Memory reactivation in slow wave sleep enhances relational learning

Lorena Santamaria [1], Ibad Kashif[1], Niall McGinley[1] and Penelope A. Lewis[1]
[1]Cubric, Cardiff University, UK

Abstract

Memories are often strengthened by sleep, which can also boost integration and relational learning. This process can be facilitated by a technique called targeted memory reactivation (TMR), which involves re-applying cues that were associated with learned material during wake during subsequent sleep. We tested whether TMR during slow wave sleep increases the solving of inference pairs in a transitive inference task. Because the slow oscillation up-state is thought to be more important for plasticity, we also asked whether stimulation is more beneficial at this phase. Our data show that TMR can boost inference on the most difficult pairs, but only when presented during the down-to-up transition of the slow oscillation. Such stimulation was associated with classifiable replay, whereas stimulation of the up-to-down transition produced no apparent replay and led to below-chance performance. These findings demonstrate that targeted memory replay in sleep can play a role in integration and relational learning.

(150 words, max 150)

1 Introduction

Relational memory is the ability to integrate multiple sources of knowledge, infer indirect associations between stimuli, and make decisions when presented with novel situations[1,2]. One typical example of such integration is transitive inference (TI), or the deduction of the rankings of non-adjacent members of a linear hierarchy which has been presented in terms of adjacent members. In simpler words, knowing the ranking of adjacent pairs, e.g. A>B and B>C can allow deductions about non-adjacent pairs: that A>C in an A>B>C hierarchy.

Despite being studied for many decades in humans[3] and in numerous other species[4–6], the mechanisms for TI remain elusive[7,8]. It was not until relatively recently that the dependence of TI on sleep was studied in humans[9]. Thus, Ellenbogen 2007, and its replication[10], demonstrated that sleep leads to increased inference ability compared with wake, but only for the most difficult inference pairs. Another study investigated the impact of both sleep and time on TI, finding that a nap is not enough to induce benefit[8].

While its role in TI remains unclear, numerous studies suggest that sleep facilitates the integration of recent experiences with prior knowledge[11–14]. Prior learning is spontaneously reactivated during sleep, and this is thought to be important for memory consolidation[15]. Importantly, such reactivation can be directly manipulated by re-administering sensory cues that have been paired with learned information during subsequent sleep, using a technique known as targeted memory reactivation (TMR), see[16] for a meta-analysis.

Although some prior work had linked cognitive performance enhancement with rapid eye movement sleep (REM), e.g.[17], most studies focus on non-rapid eye movement sleep (NREM)[18], and particularly slow wave sleep (SWS). The main rhythms of SWS are slow oscillations (SO), spindles, and sharp wave ripples. SOs are low-frequency oscillations (0.05-4Hz) that originate in the cortex[19] and reflect alternation between hyper-polarised "down-states" and depolarised neuronal "up-states"[20]. The depolarising neocortical "up-sate" of the SO drives thalamo-cortical spindles, which are transient oscillations at 9-16Hz[21], and strongly linked to reactivation[15]. Closed-loop TMR (CL-TMR), in which the TMR cues are delivered at a specific SO phase, has demonstrated distinct roles for up and down states of the SO in boosting behavioural performance[22,23]. These studies join closed loop stimulation studies without TMR showing improved task performance after stimulation of the peak or up-going phase[22,23]. Such beneficial impacts are thought to be due to more efficient reactivation at that phase of the oscillation[15], which together with the naturally enhanced synchronisation of neural firing during the up-states[24], increases the probability of consolidation.

To investigate how a night of sleep influences consolidation of the TI task, and in particular the role of the SO-specific phases, we adapted the experimental set-up used in[9] for three hierarchies so we could apply CL-TMR in a within-subjects design. One hierarchy was stimulated during the down-to-up transitions of the SOs (Up condition), another during the up-to-down transitions (Down condition) and the remaining hierarchy was not stimulated (Control condition). In keeping with previous closed-loop literature and previous TI studies, we expected TMR to benefit only the inference pairs. Particularly, we predicted a benefit in the Up condition but not in the Down condition, since there is evidence that presenting the sounds at the onset of the down-states may block the memory benefit of TMR[22].

In order to test the robustness of the TMR effect over time, participants performed a third session after two weeks. Only a few studies have examined how long the sleep-meditated benefits of TMR last, and the results are inconsistent[25–27]. In keeping with other types of paired associate learning[28], we expected premise pair knowledge to decrease over time. Similarly, inference pair performance for the Control and Down conditions should decrease, but we

hypothesised that TMR benefit for the Up condition might remain stable given the extra consolidation that has occurred.

Our data show that TMR in the SO Up condition is associated with classifiable reactivation and better inferential reasoning, but this is specific to the most difficult inference pairs. On the other hand, stimulation in the Down condition produced no evidence of reactivation, and led to inhibition of inferential reasoning that recovered after two weeks.

2 Results

Figure 1. Experimental paradigm. A) Timeline: first, participants were wired-up (PSG). Participants in the first session completed four behavioural tasks: sound-image association, sound-image test, premise pair learning and immediate test. During learning, premise pairs of three hierarchies were presented. The Immediate test was exactly the same as the learning counterpart but without feedback, allowing us to assess initial premise pair knowledge. Then, subjects went to sleep and TMR was carried out in SWS, where two out of the three hierarchies were stimulated. After waking up, still with the electrodes on, participants completed Session 2, which comprised a sound-image association and a Late test. During the Late test, Inference pairs were presented for the first time, together with the previously learned premise pairs. After 2-weeks they returned to the lab for a third session, without EEG recording, where they completed the same tests as in Session 2 plus the awareness questionnaire. B) Example of the three categories used in the experiment: outdoor scenes, uncommon objects and female faces. C) Example of the visual stimuli arranged in a randomly determined hierarchical order. D) Resultant inference pairs from (C). E) Example of a typical trial of the learning part. A cross appearing in the screen indicates the starting of the trial (lasting between 0.5s to 1s), then two images of the hierarchy appear on the screen lasting until the user presses the up or down arrows. Feedback is provided immediately afterwards (a happy or angry face accordingly). To help sound-image consolidation, the sound of each stimulus is played while the stimulus is presented again on the screen. Finally, a purple dot is presented for 500ms to indicate the end of the trial. During the test part, the feedback and the sounds are eliminated.

2.1 Behavioural results

2.1.1 Performance of training

During learning only premise pairs were presented to the participants (n=20). The mean number of blocks needed to reach the criterion was 5.31±1.41. Non-significant differences (smallest p=0.693) were found in the number of blocks needed to reach the criteria for each stimulus type (Faces M:5.1 SE:0.20, Objects M:5.6 SE:0.39, Scenes M:5.2 SE:0.33). The number of blocks to learn to criterion was considerably lower than in previous studies[9], we theorised that this is mainly due to the type of stimuli used, e.g., real versus abstract images.

2.1.2 Premise pairs

Immediate test performance was above chance (50%) in all conditions (Down M:74.7% SE:0.02, Control M:79.2% SE:002, Up M:79.6% SE:002) but not at ceiling level (See table SM3.2) in the first session. A 1-way ANOVA showed no difference between Conditions (Up, Down and Control) at baseline (F=2.69, p=0.07). Post-hoc analysis confirmed this result: Down vs Control (p=0.095), Down vs Up (p=0.095) and Up vs Control (p=0.93). Therefore, we can be confident that participants were equally able to express their premise pair knowledge in the Immediate test for all conditions.

