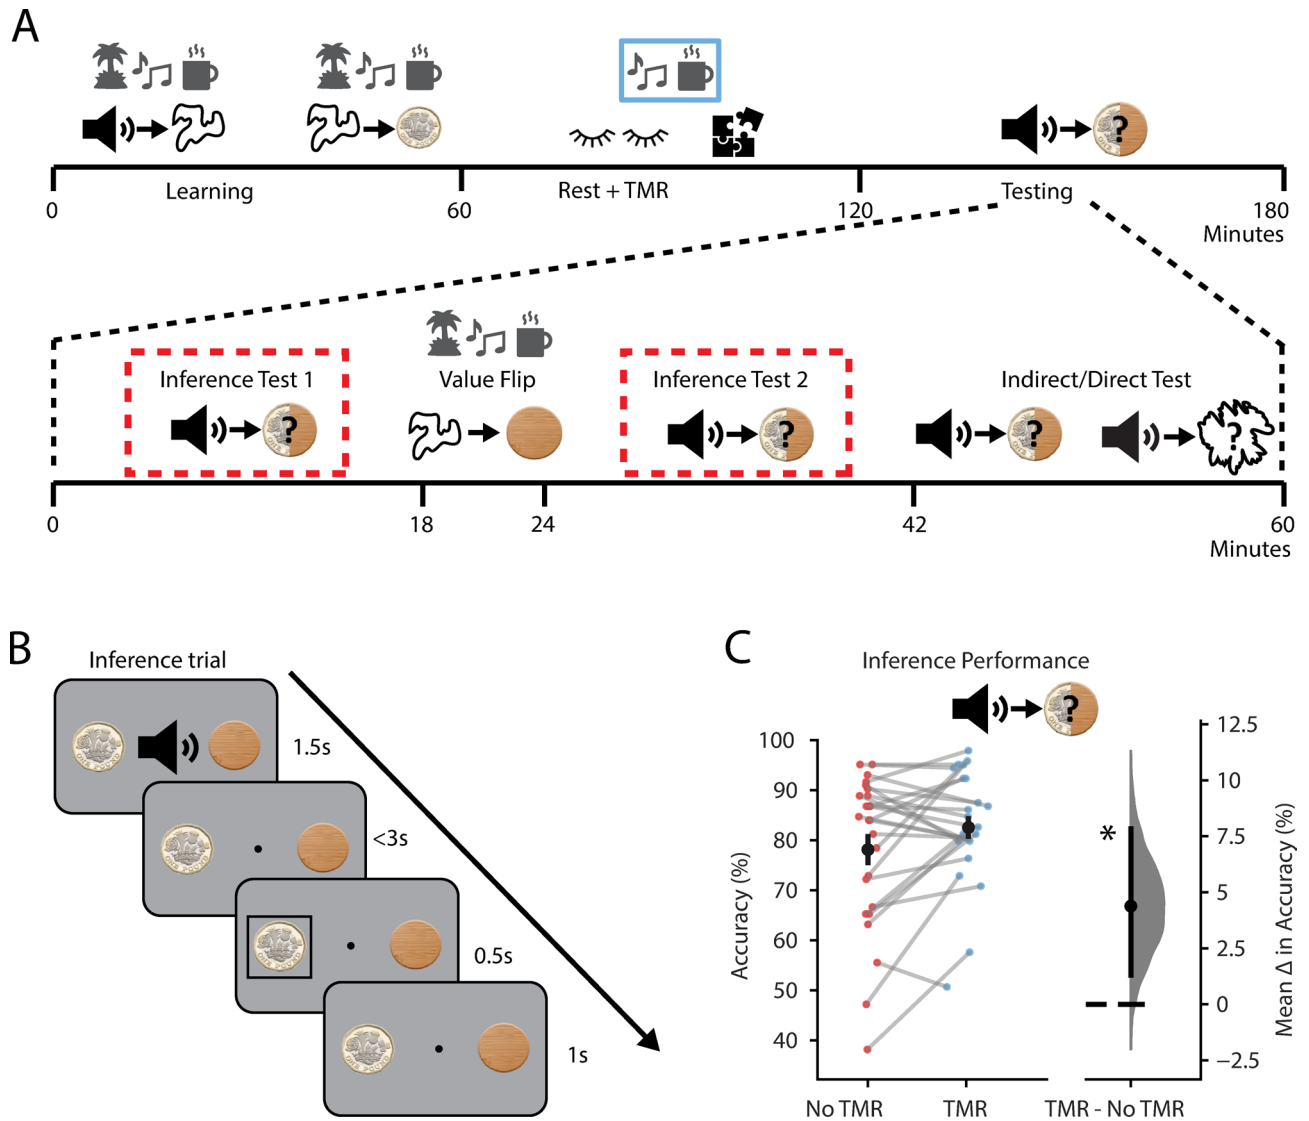
Having established that TMR facilitates inference of the unobserved  $A \rightarrow C$  link, we went on to probe the nature of that facilitation. Specifically, we asked: do participants build a shortcut ( $A \rightarrow C$ ) that is separate from the original chain of learned associations ( $A \rightarrow B$ ,  $B \rightarrow C$ ), or do they instead strengthen the component associations in the original chain ( $A \rightarrow B$ ,  $B \rightarrow C$ )? We reasoned that if TMR facilitates inference by building a new shortcut between the auditory and outcome cues ( $A \rightarrow C$ ), the resulting cognitive map may be less amenable to flexible updating when associations between cues change.

Specifically, if the visual-outcome association  $B \rightarrow C$  changes to  $B \rightarrow C^*$ , the correct chain of associations to support inference is now  $A \rightarrow B$  and  $B \rightarrow C^*$ . If participants perform inference using the chain of learned associations ( $A \rightarrow B$ ,  $B \rightarrow C$ ) then the inference  $A \rightarrow B \rightarrow C^*$  should immediately reflect any update in the component to  $B \rightarrow C^*$ . By contrast, if a direct shortcut ( $A \rightarrow C$ ) was constructed during awake rest, this would not be automatically updated to  $A \rightarrow C^*$  and behaviour may not account for the changes from  $B \rightarrow C$  to  $B \rightarrow C^*$ .

### ‘Flipping’ the component associations negates the behavioural benefit of TMR

To test this hypothesis, after the first inference test, we manipulated the association  $B \rightarrow C$  (visual-outcome), such that the outcome associated with a subset of visual cues was reversed (value-flip manipulation,  $B \rightarrow C^*$ ). Participants were exposed to this value-flip manipulation across a relatively brief 6-min task (Fig. 2A), before being asked to immediately perform the inference test once more, to test their ability to flexibly update the inferred value of auditory cues (from  $A \rightarrow C$  to  $A \rightarrow C^*$ ) (Fig. 3A). Critically, if TMR creates a shortcut ( $A \rightarrow C$ ) between auditory and outcome cues that is separate to the chain of learned auditory-visual-outcome associations ( $A \rightarrow B$ ,  $B \rightarrow C$ ), then an update in the visual-outcome leg ( $B \rightarrow C^*$ ) should not transfer to the auditory-outcome ( $A \rightarrow C$ ) shortcut. Therefore, to make correct inferences, changes to the visual-outcome ( $B \rightarrow C^*$ ) mapping would require participants to draw on the complete chain of learned associations ( $A \rightarrow B$ ,  $B \rightarrow C^*$ ) for cues in both the TMR and no TMR groups, negating the benefit of TMR. In addition, representation of an auditory-outcome ( $A \rightarrow C$ ) shortcut should compete with flexible updating in response to changes to the visual-outcome ( $B \rightarrow C^*$ ) mapping. Taken together, we predict an interaction in the effect of TMR and the value-flip manipulation on inference accuracy.