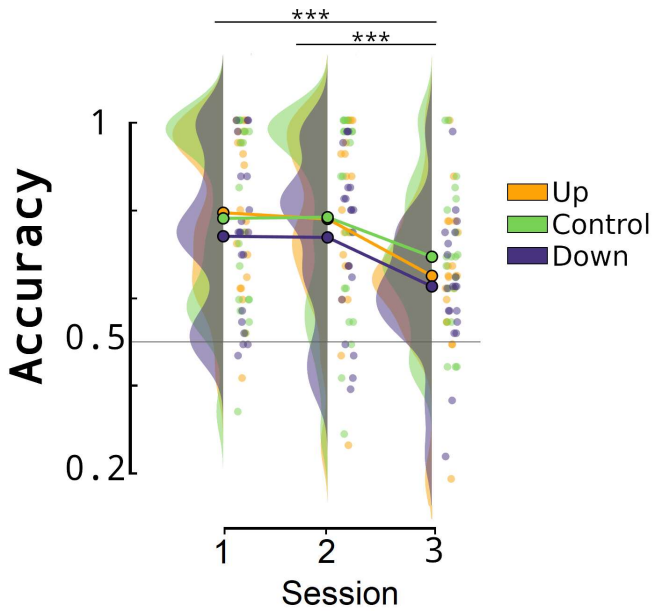

Figure 2 Premise pair accuracy results for each Session (x-axis) and for each one of the three conditions: Up (yellow), Control (green) and Down (purple). Statistically significant differences are indicated as: *p<0.05, ** p<0.01, ***p<0.001.

Performance on the sound-image association was over 90% for all the stimuli (refer to SM1 for results) and the number of times each stimulus was used for each condition remained relatively constant, as can be seen in Table SM3.1. Hence, we can rule out any bias due to stimuli or cue-memory associations.

To examine how the effect of TMR evolved over time, we performed a RM-ANOVA with Session (3 levels) and Condition (3 levels) as within subject's factors, and with accuracy as the dependent variable (see Figure 2). We found Condition ($F=3.5$, $p=0.029$) and Session ($F=32.39$, $p<0.0001$) effects, but no interaction ($p=0.72$). Nevertheless, as in Session 1, additional analysis revealed no differences (smallest $p=0.103$) between Conditions. On the other hand, clear Session effects emerged between Session 1 and Session 3 and between Session 2 and Session 3 for each Condition, as can be seen from Figure 2 (all $p<0.001$). Thus, there was a marked drop in premise pair performance after two weeks, in keeping with normal declarative forgetting. The fact that we saw no overnight improvement in premise pair accuracy is in line with[9] and its replication[10], where there were no differences between Session 1 and Session 2 in premise pair performance for either Sleep or Wake groups. This null result might be surprising given that associative memories are often strengthened by sleep[15], but it is also in line with the idea that sleep mainly facilitates more weakly encoded memories[30], though this could also depend on other factors such as the type of task or the strength of the cue-memory associations[31].

On an additional note, the well documented serial position effect (SPE[32]) in which performance of the premise pairs usually follows a U-shaped pattern with better performance in more extreme pairs (e.g., AB and EF) is clear for the first two sessions but not for the third (see Figure SM3.1). The loss of these classic differences in accuracy across the five premise pairs, particularly for the middle pairs (BC, CD, DE) during session 3 (Figure SM3.1C) may be attributable to forgetting after the two-week gap.

2.1.3 Inference pairs

Inference pairs were introduced in the Late test performed in the morning (Session 2) and repeated in the 2-week follow up session (Session 3). To examine them, we again performed a RM-ANOVA with Session (2 levels) and Condition (3 levels). Results showed a significant interaction ($F=6.21$, $p=0.002$) but no Condition ($p=0.098$) or Session ($p=0.11$) effect. Further analysis indicated a strong Condition effect for Session 2 ($F=8.7$, $p<0.001$) with differences in both Up vs Down and Control vs Down comparisons, and with Down being the only condition that remained at under chance level (50%) (refer to figure 3 (A-left), and table SM4.1 for full

detail of accuracy results). Therefore, as predicted, inference of new pairs benefitted more from TMR in the Up condition compared with TMR in the Down condition.

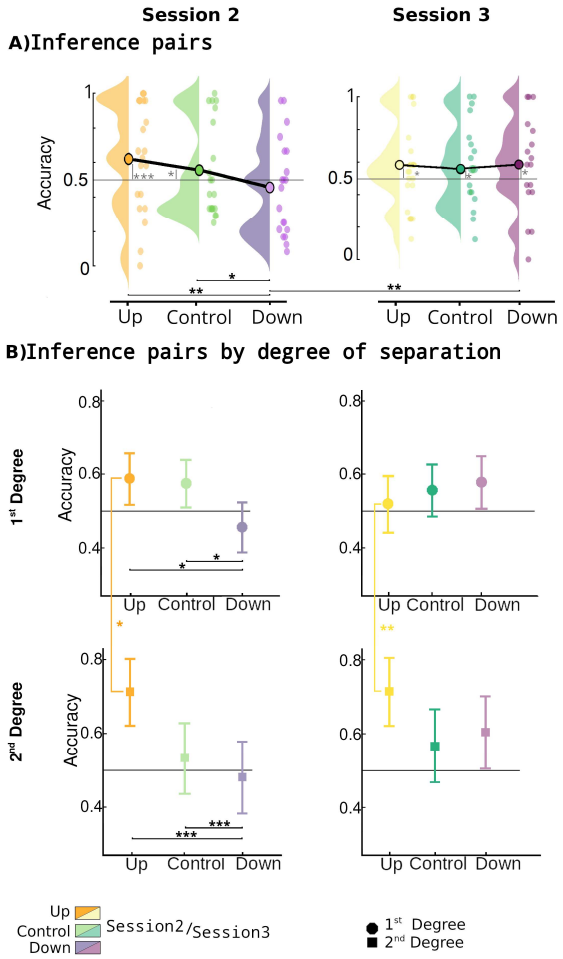

Figure 3. Inference pair performance for each condition: Up (yellow colours), Down (purple colours) and Control (green colours. A) Performance of averaged inference pairs for Session 2 (next morning) on the left and Session 3 (2-weeks-follow-up) on the right. B) Inference pairs performance separated by degree of separation: first (circles) and second (squares) for Session 2 (left) and Session 3 (right). Statistically significant differences are indicated as: *p<0.05, ** p<0.01, ***p<0.001. Statistics colour code: Grey: chance level (50%), Black: between conditions or sessions, Coloured: within condition.

However, contrary to expectations, although the average inference pair performance for the second session was numerically higher for Up (M: 62.8% SE:2.8) than Control (M:56.0% SE:2.7), this was not significant. This differs markedly from previous CL-TMR literature, where significant differences were found between Up and Control[22]. However, that study used a different behavioural task, presumably involving alternative memory mechanisms.

We also tested for a relationship between TMR benefit in the Up and Down states. There is no correlation between the performance in these two conditions (p=0.38) ruling out any relationship between Up and Down performance. Additionally, it has already been demonstrated that inference pair performance is strongly sleep dependant[9,10], so an above chance performance after sleep is not surprising even for the non-cued condition. This finding

was consistent across participants, with 65% of participants presenting higher inference pair performance for the Up than Control, 75% with Up-performance greater than Down, and 65% Control-performance larger than Down (see Figure SM4.1).

Interestingly, the difference between conditions disappeared after 2-weeks (Figure 3(A-right)), with the smallest p-value for better performance in Up than Control condition being (p=0.090). Furthermore, there was an unexpected improvement in the Down condition (F=8.96, p=0.003) which increased from under chance in Session 2 (M:46.4% SE:2.8) to over chance level in Session 3 (M:58.6%, SE:2.9). Thus, the disruption caused by down-state TMR stimulation only affected performance the next morning, and this stimulation still improved transitive inference in the long term. We speculate that subsequent nights of sleep may have counteracted the inhibition of replay initially caused by the TMR in the down states, with spontaneous up-state replay during this period acting in a similar manner to the Up stimulation.

We also examined the differences in performance between the first and second degree of separation (see Figure 1(D)) of the inference pairs. We repeated the same analysis as above, now adding Degree (2 levels) as a third factor. This revealed a significant effect of Degree (p=0.017), the interaction between Session and Condition that was discussed above, and another interaction between Condition and Degree (p=0.015). There was no interaction between Session and Degree (p=0.55) or Session, Degree and Condition (p=0.69). Following these results, we investigated the Degree effect for each condition independently. That is, we performed a RM-ANOVA with Session (2 levels) and Degree (2 levels) for each condition. Results are illustrated in Figure 3(B), left for Session 2 and right for Session 3, but refer to SM4 for full statistics. Regarding Session, and following the aforementioned results, only the Down condition (purple colours) showed a significant effect (p=0.0013), and this is driven by the first degree of separation (first degree: p=0.009, second degree: p=0.13). On the other hand, only the Up condition (yellow) showed a Degree effect (p=0.023) with better performance for the second degree of separation for both sessions (p=0.037, p=0.002 respectively). Importantly, the second degree of separation is only significantly different (p=0.005) from the Control condition (green) in Session 2. This is line with prior work[9], showing a sleep benefit to the second degree of separation, and also with other sleep literature showing that the most difficult tasks received the maximum sleep-dependant benefit[33,34]. However, this pattern was not present in the Control condition (green colours) where no significant difference was found between the degree of separation or between session (min p=0.44).

In summary, we found a lasting effect of TMR on the TI task. Stimulation on the up states only presented benefits for the most difficult pairs, that is the inference pairs with second

degree separation (e.g., BE pair), compared to Control. Importantly, these differences between the first and second degree of separation were only present in the Up condition. Behavioural performance on inference pairs remaining constant after two weeks for both Up and Control conditions. Interestingly, down-state stimulation inhibited sleep-benefits, decreasing next morning inference pair performance compared to Control. However, this picture changed after two-weeks, where performance on inference pairs (both averaged and separated by difficulty) improved to above chance levels such that it did not differ from the other two conditions, Control and Up.

2.2 EEG results

In order to examine the effect of TMR for each SO phase, we played two types of sounds, one set associated with the previously learned hierarchy (Experimental sounds) and one set of novel sounds forming an unlearned hierarchy (Control sounds) as indicated in Figure 5(A). Refer to section 4.5 for full details.

2.2.1 Efficiency of the closed-loop algorithm

To check if the online algorithm correctly differentiated between the Up and Down conditions, we calculated ERPs separately for each Condition (Up/Down) and Cue type (Experimental/Control sounds). An ERP cluster analysis showed three clusters that differed between Up and Down for experimental sounds and another three for control sounds before and during the stimulus onset (see Figure 4). After the stimulus, Up and Down conditions followed similar temporal evolution. In both cases, two of the clusters were found before the stimuli onset and one straight after, where opposite polarities were expected. The clusters involved most of the occipital and parietal electrodes. As expected, these ERPs demonstrated that the cues for the Up and Down conditions were played in highly distinctive oscillation phases, and that the differences between ERP responses to Experimental and Control sounds were minimal.

As an additional step to corroborate the accuracy of the closed-loop algorithm we calculated the phase of the cortical SO at stimulus onset for each trial and participant (see Figure 5(C)). For experimental sounds, the average values at channel F3 were 358.20° (SD:0.58) and 205.61° (SD:0.44) for the Up and Down conditions respectively. Similar values were obtained for the control sounds: 358.79° (SD:0.59) and 208.37° (SD:0.50) respectively. Circular statistics corroborate a significant difference between Up and Down conditions for

both stimulus and control (p<0.001) but not differences between these sounds for either Down or Up (p>0.1).

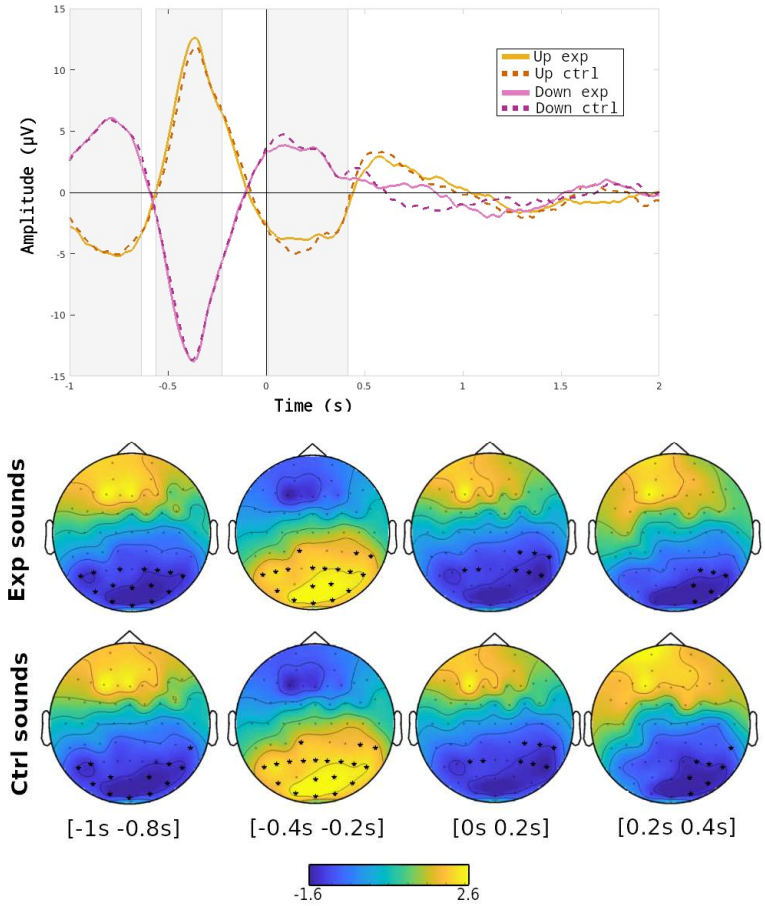

Figure 4: ERPs statistics. Top graph: ERPs for Up (yellow) and Down (purple) conditions and experimental (solid line) and control (dashed line) sounds. There are no differences between experimental and control sounds but there are differences, as expected, between Up and Down stimulations. These differences are highlighted in grey in the ERP plots and represented in the topographies (Bottom graph). Three significant clusters showing differences between Up and Down condition for experimental (top row of topographies) and control sounds (bottom row of topographies) were found. The first two columns of topographies correspond to the first two grey areas from the left in the ERP plot. The last cluster was divided into two topographies as the number of channels involved varied slightly within the cluster. These clusters are in the occipital-parietal areas and are found on the peaks of the oscillations before and up to around 400ms after the stimulation.