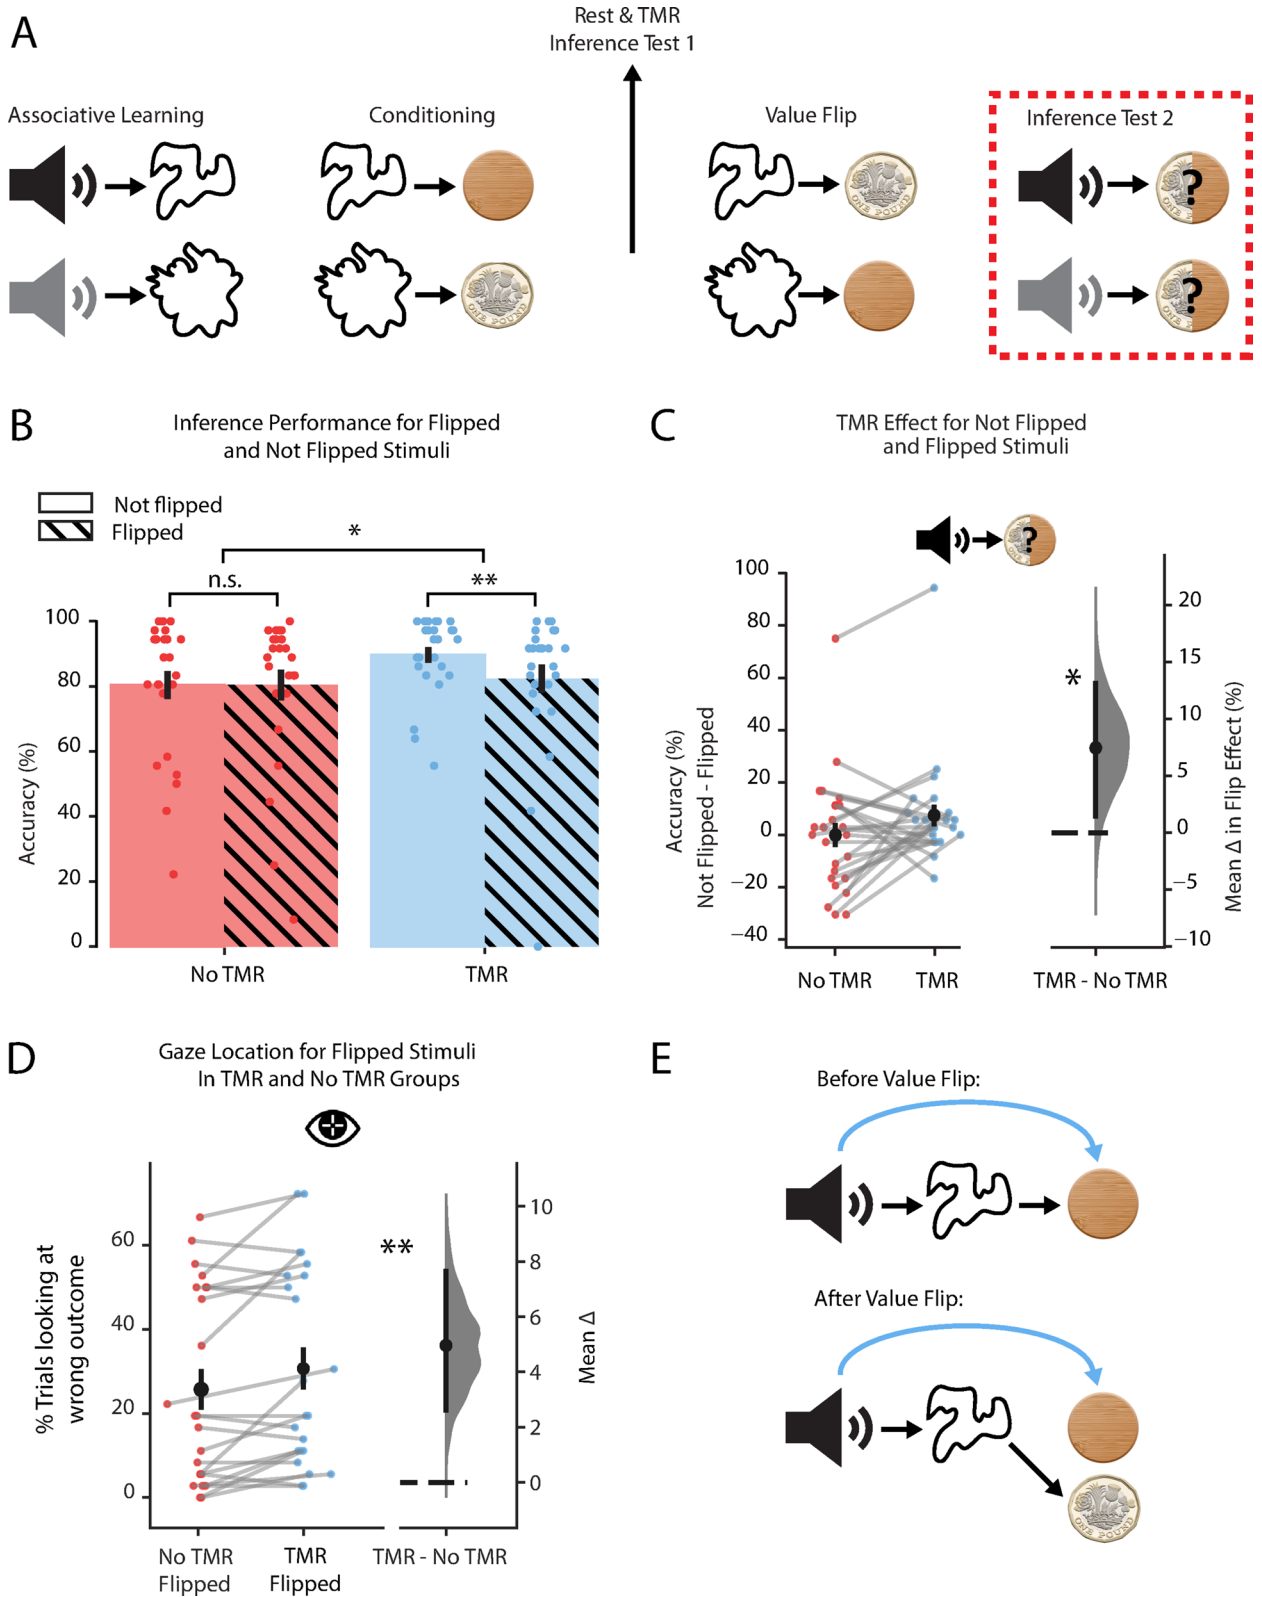
We observed evidence for a significant interaction between the effect of TMR and the value flip manipulation on inference accuracy ( $p = 0.025$ , Fig. 3B,C). We note that this interaction analysis effectively controls for the



**Fig. 2.** TMR improves inferential choice. **(A)** Structure and timings of the testing phase of the task. Tests highlighted in red were combined to examine the effect of TMR on the ability to infer an association between the auditory cues and the outcome. The remaining stages ('Value Flip' and 'Indirect/Direct Test') are detailed in Figs. 3 and 4. **(B)** Schematic: Example trial in the inference tests (I.e., 'Inference Test 1' and 'Inference Test 2'). **(C)** Inference test performance. Participants were significantly better at inferring the correct outcome for auditory cues in the TMR group ( $p=0.013$ , paired bootstrapped test, one-tailed). No significant effect was observed in reaction time data, where reaction times were not speeded. Left: raw data points for accuracy on auditory cues in the No TMR group (red; left) and TMR group (blue; right); each data point: mean accuracy for a given participant; black dot, mean; black ticks  $\pm$  SEM. Right: difference in mean inference accuracy between auditory cues in the No TMR and TMR groups shown using bootstrap-coupled estimation (DABEST) plots<sup>48</sup>. Effect size for the difference between No TMR and TMR groups was computed from 100,000 bias-corrected bootstrapped resamples<sup>52</sup>: black dot, mean; black ticks, 90% confidence interval; filled-curve, sampling-error distribution. \*indicates  $p < 0.05$ .

difference in baseline inference accuracy between cues in the TMR and no TMR groups. Any detrimental effect of value flipping will operate upon this baseline accuracy. We also note that overall performance accuracy remained high, suggesting that even when a detrimental effect of value flipping was observed, on the majority of trials participants were still able to use the complete chain of learned associations to perform inference ( $A \rightarrow B$ ,  $B \rightarrow C^*$ ).

Using a series of post-hoc paired tests, we further demonstrate that the direction of the interaction between TMR and the value-flip manipulation was as predicted. Namely, within the TMR condition, which inherently controls for the effect of TMR on inference accuracy, a significant effect of the value-flip manipulation can be observed ( $p=0.003$ , Fig. 3B). Moreover, for the Not Flipped but not the Flipped cues, a significant effect of TMR was observed (Not Flipped:  $p=0.001$ ; Flipped:  $p=0.234$ ). Therefore, for visual-outcome mappings ( $B \rightarrow C$ ) where



there was no change in the association (Not Flipped associations), participants were better at inferring the correct outcomes for auditory cues in the TMR compared to the no TMR group ( $p=0.001$ , Fig. 3B,C). However, for flipped associations ( $B \rightarrow C^*$ ), there was no significant difference between cues in the TMR and no-TMR groups ( $p=0.234$ , Fig. 3B,C). Taken together, these results suggest that the value flip manipulation impairs inference, when controlling for the effect of TMR, and the benefit of TMR is negated by the value-flip manipulation. This is consistent with the proposal that, after rapid updating in the visual-outcome ( $B \rightarrow C^*$ ) mapping, participants had to perform inferential choice by drawing on the complete chain of learned associations ( $A \rightarrow B$ ,  $B \rightarrow C^*$ ), rather than the shortcut ( $A \rightarrow C$ ), thus negating the benefit of TMR.