2.2.2 Time-frequency analysis

We performed time-frequency analysis independently for the Up and Down conditions, examining the differences between experimental and control sounds and the interaction effect. We focussed on two frequency bands[22,35,36], theta (5-8Hz) and spindle (9-16Hz) for negative clusters only (1-tailed)[22]. After

cues played on the Up condition we found a significant decrease in theta power between 1900ms and 2200ms for experimental compared to control sounds involving a cluster of 23 channels mainly centred in the left hemisphere (p=0.045, see Figure 6(A) and (B)). We

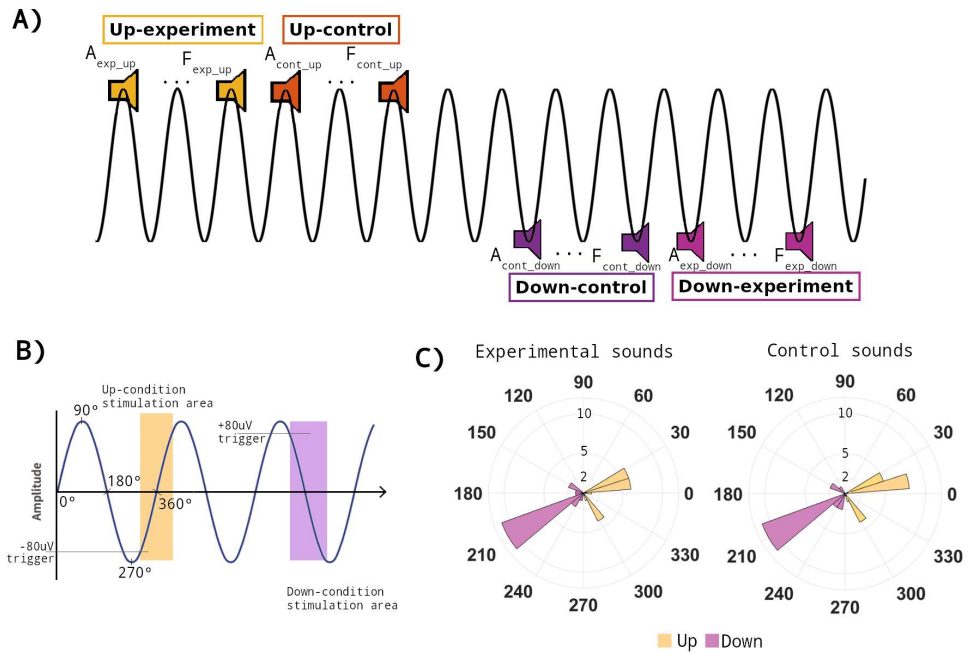

Figure 5. TMR protocol. A) Each element of the hierarchy was played in the right order for the experimental condition, control sounds were assigned to a random hierarchy. The order of Up/Down and Experimental/Control sounds was randomised by block. B) The slow-wave detection algorithm for the up (yellow) and down (purple) transition areas. C) Phase angle analysis at stimulus release. The left polar histogram represents the experimental sounds with Up condition in yellow and Down condition in purple. Similarly, the right histogram represents the Control sounds. Both graphs are the values for electrode F3.

observed a similar pattern in the spindle frequency band, where the power was lower for experimental compared to control sounds. This cluster was larger, ranging from 1350 to 2400ms and involving 58 channels (p=0.0430, see Figure 6(A) and (C)). These temporally late differences are in line with previous work. For instance, one study[37] showed a difference between old cues and controls at 1700-2300ms from stimulus onset on the fast spindle frequency band. Similarly, another study[22] found a significant cluster lasting from 850 to 1750ms also in the spindle band when comparing recall performance in a CL-TMR set-up. Interestingly, we found no significant difference between experimental and control sounds during Down state stimulation, with no clusters for theta band and a lowest probability of p=0.28 for the spindle band (see Figure 6(D-F)). Finally, we directly contrasted the interaction between Up (experimental vs control sounds) and Down (experimental vs control sounds) (see Figure 6(G)). No significant clusters emerged for theta (p=0.070, 6(H)) but a similar cluster to that seen in the Up state did emerge for the spindle band (p=0.0160, Figure 6(I)), also involving 58 channels and likely driven by the Up-condition cluster. We also performed a similar analysis in a broader band (4 to 20Hz), see results in SM5.

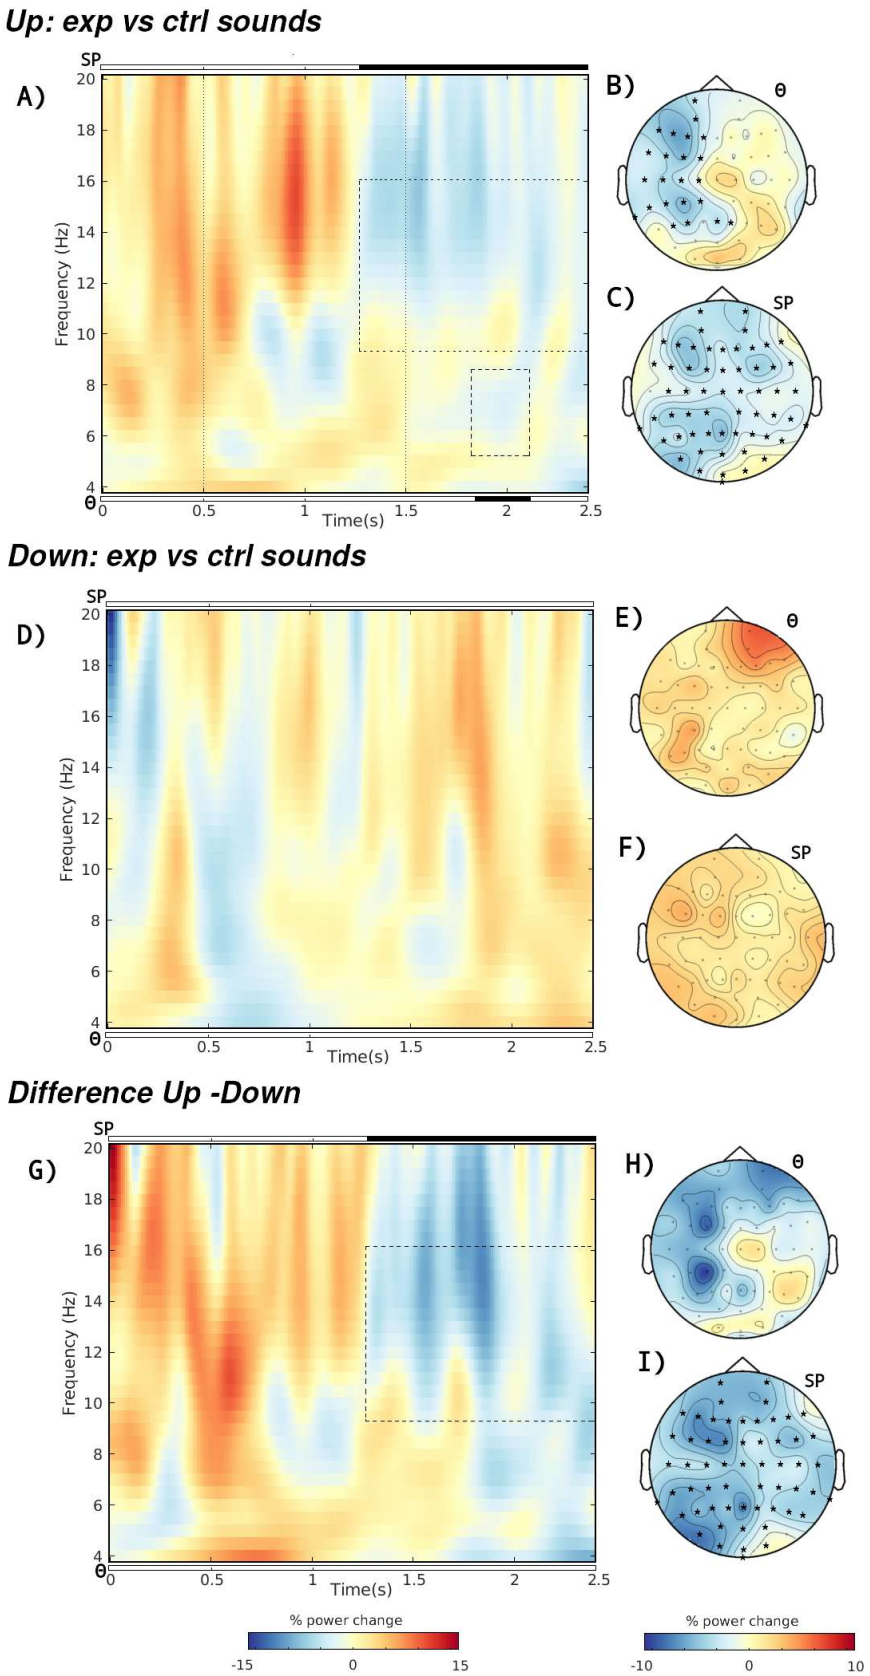

Figure 6. Time-frequency results. Contrast between experimental and control sounds for (A) Up stimulation, (D) Down stimulation and Up-Down interaction (G). Black bars indicate time of significant clusters in the time-frequency analysis, with spindles band (SP) on top and theta band (θ) on the bottom of each graph. That time window was used to plot the topographies on the right side for the Up (B, D), Down (E, F) and Interaction (H,I). The topographies show the mean averaged (over time) power spectrum within the significant cluster for the theta (θ, top) and spindle band (SP, bottom) for each case. Channels involved in the significant cluster are highlighted with a star.

2.3 Detection of reactivation

To determine if memory-related neural activity would be distinctively reactivated during the night, we trained two machine learning algorithms to differentiate between experimental and control sounds for each condition (Up/Down), refer to Section 4.8 for further details.

After calculating the classification ability for each participant, we performed cluster statistics at the group level for each condition. Only the Up condition presented a significant cluster of above-chance classification. This cluster (p=0.037) ranges from 1204 to 1298ms relative to the stimulus onset as shown in Figure 7 for the SVM classifier with AUC as metric. Similar values were obtained for the rest of the tested combinations, refer to table SM7.1 for full details. There were no significant clusters for the Down condition (lowest p=0.074), refer to table SM7.2 for each classifier result and figure SM7.1 for a grand average plot of the performance of both conditions. This is in line with the time-frequency analysis results where only the Up condition showed differences between experimental and control sounds. The absence of difference between these sounds could indicate that the brain is not able to activate the corresponding memory trace above chance when the cues are presented in the down state. Hence, no behavioural benefit is obtained the following morning. However, such speculation does not explain the performance improvement of the participants two weeks later.

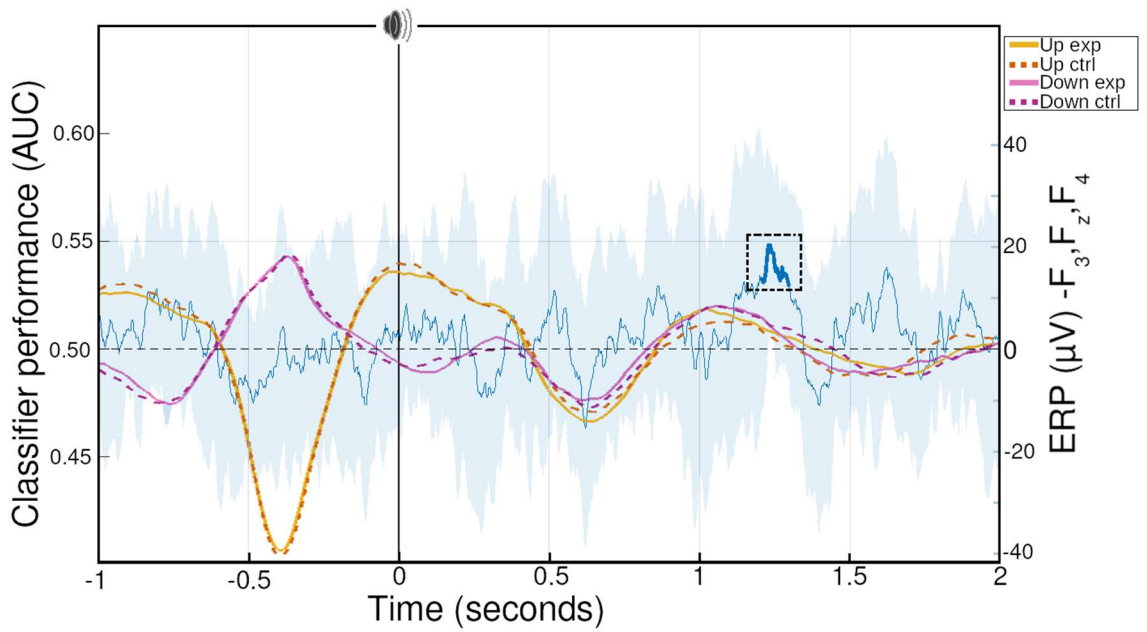

Figure 7 Classification performance. SVM classifier (blue line) presented a statistically significant cluster centred around 1.3 seconds (thicker line). On top are superimposed the grand average ERP values for Up (yellowish colours) and Down (purple colours) separated by experimental (solid lines) and control (dotted lines) sounds.

To better understand the relationship between replay and consolidation, we performed a series of correlations between classification performance within the significant cluster (mean and peak) and the behavioural metrics described above. This revealed a significant negative

correlation (R=-0.65, p_{unc}=0.002, p_{corr}=0.022 for the SVM+AUC combination) between mean classification and behavioural performance on the second degree of separation of the inference pairs for the morning after sleep (Session 2), see Figure 8. This result fits well with the fact that only the second-degree separation pair of the Up condition presented a significant difference in overnight improvement when compared with the control condition for Session 2, and similarly only Up showed differences between the degree of separation on inference pairs between conditions. The correlation was constant across all tested machine learning algorithms, see table SM7.4 for the full list. No other correlation with any behavioural metric was significant before correction (all p>0.05). However, the fact that this correlation is negative, e.g., the better the algorithm can classify the worse participants performed in that particular pair is perhaps surprising. Similar negative correlations have been reported before, e.g. between classification performance and behavioural metrics[38] and between EEG-microstates and behaviour[39].One possible hypothesis, it is these "good" performers try to fit the control sound into the previously-learned hierarchy. That is, when they hear the control sound they may replay the cued hierarchy again, making it more difficult for the classifiers to differentiate between control from learned sounds.