◀ **Fig. 3.** TMR creates a shortcut which cannot be readily updated. **(A)** Schematic illustrating the value flipping manipulation. After the learning and awake rest phases, the outcome associated with half of the visual cues was flipped ('Value Flip'). During Inference Test 2 (highlighted in red), we tested to see if participants could correctly infer the new outcome that was indirectly associated with the auditory cue. **(B)** Inference test performance after the value flip manipulation. A significant interaction was observed in inference accuracy in response to auditory cues in the Flipped vs. Not Flipped condition for the TMR vs. no TMR group ( $p=0.025$ , paired bootstrap test, one-tailed). The predicted directionality of this effect was confirmed using post-hoc tests. Within the TMR condition, a significant difference in inference accuracy was observed between the Flipped and Not Flipped cues ( $p=0.003$ , post-hoc paired bootstrap test, one-tailed), while no significant difference was observed within the no TMR condition ( $p=0.490$ , post-hoc paired bootstrap test, one-tailed). Within the Not Flipped condition, but not in the Flipped condition, a significant difference in inference accuracy was observed for the TMR vs. no TMR cues (Not Flipped:  $p=0.001$ ; Flipped:  $p=0.234$ ; post-hoc paired bootstrap test, one-tailed). No significant effects were observed in reaction time data (reaction times were not speeded). Bars: mean percentage accuracy for No TMR (red; left two bars) and TMR (blue; right two bars) groups split by whether the visual cue-outcome association was Not Flipped (plain; left bar of each group) or Flipped (hatched; right bar of each group); black ticks:  $\pm$  SEM; each data point: mean accuracy for one participant. **(C)** As shown in B, a significant interaction was observed in inference accuracy in response to auditory cues in the Flipped vs. Not Flipped condition for the TMR vs. no TMR group ( $p=0.025$ , paired bootstrap test, one-tailed). No significant effect was observed in reaction time data (reaction times were not speeded). Left: raw data points for No TMR cues (left) and TMR cues (right); each data point: mean effect of flipping (Not Flipped–Flipped accuracy) for a given participant; black dots, mean; black ticks  $\pm$  SEM. Right: difference in means between TMR and No TMR cues shown using bootstrap-coupled estimation (DABEST) plots as in Fig. 2C: black dot, mean; black ticks, 90% confidence interval; filled-curve, sampling-error distribution. **(D)** For auditory cues where the visual cue-outcome association was flipped, gaze data revealed a significant difference in percentage of trials spent looking at the wrong outcome for auditory cues in the TMR compared to the No TMR group ( $p=0.002$ , paired bootstrap test, one-tailed). Analysis applied to correct trials only. Left: raw data points for No TMR group (red; left) and TMR group (blue; right); each data point: mean percentage of trials looking at wrong outcome for a given participant; black dots, mean; black ticks  $\pm$  SEM. Right: difference in means between No TMR and TMR groups shown using bootstrap-coupled estimation (DABEST) plots as in Fig. 2C: black dot, mean; black ticks, 90% confidence interval; filled-curve, sampling-error distribution. See Figure S3 for equivalent plot for Not Flipped cues. **(E)** Schematic showing how a shortcut may conflict with a chain of direct links after the value-flip manipulation. Top: before value-flip, both the chain of direct links (black arrows) and the shortcut formed by TMR (blue curved arrow) lead to the same outcome. Bottom: After value-flip, the chain of direct links has been updated but now conflicts with the shortcut which cannot be rapidly updated. \*indicates  $p < 0.05$ ; \*\* indicates  $p < 0.01$ ; n.s. indicates no significant difference.

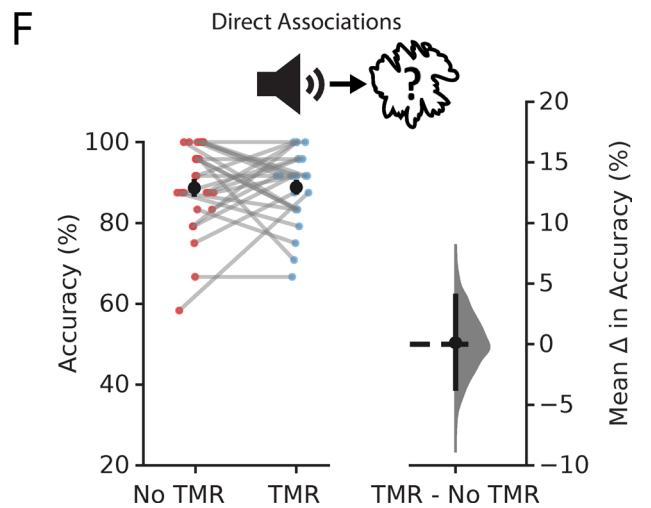
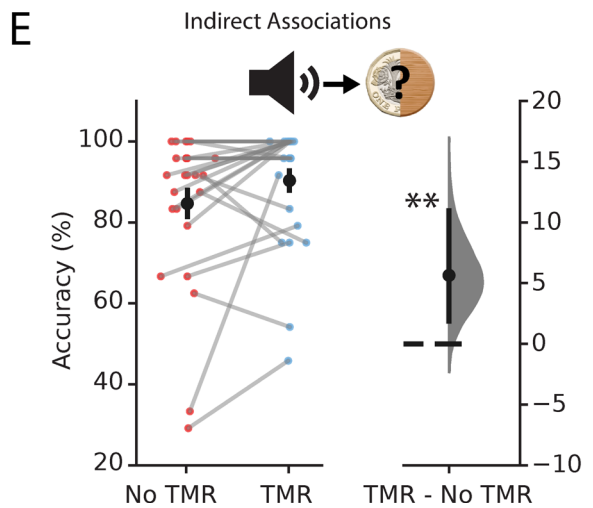
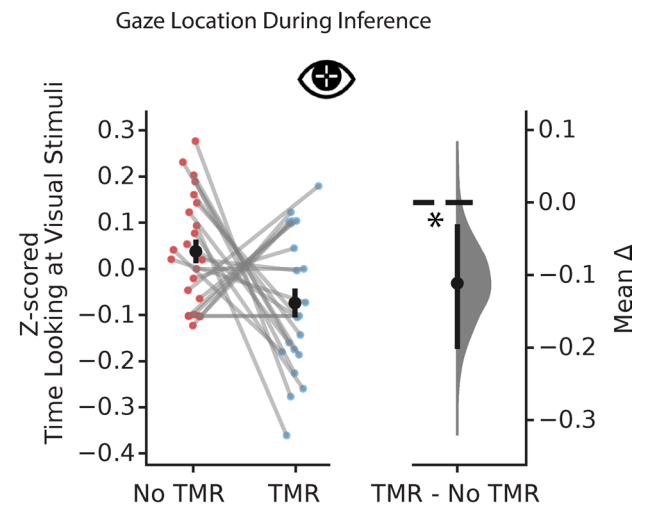


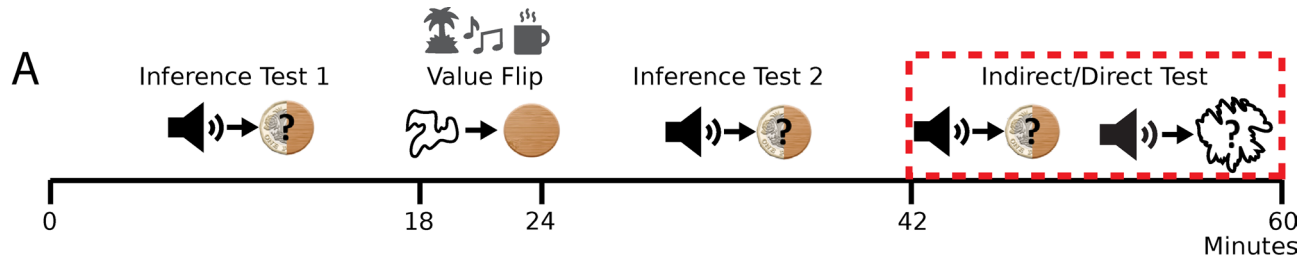
### Gaze data suggest that the outdated shortcut competes with the flipped association

Next, we reasoned that if the outdated shortcut  $A \rightarrow C$  continues to be represented, it might in fact compete for behavioural control with the modified chain of learned associations ( $A \rightarrow B, B \rightarrow C^*$ ). Such competition would strengthen the conclusion that the inferred shortcut  $A \rightarrow C$  exists as a separate entity to the chain of observed associations (evidence that rest qualitatively changes the cognitive map). Therefore, during the inference test performed after the value-flip manipulation, we used eye-tracking data to ask whether forming a shortcut ( $A \rightarrow C$ ) biases participants to look at the previously associated (but now incorrect) outcome (C) in the inference test (Fig. 3E). To do this, we took all trials where participants performed correct inference and then calculated the percentage of these trials where participants looked at the incorrect outcome, noting that for Flipped cues the incorrect outcome corresponded to the outcome cue (C) that was previously correct prior to the value-flip manipulation.

We hypothesised that for auditory cues where the visual-outcome association was flipped ( $B \rightarrow C^*$ ), in the TMR condition participants may still retain a shortcut ( $A \rightarrow C$ ) from the auditory cue to the (now incorrect) outcome (Fig. 3E). In these cases, participants would be more likely to look at outcome C rather than  $C^*$  during the decision process. Thus, relative to auditory cues in the TMR group, we predicted that this effect should be lower, or absent for cues in the no-TMR group where the shortcut ( $A \rightarrow C$ ) was weaker or not formed. This was indeed the case. For Flipped cues, participants were significantly more likely to look at the incorrect outcome cue (C) in the TMR group compared to the no TMR ( $p=0.002$ , Fig. 3D; see Figure S3 for Not Flipped cues). This suggests that when a shortcut ( $A \rightarrow C$ ) forms due to TMR, knowledge of this shortcut competes with knowledge of relevant updated associations. This competitive effect supports the view that the shortcut ( $A \rightarrow C$ ) is separate from the chain of learned associations ( $A \rightarrow B, B \rightarrow C$ ). Consequently, TMR may be considered to support efficient behaviour by facilitating inference through the formation of a shortcut between auditory and outcome cues even when these cues have not been directly experienced together. However, building a shortcut in the underlying cognitive map appears to have limitations—the shortcut cannot be rapidly updated in response to environmental change.