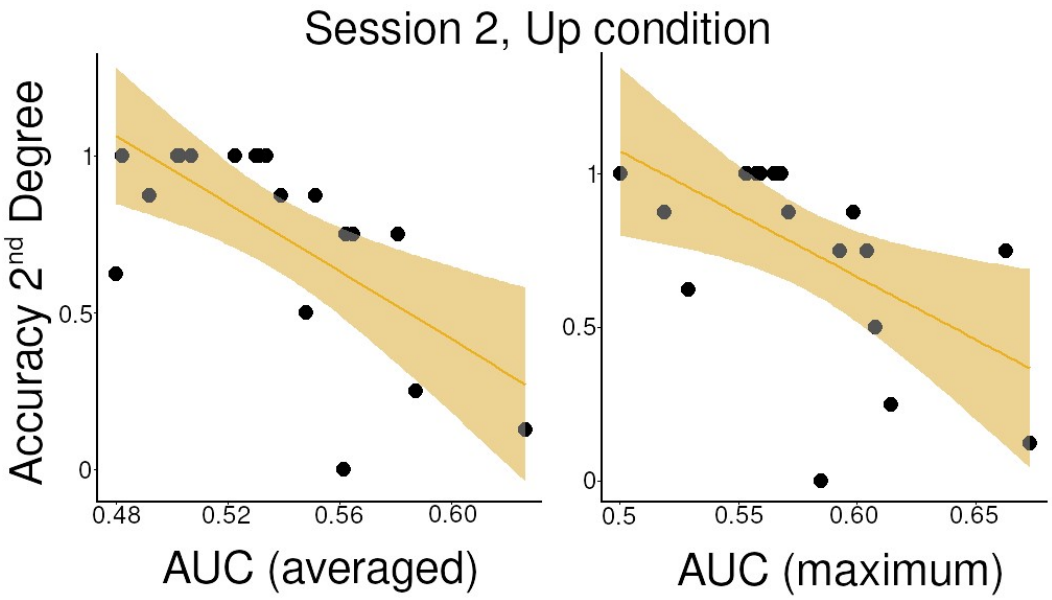

Figure 7 Correlations between the SVM+ AUC mean -left- and peak performance -right- of the significant cluster for the Up condition with the behavioural accuracy of the second degree of difference-inference pair of the Up condition. Shadow areas represents 95% confident intervals.

2.4 Awareness

When participants are aware of the linear hierarchy, explicit logical reasoning can explain their performance[40]. It is therefore important to assess awareness. We did this in two different ways,

using confidence ratings and through a questionnaire (See SM6 for full results). Neither of the methods revealed an explicit awareness of the hierarchy, hence we can discard the idea that the participant's ability of make inferences is due to this phenomenon.

3 Discussion

This paper explores the roles of both time and sleep in the development of relational memory. Previous studies have demonstrated that an offline period of sleep facilitates the evolution of inferential abilities from knowledge of premise pairs[9,10]. Here, we go further by manipulating the replay of premise pairs through auditory TMR in sleep in order to investigate the impact of such replay on subsequent ability to make inferences. We also investigate the differential influence of applying this TMR at different SO phases. Only cues presented during the down-to-up transition of the SO were associated with detectable memory replay. Such cues also led to better knowledge of the most difficult inference pairs compared to control. By contrast, cueing during the up-to-down transition was not associated with measurable replay, and did lead to significantly impaired performance next morning. The difference between cueing in Up and Down conditions was also present in the oscillatory fingerprint, with significantly different traces for the two phases (see Figure 6). Overall, our data suggest that the up-going phase of the SO may represent a prime window for triggering memory replay and boosting performance on distant inferences immediately after sleep[22].

Phase of the slow oscillations and memory

Numerous studies have demonstrated that auditory stimuli delivered in SWS are more likely to have a positive impact when they occur during SO up-states. This is true for both CL-TMR[22,23] and closed loop auditory stimulation (CLAS)[41,42]. The concept of an optimal phase or window for TMR stimulation supports the active system consolidation framework[15,22] under which SOs serve to synchronise memory consolidation across different brain structures. Depolarising up-states are more likely to activate larger groups of neurons synchronously[24] and to drive thalamocortical spindles and sharp-wave ripples in the hippocampus[43], hence facilitating hippocampal memory reactivation. This could explain why our classifiers only detect memory replay after Up stimulation, as well as the next-morning behavioural benefit of the Up condition.

While we hypothesised that stimulating the Up condition is more likely to trigger memory replay, and indeed our classifiers only detect replay after Up condition TMR, we do not mean

to suggest that replay never occurs during the down-states. One study[43] found that the optimal stimulation phase was the SO down-state. Thus, TMR cues presented during the first half of the down state were associated with lower forgetting. However, our results show a clear difference between cueing in Down and Control conditions and between Down and Up conditions for the inference pairs the next morning, both when data were averaged and when they are sorted by degree of separation. These behavioural results are strongly supported by our EEG and classifier analyses which show evidence of memory replay after TMR in Up but not Down. In fact, we found no significant differences between TMR and control sounds for the Down condition, suggesting that memory replay due to TMR cueing may be blocked during the up-to-down transition. However, differences in behavioural tasks, TMR sound lengths and CL-TMR algorithms could all potentially explain these discrepant findings.

TMR time course

In recent years many studies have suggested the importance of sleep-based memory benefits but not many have tested whether such benefits are long-lasting[44]. One study[45] demonstrated that a brief period of sleep immediately following learning enhanced memory for emotional text after four years compared with wakefulness. Another study found that taking a nap between encoding and retrieval enhanced memory to the same extent as when the same time was spent cramming, but after a week the benefits were only present in the nap group not in the cramming group[29]. Turning specifically to TMR effects, one study reported that TMR benefit to reducing implicit bias was retained after a week[26], while another study reported TMR benefit to a serial reaction time task peaked after ten days[25]. Overall, these data seem to suggest that the benefits of sleep can be long lasting, and can even increase gradually over several days. In the present study, accuracy for inference pairs in the Up and Control conditions remained at the same level two weeks later as it was right after sleep (Figure 3). This is likely due to a more general effect of sleep rather than to TMR, since the Control condition was not stimulated during the night, and Up was not statistically different from Control at the morning session. However, and against our initial hypothesis, behavioural performance in the Down condition increased from under chance level to over chance level, reaching similar overall accuracy to the other two conditions after two-weeks (Figure 3(A-right)). This result suggests that despite an initial inhibition because of TMR in the Down state, the brain was able to recover over time. We hypothesised that subsequent nights of sleep without manipulation may have allowed spontaneous replay of the memories equally for all conditions. While Up and Control conditions, which had replayed successfully in the first night, derived no benefit from this

additional replay, the Down condition did benefit, and the associated consolidation allowed this hierarchy to catch up with the other hierarchies. Additionally, the next morning test in Session 2, which included presentation of inference pairs, may also have helped to trigger subsequent replay from which Down condition benefited.