### Gaze data show that participants use intermediary cues more for non-TMR associations

In a final set of analyses, we again used eye-tracking data to ask whether TMR facilitates inference by creating a shortcut in memory between the auditory and outcome cues. We reasoned that if a shortcut ( $A \rightarrow C$ ) is represented, we would expect participants to be less likely to use the intermediary visual cues (B) to help make inferential choices for cues in the TMR group. To test this, we analysed eye-tracking data acquired during correct



trials on a third inference test ('indirect') and during a memory test for learned associations ('direct') ('indirect/direct' test, Fig. 2A, 4A,B). During the first half of this testing phase, participants again had to decide whether the auditory cues were associated with the rewarding or neutral outcome (inference test, 'indirect'). Replicating our previous findings, we found that participants were again significantly more accurate at inferring the correct relationships for auditory cues in the TMR group compared to the no TMR group ( $p = 0.009$ , Fig. 4D). However, unlike in the previous inference tests, all of the intermediary visual cues were now presented on the screen (Fig. 4B). We predicted that if participants had formed a shortcut between auditory and outcome cues, they would be less likely to look at the visual cues to help them infer a value for each auditory cue. Thus, despite the value-flip manipulation, we expected cues in the TMR condition (relative to the no TMR condition) to be linked



◀ **Fig. 4.** TMR Creates a shortcut which does not rely on Directly Learnt Associations (A) Structure of the testing phase of the task. Data presented in (C–F) is from the ‘Indirect/Direct Test’ highlighted in red. In the indirect/direct test, participants were tested on both the indirect (auditory–outcome) and direct (auditory–visual) associations for each auditory cue. (B) Schematic: example indirect trial (schematic) where auditory cues were now presented while the visual cues were on screen. Participants were required to infer the outcome cue in response to the auditory cue. Direct trials followed the same structure except, after hearing the auditory cue, participants had to choose the associated visual cue rather than the associated outcome. (C) Gaze location from eye tracking data acquired during indirect trials in the indirect/direct test. Participants spent significantly less time looking at the visual cues in response to auditory cues in the TMR group ( $p=0.014$ , paired bootstrap test, one-tailed). Analysis applied to correct trials only. Top: example gaze trajectories from a single indirect trial for one participant in response to an auditory cue from the No TMR group (red; left) and an auditory cue from the TMR group (blue; right). Darker colours indicate later time in the trial. Bottom left: raw data points for No TMR group (red; left) and TMR group (blue; right); each data point: mean z-scored time spent looking at visual cues during inference (correct trials only); black dots, mean; black ticks  $\pm$  SEM. Bottom right: difference in means between No TMR and TMR groups shown using bootstrap-coupled estimation (DABEST) plots as in Fig. 2C: black dots, mean; black ticks, 90% confidence interval; filled-curve, sampling-error distribution. (D, E) The effect of TMR on indirect associations (D) and direct associations (E). There was a significant effect of TMR on accuracy for indirect associations ( $p=0.009$ , paired bootstrap test, one-tailed, for both Flipped and Not Flipped cues) but not for direct associations ( $p=0.470$  paired bootstrap test, one-tailed, for both Flipped and Not Flipped cues). No significant effect was observed in reaction time data, where reaction times were not speeded. Left of each panel: raw data points for No TMR group (red; left) and TMR group (blue; right); each data point: mean accuracy for a given participant; black dots, mean; black ticks  $\pm$  SEM. Right of each panel: difference in means between No TMR and TMR groups shown using bootstrap-coupled estimation (DABEST) plots as in Fig. 2C: black dot, mean; black ticks, 90% confidence interval; filled-curve, sampling-error distribution. \*indicates  $p<0.05$ , \*\*indicates  $p<0.01$ .

via a short-cut (A  $\rightarrow$  C) that influences gaze location, for both the Flipped and Not Flipped cues, by reducing time spent looking at the intermediary visual cue (B). Consistent with our prediction, on correct trials participants spent less time looking at the visual cues for auditory cues in the TMR group compared to the no TMR group ( $p=0.014$ , Fig. 4C). Therefore, for cues in the TMR group, there was a reduced need to look at the intermediary visual cue to help make a correct inference. Together with our previous findings, this result suggests that TMR facilitates the formation of a memory shortcut between the indirectly linked auditory and outcome cues.

### The effect of TMR on inference is not explained by improved memory for directly learnt associations

Finally, we asked whether awake contextual TMR also facilitates memory accuracy for directly associated cues. Previous work has shown that TMR protocols that use punctate auditory cues improve memory for directly learnt associations<sup>35</sup>. However, since our TMR protocol uses contextual cues, we wanted to investigate whether this protocol would have similar or different effects on memory. During the indirect/direct test (Fig. 4A,B), participants were also given a memory test for learned associations between auditory and visual cues (‘direct associations’). By contrast to performance on the inferred relationships between auditory and outcome cues (‘indirect associations’) reported above (Fig. 4D), when participants were tested on direct associations, they showed no significant difference in accuracy for auditory-visual (A  $\rightarrow$  B) cue mappings in the TMR group compared to the no TMR group ( $p=0.470$ , Fig. 4E).

Since participants were trained to ceiling on the direct associations, it is possible that we did not see an effect of TMR due to ceiling performance. To investigate this potential confound, we performed a split based on combined performance for TMR and no TMR groups, and applied this to both the indirect association (inference) test (in which we saw an effect of TMR) and direct association test (in which we did not see an effect of TMR) (Figure S2). For direct associations, even for participants who were not at ceiling, we continued to observe no significant effect of TMR ( $p=0.389$ , top 50%;  $p=0.368$ , bottom 50%), suggesting that the lack of improvement under TMR was probably not due to a ceiling effect. In contrast, for indirect associations (where we had observed an overall effect of TMR), the TMR trend remained for both the bottom and top half of participants ( $p=0.071$ , top 50%;  $p=0.034$ , bottom 50%). Together these findings suggest that the reported difference in effect of TMR on indirect and directly learned associations is likely not due to difference in task difficulty or due to a ceiling effect, but instead indicates that contextual TMR applied during awake rest prioritises formation of shortcuts that support inference, rather than strengthening learned memory content.

Together, these results suggest that while contextual TMR improves inferential choice across indirectly associated cues, this improvement cannot be explained by improved performance on directly learned associations. Therefore, the benefits of contextual TMR on inferential choice must be explained by a mechanism that goes beyond merely strengthening learned associations. This provides further evidence to suggest that periods of awake rest facilitate inference by restructuring knowledge and creating shortcuts between items in memory that extend beyond direct experience.

### Discussion

Previous studies have shown the benefit of rest/sleep for inference<sup>9–13,36,37</sup>. However, the nature of the changes that occur during rest/sleep to facilitate inference remain unclear. In this study, the central question we ask is: *How does memory reactivation reorganise knowledge during awake rest, to guide adaptive behaviour?* We sought to

tease apart two ways in which rest/sleep could affect memories: an improvement in behavioural performance on an inference task could be supported by either a quantitative strengthening of learned, component associations ( $A \rightarrow B$  and  $B \rightarrow C$ ), or a qualitative change, involving the formation of a new, 'shortcut' link ( $A \rightarrow C$ ). The latter, qualitative change would reflect an essential pre-requisite for the organization of knowledge into a structured cognitive map that is more than the sum of its parts.

Using a multi-stage inference task with awake auditory contextual TMR<sup>29</sup>, we show that memory reactivation during awake rest benefits inference by building a new shortcut (e.g.  $A \rightarrow C$ ), with no effect on directly trained associations (i.e.  $A \rightarrow B$ ). We demonstrate this result using two approaches. First, we show that TMR improves behavioural accuracy when participants are required to perform inferential choice ( $A \rightarrow C$ ). Second, using eye-tracking data we show that TMR reduces the time spent looking at the intermediary visual cues (B) during inferential choice. Together these results suggest that promoting memory reactivation during awake rest facilitates the formation of novel shortcuts in the cognitive map. This may explain why periods of rest and sleep facilitate behavioural readouts of inference, but also other cognitive processes such as insight, generalisation and abstraction<sup>9–14</sup>. Interestingly, both our behavioural and eye-tracking data show that these shortcuts cannot be as readily updated in response to subsequent changes in the environment. Rather, shortcuts (e.g.  $A \rightarrow C$ ) may limit behavioural flexibility when optimal behaviour requires participants to revert to a chain of learned associations (i.e.  $A \rightarrow B$ ,  $B \rightarrow C$ ). Together, our results reveal how memory reactivation during awake rest plays a causal role in *restructuring* memories, while highlighting how the organisation of knowledge sets an important trade-off between efficient and flexible behaviour. Thus, re-organising knowledge may improve behavioural flexibility, but only if the environment remains constant and the expected statistical relationships between events/cues are persevered.

While our ability to make novel inferences can clearly support efficient behaviour, representing inferred relationships as a shortcut ( $A \rightarrow C$ ) rather than a chain of learned associations ( $A \rightarrow B$ ,  $B \rightarrow C$ ) may limit behavioural flexibility. Our data demonstrates the limitation of building a shortcut in two key analyses. These two analyses are applied to data acquired after the value-flip manipulation, where half of the visual-outcome mappings are 'flipped' ( $B \rightarrow C$  becomes  $B \rightarrow C^*$ ). First, we show that the TMR benefit on performing inference is no longer observed for the Flipped cues. Second, for the Flipped cues, eye-tracking data shows that participants spend more trials looking at the previously correct, but now incorrect outcome ( $C^*$ ) for TMR compared to no-TMR cues. We note that after the value-flip manipulation overall performance remains high, suggesting that the full chain of learned associations ( $A \rightarrow B$ ,  $B \rightarrow C^*$ ) are still recalled on the majority of trials. Together these findings suggest the following: first, the outdated shortcut ( $A \rightarrow C$ ) may compete with the updated chain of associations ( $A \rightarrow B$ ,  $B \rightarrow C^*$ ); second, the shortcut ( $A \rightarrow C$ ) is indeed separate from the chain of associations ( $A \rightarrow B \rightarrow C/C^*$ ). Overall, these data demonstrate that shortcuts in memory (e.g.  $A \rightarrow C$ ) cannot be readily updated in response to rapid changes in the environment and may limit behavioural flexibility when optimal behaviour requires participants to revert to a chain of learned associations (i.e.  $A \rightarrow B$ ,  $B \rightarrow C^*$ ). We speculate that if a further rest period was included after the change in value of outcome (C), a new updated shortcut ( $A \rightarrow C^*$ ) would have been formed and used. Indeed, there is evidence to suggest that rest and sleep may play a role in integrating new memories into existing knowledge<sup>38–41</sup>. Our findings complement this work by suggesting that adaptive re-organisation of the cognitive map cannot always occur 'on the fly' and may instead rely on offline periods of rest or sleep. Our study therefore provides insight into how new information is incorporated into a cognitive map, to guide adaptive behaviour.

A candidate neural mechanism that may explain our results is replay. Previous studies suggest that auditory TMR leads to memory reactivation and replay in hippocampus<sup>18,30</sup>, but also in brain regions typically recruited during the initial learning phase<sup>31,33,42</sup>. Moreover, the power of particular oscillatory dynamics associated with memory consolidation are typically enhanced during TMR, along with coupling between encoding regions and hippocampus. In this study, memory reactivation in response to TMR may similarly drive memory reactivation in brain regions that represent cues in the inference task, together with hippocampal circuits. During an inferential learning task in mice, similar to the one employed here, hippocampal replay during rest/sleep was shown to go beyond direct experience by "joining-the-dots" between indirectly linked cues<sup>26</sup>. Specifically, during awake sharp-wave ripples (SWRs), hippocampal spiking activity increasingly represented the inferred relationship ( $A \rightarrow C$ ) in the absence of the intermediary cue (B). This contrasted with activity in the hippocampus at the time of inferential choice, where hippocampal activity appeared to draw directly from learned associations, representing the intermediary cue (B) but not the outcome (C)<sup>26</sup>. During inferential choice, the inferred outcome was found to instead be represented by prefrontal cortex (PFC) and putative ventral tegmental area (VTA)<sup>26</sup>. Taken together, this suggests a mechanism whereby spiking sequences in hippocampal SWRs may both strengthen directly learned component associations while also providing a training signal to other brain regions, to assign credit to cues that are indirectly associated with reward, thereby generating a short-cut.

However, in our previous work<sup>26</sup>, the behavioural consequences of this candidate mechanism were not tested. Here, in humans, we use a contextual TMR manipulation<sup>29</sup> to bias the content of memory reactivation during awake rest. This causal manipulation suggests that memory reactivation during awake rest leads participants to form shortcuts between indirectly related cues. Furthermore, these shortcuts reduce the need to rely on directly learned associations, by reducing dependence on the intermediary cue. Together these findings suggest that our TMR manipulation promotes the formation of a shortcut in the underlying cognitive map by biasing the content of hippocampal replay such that inferred relationships in the TMR group are prioritised. Yet, we speculate that the shortcut may not be represented in hippocampus, despite offline hippocampal activity likely being necessary for construction of the shortcut. Since this study did not directly measure replay, future investigations will need to validate this interpretation using neuronal recordings.

When designing the multi-stage inference task, we leveraged a sensory preconditioning paradigm<sup>2</sup>. Sensory preconditioning is a well-established behavioural model of associative inference in both humans and animals<sup>26,43–47</sup>

and causal manipulations demonstrate hippocampal dependence<sup>26</sup>. Thus, while sensory preconditioning does not provide a model for episodic memory per se, the paradigm does aim to capture relevant features of episodic memory, such as associative relationships between stimuli, while showing necessary involvement of a common hippocampal circuit. Importantly, performance on a sensory preconditioning task captures the quintessential features of behaviours that draw from a cognitive map, by testing both recall of learned associations but also inference across sensory cues that have not been experienced together. Behavioural performance may therefore be thought to draw from a cognitive map that links preconditioned cues (A) to reward (C). The task design is parsimonious given that inference is performed across groups of 3 cues, arguably the minimum number of cues that allow for inference. To further test the effect of TMR on both inference and value updating with sufficient controls in place, we included 12 sets of 3 cues. Overall, our task design ensures participants show reliable and efficient learning, while minimising unwanted effects attributed to memory load or forgetting.

It was beyond the scope of this study to test the effect of memory reactivation on formation of larger cognitive maps or maps that include sensory cues deriving from alternative sensory modalities. However, we predict that our findings should generalise to inferences that occur across longer sequences of cues, including those formed from cues that derive from modalities that differ to those included here (i.e. auditory, visual, monetary). We note that electrophysiological data from V4 in non-human primates suggests that the temporal position within a 7-element visual sequence influences the extent of memory reactivation, where more robust memory reactivation emerges for cues at the start and end of the sequence<sup>48</sup>. One possibility is that these temporal-position effects can be explained by formation of short-cuts in the hippocampal cognitive map. To address this knowledge gap there is a need for future studies to record from both hippocampus and sensory cortices on tasks that include longer sequences of cues.

The specifics of our TMR manipulation must be considered when comparing our findings with that of others. Previous studies demonstrate that TMR during periods of post-encoding sleep elicits a robust effect on memory retention (for example<sup>30,35,49–55</sup>). More variable results are reported when using TMR during periods of awake rest<sup>29,30,35,49,50,56–60</sup>, which may in part be explained by variation in the level of task engagement during the rest period<sup>50,58</sup>. Specifically, application of TMR in the awake state appears to show no benefit for memory when cueing is coupled with externally-orientated, attention demanding tasks<sup>35,50</sup>, and appears to prioritize weakly learned information<sup>50,58,61</sup>. Our results are congruent with these previous findings, in suggesting that strongly learned direct associations do not reliably benefit from memory reactivation in awake rest, while weaker non-directly trained associations that extend beyond direct experience are facilitated by TMR<sup>29</sup>.

Moreover, we note that our TMR protocol differs from the majority of TMR manipulations by including a contextual rather than punctate auditory cues to trigger memory reactivation. Punctate auditory cues may have several drawbacks, including the need to precisely time the cue to specific phases of sleep oscillations to see the best effects<sup>62</sup>. Furthermore, since punctate cues are associated with specific stimuli, they provide a limited tool to investigate how links across multiple stimuli are consolidated. The contextual TMR protocol used here provides a unique opportunity to investigate the effect of memory reactivation on a group of cues that together form the primitive for a map of associations. As described previously<sup>29</sup>, the aim of this TMR protocol is to capitalise on the endogenous neural dynamics of reactivation within a cognitive map, without interrupting or altering the natural sequence of replay, as may occur with the introduction of punctate cues. Arguably, this approach is more comparable to TMR studies that have used contextual odour cues in sleep<sup>30,63,64</sup>, however we apply the contextual cue during awake rest. Our data replicate our previous findings: memory reactivation during awake rest improves performance on new behaviours, such as associative inference, with no effect on behaviours that draw directly from learned information<sup>29</sup>. In addition, we demonstrate *how* memory reactivation during periods of awake rest influences the organisation of memories within the cognitive map, where memory reactivation plays a causal role in forming novel links that go beyond direct experience.

Overall, our results demonstrate that knowledge is subject to qualitative restructuring during offline periods of rest/sleep. During awake rest, stimuli that are reactivated (due to contextual TMR) are organized into cognitive maps through the creation of novel shortcuts between indirectly linked cues. These shortcuts are distinct from the original chain of learned associations and the behavioural benefit of forming a shortcut cannot be explained by quantitatively strengthening the component learned associations (since there was no benefit of TMR on these learned associations even in participants whose performance was well below ceiling). Furthermore, when the component learned associations are rapidly modified in response to changes in the environment, shortcuts are not updated. Indeed, an outdated shortcut can conflict or compete with an updated chain of associations to control decision making, illustrating that the two sets of associations are separate. Together our findings provide new insight into how integrated cognitive maps are formed during periods of awake rest. Moreover, we demonstrate the effectiveness of using awake contextual TMR as a tool to reveal how the resting brain re-organises memory to shape adaptive behaviour.