Premise versus inference pairs

The ability to make transitive inferences has been shown in several animal species, suggesting that it may have some form of evolutionary relevance. A range of models has been proposed to explain the mechanisms underlying this task across species (see[8] for a summary). Some of these focus on the preference of B over D in order to examine inference pairs. This occurs in both associative[46] and conceptual models[5] and the key question is how participants are able to identify the linear order and select B over non-adjacent D when only adjacent pairs (A-B, B-C and C-D) of a hierarchy, i.e. A>B>C>D have been presented. Other models argue that transitive inference relies on cognitive processes, considering that the preference of B over D is not based on logic or reasoning, that is rejecting the idea of "inference" per se[2,47]. Determination of which of the existing models is more appropriated to describe TI in humans is outside the scope of this work, but we can argue based on our results that premise and inference pairs rely on different processes. It is true that only premise pairs were used for training, but this is unlikely to explain the observed differences. Thus, unlike inference pairs, premise pairs showed no benefits from either TMR or sleep, with no overnight improvement for any condition. Furthermore, stimulating in the Up condition led to a difference between the first and second degree of separation in inference pairs. This was not present in either of the other two conditions and remained after two weeks. Furthermore, both degrees of inference separation were performed better in the Up than the Down condition (Figure 3(B)) in the morning session. These results suggest underlying differences in relational memory between inference and premise conditions and between inference pair types. In contrast to previous studies such as[9], this is a within-subject design so these differences cannot be explained by circadian influences or any other similar effect.

In sum, our results show that the complex process of making indirect inferences can be facilitated by cueing replay during slow wave sleep. Importantly, this was only true for the most difficult inference pairs, and only cueing during the down-to-up phase was effective. Also, these effects were not transitory, but instead lasted for the full two-week period examined here. These results provide strong support for the idea that memory replay in sleep is important for high-level qualitative changes in memory, such as integration and relational

memory. Our findings also hold promise for the application of sleep-based interventions to drive improvement in such complex memory, and its application to real-word tasks.

4 Methods

4.1 Participants

Thirty adults (10 males, mean age 27 ±3.72) participated in the overnight experiment. All of them with no-self reported history of neurological, psychiatric, sleep or motor disorders. All participants completed a screen questionnaire before selection, provided written informed consent, and were reimbursed for their time. The experiment was approved by the School of Psychology Ethics Committee at Cardiff University. All participants agreed to abstain from caffeine and alcohol during the study and for 24-hours before.

4.2 Materials

The behavioural tasks were presented in a quiet room, participants were comfortably seated in front of the computer and stimuli were presented using Matlab©, Psychtoolbox[48] and Cogent 2000 (www.vislab.ucl.ac.uk/cogent.php). Three types of visual stimuli were presented to the participants: female faces[49], outdoor scenes (taken from internet) and unusual objects[50], see Figure 1(B). Each stimulus was easily distinguishable from the others within and between categories. All of them were presented in grey scale and matched for luminance.

Each image was associated with an exclusive sound, congruent with the image, to be able to use them as TMR stimuli during the sleep part of the experiment. Sounds were taken from the internet and truncated into two different lengths, 2 seconds and 200ms, and pitch normalised. We used the longer sound in behavioural training, facilitating sound-image association. The shorter version of the sound was used for the behavioural testing and for TMR during sleep. The tones were played through noise-cancelling headphones (Sony MDR-ZX110NA) during the behavioural tasks and through speakers (Dell A225) during sleep.

The order of presentation of each stimulus category was counterbalanced across participants and the order of stimuli within each category was completely randomised for each subject. Hence, the experimenter was completely blind to which stimuli form the hierarchy and its order within each condition (Up, Down, Control), and also which type of stimuli (faces, objects or scenes) was selected for each condition, so as not to influence the results.

Each of the three hierarchies (faces, object and scenes) comprised 6 images, each one with an associated (highly discriminable) sound. We prepared a set of 12 images and 12 sounds per hierarchy, that is a total of 36 images and sounds. At the beginning of the

experiment, for each one of the three hierarchies, 6 of these images with their corresponding sounds were selected to be learned and the remaining 6 sounds were used as controls to be played during the TMR stimulation. Before participants started the first task, the 6 images which would be used during the experiment and the other 6 which would be used as control were randomly selected for each category (blind for the experimenter). The control images were never used, but the control sounds were played during TMR as explained in section 4.5.

4.3 Procedure

Participants arrived at the laboratory at 7 or 8pm and changed into their sleepwear. They reported their alertness by completing the KSS[51] and SSS[52] questionnaires. Afterwards, they were fitted for PSG recording and performed the initial training and the immediate test explained in 4.4 and Figure 1(A). Participants were ready for bed around 11pm. During the night, the previously learned tones were replayed during SWS. From the three stimulus categories, one was kept as a control (Control) and was not replayed during the night, allowing us to compare it against the other two cued during the Up and Down conditions respectively.

After 7-8 hours of sleep, participants were woken at the agreed time and allowed at least 20 min to overcome sleep inertia. During this time, they were given time to go to the toilet, eat and drink before completing the sleep quality, KSS and SSS questionnaires. Participants then completed the Late test and another Sound-Image association test (Section 4.4 and Figure 1(A)). Afterwards, the electrodes were removed, and participants were allowed to have a shower or go home.

Finally, participants came back to the laboratory two weeks later (± 2days) to complete a second Late test and Sound-Image association test, identical to the previous one, in order to test how robust the sleep-TMR mediated benefits of the task were. However, this time no EEG recordings were performed.

From the thirty participants that completed the task, 10 were eliminated because of technical problems (n=3) or because they did not have enough stable SWS (n=7) to perform enough stimulations (we required 12 rounds, these participants were mostly in light sleep-N2 stage). From the 20 participants left, 3 could not finish Session 3 due to the pandemic. We therefore included 20 participants for the first two sessions and only 17 for the last one.

4.4 Experimental tasks

The experiment was composed of three sessions: evening (Session 1), next morning (Session 2) and a follow up session two-weeks after (Session 3). Each one of the sessions was divided into different task as explained below.

4.4.1 Session 1

This will be performed in the evening after participants changed into their sleepwear and were fitted with the EEG cap.

Sound-Image association learning: In the first session, participants started with a training period. Firstly, they were asked to perform a sound-image association task. For each one of the three categories, participants were shown one by one each of the six items forming the category. At the same time the associated sound was played. Each sound-image pair was shown to the participants 4 times. Each category was presented to the participant in order, for instance, 6 scenes followed by 6 faces and then 6 objects. However, the items within category were presented randomly and the order of the categories was counterbalance across participants to avoid any bias. The sounds for this training session were two seconds long to facilitate the learning.

Sound-Image association test: Immediately after the training, all participants performed a test session to determine their retention level for such associations. Three images were presented on the screen and a sound was played simultaneously. Participants were asked to select as quickly and accurately as possible the image corresponding to the sound using the arrow keys of the keyboard. Straight after they pressed the corresponding key, visual feedback appeared on the screen in the shape of a rectangle surrounding the image corresponding with the sound, e.g., the right answer. The rectangle was green if the participant's selection was correct or red if it was wrong. The aim of adding feedback during training is to strengthen the sound-image association.