## Methods

### Participants

A total of 32 healthy participants were recruited for this study. 7 participants were excluded from analysis as they failed to reach the performance criterion during the initial associative learning phase (criterion: > 85% performance accuracy for each contextual group of cues). The remaining 25 participants were included in the analysis (mean age of 24.72 ± 6.35 years (standard deviation), range 19–40, 4 males). All participants had normal or corrected-to-normal vision. All experiments were performed in accordance with the Declaration of Helsinki and all relevant guidelines and regulations, and were approved by the University of Oxford ethics committee (reference number: R43593/RE013). All participants gave informed consent prior to participating in the study.

### Experimental set-up

Cues were presented on a 24-inch screen with a spatial resolution of 1920×1080 pixels (width×height), a background luminance of 0.5 (grey), and a refresh rate of 100 Hz. The approximate distance of participants from the screen was 64cm, which meant that one degree of visual angle corresponded to 40 pixels on the screen. Stimulus presentation was controlled using Psychophysics Toolbox-3<sup>65</sup> in MATLAB (version R2015b). Eye movements were recorded in the testing phase of the experiment (after the rest period) using an eye-tracking camera (EyeLink®, SR Research) tracking both eyes at a rate of 1kHz.

### Inference task—overview

Participants performed an inference task (Fig. 1A,B). In the first stage of the task ('associative learning'), participants learnt to associate punctate auditory cues with visual cues. The punctate auditory cues were naturalistic sounds from the BBC sound library (<https://sound-effects.bbcrewind.co.uk/>) and the visual cues were unsymbolic shapes<sup>66</sup>. Each visual cue was paired with three auditory cues. In the second stage of the task ('conditioning'), these visual cues were then associated with either a pound coin (rewarded outcome) or a wood coin (neutral outcome, of no value). The cues were split into two sets, grouped according to the contextual background music that played throughout learning. For one of the two sets of cues, participants heard jungle music in the background and the punctate auditory cues were sounds you would hear in a jungle (e.g., animal noises). For the other group, participants heard café music in the background and the punctate auditory cues were sounds you would hear in a café (e.g., a cash register). In total, 12 auditory cues and 4 visual cues were included in each group and the relationship between the conditions (i.e., which context was learnt first, and which context was subject to TMR) was fully counterbalanced across participants. Since the visual cues had no direct meaning, they were randomly assigned to each group across participants.

### Behavioural protocol

The associative learning stage was designed to allow participants to learn associations between auditory and visual cues. This stage of the task included passive exposure and testing trials (Fig. 1C). In the exposure phase, on each trial a punctate auditory cue was presented for 1.5 s, immediately followed by presentation of a visual cue for 1.5 s, followed by an inter-trial interval (ITI) of 2 s. Each mini-block of exposure consisted of 12 trials (one presentation of each auditory-visual pair). Each mini-block of exposure was followed by a testing mini-block. On each trial in the testing mini-block participants heard an auditory cue for 1.5 s before being presented with 3 possible visual cues. Participants were then required to select the visual cue that was paired with the auditory cue they just heard, responding within 5 s using the keys 'j', 'k', and 'l'. Once a choice had been selected, a black rectangle would highlight this choice for 0.5 s, followed by an ITI of 2 s. At the end of each testing mini-block, participants were given feedback on their average performance. Each block in the associative learning phase was made up of one exposure mini-block followed by one testing mini-block and consisted only of cues from one of the two groups of cues, with the relevant contextual background music playing throughout the block. Blocks alternated between the two contexts with a total of 8 blocks per context. The relationship between the conditions (i.e., which context was learnt first, and which context was subject to TMR) was fully counterbalanced across participants.

The conditioning stage was designed to allow participants to learn associations between visual and outcome cues. The blocks followed the same structure as the associative learning stage, including both passive exposure and testing trials (Fig. 1D). In each trial of the exposure phase, a visual cue was first presented for 1.5 s, immediately followed by the associated outcome for 1.5 s, followed by an ITI of 2 s. Each exposure mini-block consisted of 4 trials (one presentation of each visual cue), which was followed by a testing mini-block. In each trial of the testing mini-block, a visual cue was presented together with the two outcome cues for 1 s. After this 1 s, the visual cue would disappear but the outcome cues would remain on the screen and the participants had a further 3 s to make their choice (total of 4 s response period). Participants had to select which outcome they thought was associated with the visual cue using the keys 'j' and 'l'. Once a choice had been selected, a black rectangle would highlight their choice for 0.5 s, followed by an ITI of 2 s. Participants were given feedback on their average performance at the end of each mini-block of 4 trials. Each block in the conditioning phase was made up of one exposure mini-block followed by one testing mini-block and consisted only of cues from one of the two groups of cues, with the relevant contextual background music playing throughout the block. Blocks alternated between each context with a total of 2 blocks per context. The relationship between the conditions (i.e., which context was learnt first, and which context was subject to TMR) was fully counterbalanced across participants.

The TMR manipulation was performed after the conditioning stage. During TMR, participants were instructed to rest for up to 60 min, during which participants were asked to alternate between resting with their eyes closed for 10 min, while remaining awake, and doing a jigsaw puzzle for 20 min. In total, they spent 20 min with their eyes closed and 40 min doing a jigsaw puzzle. During this awake rest period, the contextual soundtrack associated with the TMR group was played throughout. The contextual soundtrack used for TMR was fully counterbalanced across participants. To encourage participants to attend to the soundtrack, the soundtrack was turned on and off for periods of between 9 and 16 s, alternating with periods of silence. The length of each on and off period was selected from a random distribution.

After the TMR manipulation, participants were required to perform an inference test. During the inference test, participants were asked to infer the relevant outcome in response to each auditory cue.

Before beginning the first inference test, participants practised performing two tasks together using only two auditory cues, not taken from the task, and a 1-back task using digits rather than shapes. The practice consisted of up to 4 blocks with 24 trials per block. The timings of these practice trials matched the timings of the main inference test described below. The participants were not required to perform under time pressure as reaction times were not speeded.

We designed three different versions of the inference test which included distractor tasks designed to increase the load on visual working memory or general cognitive resources (Figure S1A). For all versions, participants were presented with an auditory cue and had to choose which of the two outcomes they thought was associated with the auditory cue (Fig. 2B). The auditory cue was presented together with the two outcome cues for 1.5 s. After 1.5 s, the auditory cue would stop but the outcome cues would remain on the screen and the participants had a further 3 s to make their choice (total of 4.5 s response period). Participants had to select which outcome they thought was associated with the auditory cue using the keys ‘j’ and ‘l’. Once a choice had been selected, a black rectangle would highlight this choice for 0.5 s, followed by an ITI of 1 s. Participants were told they would receive a bonus reward of up to £5 based on how many cues they correctly inferred. Participants were given feedback on their average performance at the end of each block. Auditory cues from both contextual groups were randomly interleaved in each block and there was no background music during the inference tests.

For 2 of the 6 blocks, the participants also performed a visual memory distractor task (Figure S1A). The purpose of this distractor task was to modulate task difficulty: a visual memory distractor task could impair recall of the intermediate visual cue when participants perform inference. The visual memory distractor task involved using a 1-back task with shapes similar to those seen in the task. On each trial, between presentations of each auditory cue, a shape would appear for 2 s, and the participants had to respond based on whether it was the same or different from the previous shape using the keys ‘a’ and ‘d’. The shape would become darker (participant chose ‘same’) or lighter (participant chose ‘different’) after the participant had made their choice. The shape remained on screen until the 2 s was complete. The shapes would appear in a random orientation which participants were instructed to ignore. Following a pause of 1 s, the auditory cue would begin, and the rest of the inference trial would commence as described above.

For another 2 of the 6 blocks, participants performed a non-memory-based distractor task (Figure S1A). The purpose of this distractor task was again to modulate task difficulty, but without the interfering effect of a visual memory load. On each trial, between presentations of each auditory cue, an image would appear for 2 s, and the participants had to respond based on whether the image depicted something living (e.g., a hedgehog) or something non-living (e.g., balloons). The image would become darker (participant chose ‘living’) or lighter (participant chose ‘non-living’) after the participant had made their choice. The image remained on screen until the 2 s was complete. Following a pause of 1 s, the auditory cue would begin, and the rest of the inference trial would commence as described above.

For the remaining 2 blocks, there was no distractor task. After the normal ITI of 1 s, there would be an additional pause of 3 s to ensure trials across all versions of the inference task were the same length. Following this pause, the auditory cue would begin, and the rest of the inference trial commenced as described above. These 2 blocks which did not include distractor tasks always occurred at the end to minimise order effects of participants practising the inference task on its own.

After all 6 blocks of the first inference test, participants redid the conditioning stage of the experiment but with half of the associations between visual and outcome cues flipped (‘value flip’) (Fig. 3A). The protocol for the value-flip stage was otherwise the same as for the conditioning stage (Fig. 1D) except participants completed 1 block per context (rather than the 2 blocks per context done in the initial conditioning stage). After the value-flip, a second inference test (‘inference test 2’) was conducted using the same procedure as previously described and also included 6 blocks (Fig. 2B; Figure S1A).

In the indirect/direct test, we tested both the indirect (auditory–outcome) and direct (auditory–visual) associations paired with each auditory cue (Fig. 4B). All visual cues and outcomes were present on the screen. By having all cues on the screen, we could test whether there were differences in how long participants spent looking at the visual cues during inference in the TMR and no TMR groups. However, by displaying the visual cues during the indirect/direct test participants were provided with an opportunity to solve inference for all auditory cues (A) using the intermediary visual cues (B), rather than relying on a short-cut (A → C) generated during memory reactivation. Therefore, to test the effect of TMR on inferential choice it was necessary to only test inferential choice in the presence of the visual cues at the very end of the experiment, after the value-flip manipulation. For the first 2 blocks of this test participants completed the indirect test: on each trial an auditory cue was presented, and participants had to select one of the two outcomes using the mouse. For the final 2 blocks of this test participants completed the direct test: on each trial the auditory cue was presented, and participants had to select one of the 8 possible visual cues which was paired with the sound. This allowed us to assess performance on the directly learnt associations as well as the indirect auditory–outcome association. In each block a total of 24 trials were included, with equal number of trials in the TMR and no TMR condition. For all blocks, the auditory cue was presented for 1.5 s. Following this, participants had 30 s to select either the outcome cue (first 2 blocks) or the visual cue (last 2 blocks) associated with the auditory cue. Participants made their selection by moving the cursor over the chosen cue and pressing ‘spacebar’ to confirm. Once a choice had been selected, a black rectangle would highlight this choice for 0.5 s, followed by an ITI of 1 s. Participants were given feedback on their performance at the end of each block.

### Eye-tracking data preprocessing

For 2 participants, no usable eye-tracking data was collected. The remaining 23 participants were all included in eye-tracking analyses. The conversion of the EyeLink® 1000 Edf files was done with the Edf2Mat Matlab Toolbox designed and developed by Adrian Etter and Marc Biedermann at the University of Zurich. Gaze position and pupillometry data were smoothed using a 50 ms Gaussian kernel. Blinks were detected from the smoothed pupillometry data and removed from the gaze position data. For each participant, the eye with the fewest missing samples was selected and used for all further analyses. The gaze position was epoched around the auditory cue onset for trials in the second inference test and for indirect trials in the indirect/direct test. For the comparison of Flipped cues (or Not Flipped cues), we analysed the percentage of trials where participants looked

at the incorrect outcome cue, using a binary label to establish whether the incorrect outcome had been viewed on any given trial. To investigate whether participants relied on knowledge of the intermediary visual cue, we analysed the time spent looking at the visual cues in response to auditory cues in the TMR and no TMR groups on correct trials of the indirect test. Due to the large variation in the average time spent looking at the visual cues, we z-scored the fixation lengths across all trials for each participant. We further analysed the proportion of trials where participants looked at the visual cues in response to auditory cues in the TMR and no TMR groups on correct trials of the indirect test.

### Analysis and statistical tests

Analysis was conducted in MATLAB (version R2022a) and in Python v.3.6 (<https://www.python.org/downloads/release/python-363/>), using the Python packages DABEST<sup>67</sup>, scipy<sup>68</sup>, numpy<sup>69</sup>, matplotlib<sup>70</sup>, seaborn<sup>71</sup>, pandas<sup>72</sup>. Throughout this study, we used a bootstrap-coupled estimation of effect sizes<sup>73</sup>. For each estimation plot showing a difference between TMR and No TMR groups: the left panel shows the distribution of raw data points for the entire dataset, superimposed on plots reporting group means  $\pm$  SEM; and the right panel displays the difference between the TMR group and the No TMR group, computed from 100,000 bootstrapped resamples. For each estimation plot: black dot, mean; black ticks, 90% or 95% confidence interval; and filled curve: bootstrapped sampling error distribution. All p-values were estimated using bootstrapping, across 100,000 bootstrapped resamples. We used 90% confidence intervals (one-sided) for tests of difference between TMR and No TMR groups, with the prediction that TMR improves accuracy, unless noted otherwise. For any other tests, we used 95% confidence intervals (two-sided). All tests are performed on  $n = 25$  participants (degrees of freedom (d.f.) = 24), except for all tests on eye data which were performed on  $n = 23$  participants (d.f. = 22).

### Data availability

The data and code used in this study will be made available via the MRC BNDU Data Sharing Platform (<https://data.mrc.ox.ac.uk>) upon publication: the contact for this platform is Ben Micklem, Research Support Manager, at [ben.micklem@bndu.ox.ac.uk](mailto:ben.micklem@bndu.ox.ac.uk).

Received: 12 February 2025; Accepted: 10 June 2025

Published online: 09 July 2025

### References

1. Tolman, E. C. Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189–208 (1948).
2. Brogden, W. J. Sensory pre-conditioning. *J. Exp. Psychol.* **25**, 323 (1939).
3. Bottini, R. & Doeller, C. F. Knowledge across reference frames: Cognitive maps and image spaces. *Trends Cogn. Sci.* **24**, 606–619 (2020).
4. Behrens, T. E. J. et al. What is a cognitive map? Organizing knowledge for flexible behavior. *Neuron* **100**, 490–509 (2018).
5. Peer, M., Brunec, I. K., Newcombe, N. S. & Epstein, R. A. Structuring knowledge with cognitive maps and cognitive graphs. *Trends Cogn. Sci.* **25**, 37–54 (2021).
6. Whittington, J. C. R., McCaffary, D., Bakermans, J. J. W. & Behrens, T. E. J. How to build a cognitive map. *Nat. Neurosci.* **25**, 1257–1272 (2022).
7. Wamsley, E. J. Offline memory consolidation during waking rest. *Nat. Rev. Psychol.* **1**, 441–453 (2022).
8. Diekelmann, S. & Born, J. The memory function of sleep. *Nat. Rev. Neurosci.* **11**, 114–126 (2010).
9. Coutanche, M. N., Gianessi, C. A., Chanales, A. J. H., Willison, K. W. & Thompson-Schill, S. L. The role of sleep in forming a memory representation of a two-dimensional space. *Hippocampus* **23**, 1189–1197 (2013).
10. Werchan, D. M. & Gómez, R. L. Generalizing memories over time: Sleep and reinforcement facilitate transitive inference. *Neurobiol. Learn. Mem.* **100**, 70–76 (2013).
11. van Kesteren, M. T. R., Fernández, G., Norris, D. G. & Hermans, E. J. Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. *Proc. Natl. Acad. Sci.* **107**, 7550–7555 (2010).
12. Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D. & Walker, M. P. Human relational memory requires time and sleep. *Proc. Natl. Acad. Sci.* **104**, 7723–7728 (2007).
13. Lau, H., Alger, S. E. & Fishbein, W. Relational memory: A daytime nap facilitates the abstraction of general concepts. *PLoS ONE* **6**, e27139 (2011).
14. Craig, M., Dewar, M., Harris, M. A., Della Sala, S. & Wolbers, T. Wakeful rest promotes the integration of spatial memories into accurate cognitive maps. *Hippocampus* **26**, 185–193 (2016).
15. Foster, D. J. Replay comes of age. *Annu. Rev. Neurosci.* **40**, 581–602 (2017).
16. Joo, H. R. & Frank, L. M. The hippocampal sharp wave–ripple in memory retrieval for immediate use and consolidation. *Nat. Rev. Neurosci.* **19**, 744–757 (2018).
17. Buzsáki, G. Hippocampal sharp wave–ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus* **25**, 1073–1188 (2015).
18. Bendor, D. & Wilson, M. A. Biasing the content of hippocampal replay during sleep. *Nat. Neurosci.* **15**, 1439–1444 (2012).
19. Kudrimoti, H. S., Barnes, C. A. & McNaughton, B. L. Reactivation of hippocampal cell assemblies: Effects of behavioral state, experience, and EEG dynamics. *J. Neurosci.* **19**, 4090–4101 (1999).
20. O'Neill, J., Senior, T. J., Allen, K., Huxter, J. R. & Csicsvari, J. Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nat. Neurosci.* **11**, 209–215 (2008).
21. Ólafsdóttir, H. F., Bush, D. & Barry, C. The role of hippocampal replay in memory and planning. *Curr. Biol.* **28**, R37–R50 (2018).
22. Schuck, N. W. & Niv, Y. Sequential replay of nonspatial task states in the human hippocampus. *Science* **364**, eaaw5181 (2019).
23. Dragoi, G. & Tonegawa, S. Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* **469**, 397–401 (2011).
24. Ólafsdóttir, H. F., Barry, C., Saleem, A. B., Hassabis, D. & Spiers, H. J. Hippocampal place cells construct reward related sequences through unexplored space. *Elife* **4**, e06063 (2015).
25. Liu, Y., Dolan, R. J., Kurth-Nelson, Z. & Behrens, T. E. J. Human replay spontaneously reorganizes experience. *Cell* **178**, 640–652.e14 (2019).
26. Barron, H. C. et al. Neuronal computation underlying inferential reasoning in humans and mice. *Cell* **183**, 228–243.e21 (2020).
27. Wu, X. & Foster, D. J. Hippocampal replay captures the unique topological structure of a novel environment. *J. Neurosci. Off. J. Soc. Neurosci.* **34**, 6459–6469 (2014).

28. Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S. & Redish, A. D. Hippocampal replay is not a simple function of experience. *Neuron* **65**, 695–705 (2010).
29. Rawson, A. B., Nalluru, S., O'Reilly, J. X. & Barron, H. C. Memory reactivation generates new, adaptive behaviours that reach beyond direct experience. *Sci. Rep.* **14**, 30097 (2024).
30. Rasch, B., Büchel, C., Gais, S. & Born, J. Odor cues during slow-wave sleep prompt declarative memory consolidation. *Science* **315**, 1426–1429 (2007).
31. Carbone, J. & Diekelmann, S. An update on recent advances in targeted memory reactivation during sleep. *NPJ Sci. Learn.* **9**, 31 (2024).
32. Oudiette, D. & Paller, K. A. Upgrading the sleeping brain with targeted memory reactivation. *Trends Cogn. Sci.* **17**, 142–149 (2013).
33. Lewis, P. A. & Bendor, D. How targeted memory reactivation promotes the selective strengthening of memories in sleep. *Curr. Biol.* **29**, R906–R912 (2019).
34. Tambini, A. & Davachi, L. Awake reactivation of prior experiences consolidates memories and biases cognition. *Trends Cogn. Sci.* **23**, 876–890 (2019).
35. Rudoy, J. D., Voss, J. L., Westerberg, C. E. & Paller, K. A. Strengthening individual memories by reactivating them during sleep. *Science* **326**, 1079–1079 (2009).
36. Tadros, T. & Bazhenov, M. Role of sleep in formation of relational associative memory. *J. Neurosci.* **42**, 5330–5345 (2022).
37. Abdou, K. et al. Prefrontal coding of learned and inferred knowledge during REM and NREM sleep. *Nat. Commun.* **15**, 4566 (2024).
38. Tamminen, J., Payne, J. D., Stickgold, R., Wamsley, E. J. & Gaskell, M. G. Sleep spindle activity is associated with the integration of new memories and existing knowledge. *J. Neurosci.* **30**, 14356–14360 (2010).
39. Tamminen, J., Ralph, M. A. L. & Lewis, P. A. The role of sleep spindles and slow-wave activity in integrating new information in semantic memory. *J. Neurosci.* **33**, 15376–15381 (2013).
40. Tamminen, J., Lambon Ralph, M. A. & Lewis, P. A. Targeted memory reactivation of newly learned words during sleep triggers REM-mediated integration of new memories and existing knowledge. *Neurobiol. Learn. Mem.* **137**, 77–82 (2017).
41. Landmann, N. et al. The reorganisation of memory during sleep. *Sleep Med. Rev.* **18**, 531–541 (2014).
42. Hu, X., Cheng, L. Y., Chiu, M. H. & Paller, K. A. Promoting memory consolidation during sleep: A meta-analysis of targeted memory reactivation. *Psychol. Bull.* **146**, 218–244 (2020).
43. Preston, A. R., Shrager, Y., Dudukovic, N. M. & Gabrieli, J. D. E. Hippocampal contribution to the novel use of relational information in declarative memory. *Hippocampus* **14**, 148–152 (2004).
44. Robinson, S. et al. Chemogenetic silencing of neurons in retrosplenial cortex disrupts sensory preconditioning. *J. Neurosci. Off. J. Soc. Neurosci.* **34**, 10982–10988 (2014).
45. Jones, J. L. et al. Orbitofrontal cortex supports behavior and learning using inferred but not cached values. *Science* **338**, 953–956 (2012).
46. Wimmer, G. E. & Shohamy, D. Preference by association: How memory mechanisms in the hippocampus bias decisions. *Science* **338**, 270–273 (2012).
47. Sadacca, B. F. et al. Orbitofrontal neurons signal sensory associations underlying model-based inference in a sensory preconditioning task. *Elife* **7**, e30373 (2018).
48. Fernandez-Leon, J. A., Hansen, B. J. & Dragoi, V. Representation of rapid image sequences in V4 networks. *Cereb. Cortex* **28**, 2675–2684 (2018).
49. Schreiner, T. & Rasch, B. Boosting vocabulary learning by verbal cueing during sleep. *Cereb. Cortex* **25**, 4169–4179 (2015).
50. Oudiette, D., Antony, J. W., Creery, J. D. & Paller, K. A. The role of memory reactivation during wakefulness and sleep in determining which memories endure. *J. Neurosci.* **33**, 6672–6678 (2013).
51. Antony, J. W., Gobel, E. W., O'Hare, J. K., Reber, P. J. & Paller, K. A. Cued memory reactivation during sleep influences skill learning. *Nat. Neurosci.* **15**, 1114–1116 (2012).
52. Berkers, R. M. W. J. et al. Cued reactivation during slow-wave sleep induces brain connectivity changes related to memory stabilization. *Sci. Rep.* **8**, 16958 (2018).
53. Cairney, S. A., Durrant, S. J., Hulleman, J. & Lewis, P. A. Targeted memory reactivation during slow wave sleep facilitates emotional memory consolidation. *Sleep* **37**, 701–707 (2014).
54. van Dongen, E. V. et al. Memory stabilization with targeted reactivation during human slow-wave sleep. *Proc. Natl. Acad. Sci.* **109**, 10575–10580 (2012).
55. Batterink, L. J., Creery, J. D. & Paller, K. A. Phase of spontaneous slow oscillations during sleep influences memory-related processing of auditory cues. *J. Neurosci.* **36**, 1401–1409 (2016).
56. Schreiner, T. & Rasch, B. Cueing vocabulary in awake subjects during the day has no effect on memory. *Somnologie - Schlaforschung Schlafmed.* **19**, 133–140 (2015).
57. Schönauer, M., Geisler, T. & Gais, S. Strengthening procedural memories by reactivation in sleep. *J. Cogn. Neurosci.* **26**, 143–153 (2014).
58. Tambini, A., Berners-Lee, A. & Davachi, L. Brief targeted memory reactivation during the awake state enhances memory stability and benefits the weakest memories. *Sci. Rep.* **7**, 15325 (2017).
59. Hoffman, L. J., Foley, J. M., Tanriverdi, B., Chein, J. & Olson, I. R. Awake targeted memory reactivation doesn't work. *Mem. Cognit.* <https://doi.org/10.3758/s13421-024-01576-x> (2024).
60. Alm, K. H., Ngo, C. T. & Olson, I. R. Hippocampal signatures of awake targeted memory reactivation. *Brain Struct. Funct.* **224**, 713–726 (2019).
61. Schapiro, A. C., McDevitt, E. A., Rogers, T. T., Mednick, S. C. & Norman, K. A. Human hippocampal replay during rest prioritizes weakly learned information and predicts memory performance. *Nat. Commun.* **9**, 3920 (2018).
62. Ngo, H.-V.V. & Staresina, B. P. Shaping overnight consolidation via slow-oscillation closed-loop targeted memory reactivation. *Proc. Natl. Acad. Sci.* **119**, e2123428119 (2022).
63. Rihm, J. S., Diekelmann, S., Born, J. & Rasch, B. Reactivating memories during sleep by odors: Odor specificity and associated changes in sleep oscillations. *J. Cogn. Neurosci.* **26**, 1806–1818 (2014).
64. Shanahan, L. K., Gjorgieva, E., Paller, K. A., Kahnt, T. & Gottfried, J. A. Odor-evoked category reactivation in human ventromedial prefrontal cortex during sleep promotes memory consolidation. *Elife* **7**, e39681 (2018).
65. Kleiner, M., Brainard, D. & Pelli, D. What's new in Psychtoolbox-3? (2007).
66. Endo, N., Saiki, J., Nakao, Y. & Saito, H. Perceptual judgments of novel contour shapes and hierarchical descriptions of geometrical properties. *Jpn. J. Psychol.* **74**, 346–353 (2003).
67. Ho, J., Tumkaya, T., Aryal, S., Choi, H. & Claridge-Chang, A. Moving beyond P values: Data analysis with estimation graphics. *Nat. Methods* **16**, 565–566 (2019).
68. Virtanen, P. et al. SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261–272 (2020).
69. Harris, C. R. et al. Array programming with NumPy. *Nature* **585**, 357–362 (2020).
70. Matplotlib: A 2D graphics environment. <https://ieeexplore.ieee.org/document/4160265/>.
71. Waskom, M. L. seaborn: Statistical data visualization. *J. Open Source Softw.* **6**, 3021 (2021).
72. McKinney, W. Data structures for statistical computing in python, in *Proceedings of the 9th Python in Science Conference* vol. 445, 51–56 (Austin, 2010).
73. Efron, B. The bootstrap and modern statistics. *J. Am. Stat. Assoc.* **95**, 1293–1296 (2000).

## Acknowledgements

This research was supported by funding from the Medical Research Council (UKRI-MRC) (MR/W01971X/1, MR/L019639/1 to J.X.R.), UKRI (MR/W008939/1 to H.C.B.), and the Biotechnology and Biological Sciences Research Council (UKRI-BBSRC) (BB/M011224/1 to A.B.R.). The Wellcome Centre for Integrative Neuroimaging was supported by core funding from the Wellcome Trust (203139/Z/16/Z). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

## Author contributions

All authors contributed to the experimental design and editing the paper. C.M.S, H.C.B, and J.X.O.R contributed to data analysis and writing the paper. C.M.S. collected the data. J.X.O.R and H.C.B. jointly supervised the project.

## Declarations

### Competing interests

The authors declare no competing interests.

### Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-025-06742-y>.

**Correspondence** and requests for materials should be addressed to H.C.B. or J.X.O.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025