The position of the images on the screen was randomised every trial. The three images presented on the screen were pseudo-randomly selected, with the restriction that at least one of the two images was a 'lure' of the same category as the right answer. This time the sounds were only 200ms long. This was done, once more with the TMR sleep experiment in mind, to check that shorter versions of the sounds were still identifiable for the participants. Participants performed two blocks with three repetitions of each sound per block. At the end of each block, accuracy was presented to keep participants engaged with the task.

Pairs Learning: All participants learned five relational premise pairs of each one of the categories, following previous sleep related experiments[9,10]. If each category formed a 6-item hierarchy, schematically represented as A>B>C>D>E>F (see Figure 1(C)), the premise pairs would be: A>B, B>C, C>D, D>E, E>F. Where the notation "A>B" indicates "select A over B".

The pairs were presented on the screen one at a time, with images positioned vertically on top of each other (Figure1.E). Subjects were instructed to select the item "hiding" a smiley emoticon from the two presented on the screen, at first by trial and error, but after practice and feedback they were be able to learn which of the two items was correct. If the participant selected the correct item, the chosen item was replaced by a smiling-face (smiley emoticon). This is in line with[11], where a smiling-face was used as reinforcement. If the participant selected the wrong member of the pair, this item was replaced by an angry emoticon. After the feedback, participants received a second reinforcement as the pair was presented again but this time horizontally instead of vertically, and in the correct order (e.g. A-B) from left to right, with the corresponding sounds also played in the correct order at the same time. We intentionally avoid presenting the sounds from the beginning of the trial to prevent any misleading effects. For example, if hearing the sounds in the wrong order might have impacted upon the TMR results.

Items were organised into blocks. Each block contained 10 trials of each one of the hierarchies or categories. This meant a total of 30 trials per block were used. Each block presented each one of the five pairs of each hierarchy twice, counterbalancing the up-down positions (e.g A above B and B above A, being A the correct selection in both cases) to avoid any location effects. Pairs of the three hierarchies were not mixed within a block. For example, first all pairs for the “scenes” category were presented, then pairs in the “faces” category, and finally the “object” pairs. This order, as previously, was counterbalanced across participants. Within each category, pairs were presented pseudo-randomly to avoid explicitly revealing the hierarchy. Hence, a new pair cannot contain an item that was in the previous pair (e.g., A>B will never be followed by B>C). Furthermore, the order of the items within the hierarchy was randomly selected for each participant at the start of training, remaining unknown to the experimenter. At the end of each block, the overall performance for that block was shown on the screen to keep participants engaged with the task.

All subjects underwent a minimum set of three blocks of training. After the third block, and every block thereafter, only performance of the "middle pairs", meaning B-D, B-E and C-E, was saved to calculate the exit criteria. If the averaged performance of these pairs for two of

the last three blocks was over 66% for one of the hierarchies, the participant stopped receiving feedback for that particular hierarchy. However, all the premise pairs of this category still appeared on the screen to ensure we had the same number of trials/appearances for each hierarchy. This continued until the participant reached the criteria for each one of the three hierarchies or a maximum of 10 blocks. Either way, the program automatically stopped, hence the test was finished. In contrast to[11] and in[10] where the exit criterion was set to 75% accuracy for the middle premise pairs, we used a criterion of 66% not only to avoid ceiling effects but also to increase the chances of overnight improvement. On the other hand, we added a more restrictive criterion of 2 blocks out of 3 meeting the threshold, to be sure that the criterion was not achieved purely by chance. Similarly, to the above-mentioned studies, we only counted the middle premise pairs to evaluate the exit threshold because they are the necessary items for building the inferences.

Immediate test After the criterion was met, participants enjoyed a minimum of 5 minutes break before proceeding to perform the immediate test. This test aimed to assess the initial retention level of the learned pairs. A similar protocol was used for the testing and the training part with the exception that the feedback and sound cues were removed. Subjects were informed about this change and that they must select the right item based on the previous learning.

Participants performed a total of four blocks, with 10 trials per hierarchy. In between blocks, participants had to solve two easy arithmetic problems as a distractor tasks, clearing their short-term memory[53]. Furthermore, for this immediate test, the pairs from the different hierarchies were randomly interleaved, always with the restriction of not showing the hierarchy explicitly, as in the learning phase.

4.4.2 Session 2

Session 2 was undertaken in the morning after participants received the overnight stimulation (TMR) of two of the three hierarchies.

Late test: After filling the KSS and SSS questionnaires participants performed a similar test as before but this time, they were presented with previously learned premise pairs, new inference pairs, and one ‘anchor pair’ such that a total of 9 pairs were seen instead of just the 5 previously learned. The first new pairs were 3 inference pairs: B>D, B>E and C>E (see Figure 1(D)). These pairs are named inference pairs because if you know that B>C and C>D, then you can infer that B>D. The inference pairs can be divided into 1st and 2nd degree of separation. That refers to the number of items between the ‘pair’ items, for instance between B and D is

only one item (C), hence it has first degree separation, as does C-E. On the other hand, there are two items between B and E: C and D. Hence the B-E pair has a second degree of separation, and is thus the most difficult pairs to infer within a 6-item hierarchy. Additionally, we also added a 4th pair, the anchor pair (A>F) as a control since inference is not needed to obtain this relationship. This is due to the simple fact that A is always correct and F is always incorrect[53].

Participants were instructed that they might see novel combinations and if that was the case, they should try to make their best guess about them. At the end of each trial participants were asked how confident they were of their answer in a scale ranging from -2 (guessing) to +2 (certain) using the same up and down arrows of the keyboard used for selecting the up or down images. Following a similar protocol to the Immediate test, participants performed four blocks with maths exercises between them.

Sound-Image association test: Following the delayed test, and after a 5-minute break, subjects performed a new sound-image association test. The structure was similar to the test in the session 1 with the difference that the feedback was eliminated.

4.4.3 Session 3

This follow-up session consisted of the same tasks in the same order as in Session 2: KSS and SSS questionnaires, Late test, Sound-Image association. However, this time participants' brain activity was not recorded. Finally, the participants performed an Awareness questionnaire.

Awareness questionnaire: Participants at the end of the third session's tests answered a questionnaire about their awareness of the hierarchies and their strategies to remember the different items and the sound-image associations (see SM2.2). These questions were presented on paper, with one question per page and participants were asked not to turn the page until the previous question was completed. Additionally, they were asked not to go backwards in the document and change their answers. The list of questions was adapted from[8,47].

4.5 TMR protocol

The two categories that would be used for TMR and the one acting as a control were counterbalanced across the 30 participants (see SM3 for final count). The control category was not cued during the night. From the other two, one was assigned to the up-state cueing and the other to the down state cueing condition. Stimulation of the respective associated

sounds started after the participants entered stable slow wave sleep and was halted as soon as arousals or any other sleep stage were detected.

For each of the two TMR-categories, Up and Down conditions, to the hierarchy participants learned during the training part, an extra control-hierarchy of sounds were added. These control-hierarchy-sounds, also 200ms duration, were completely novel to the participants and were included to help us to identify the TMR effect from the normal sleep-brain activity. Each one of the hierarchies, experimental and control, was composed of 6 items and played in order: A, B, C, ..., F. The order of the control and experimental hierarchies were randomised and counterbalance across blocks. Similarly, the order of the Up and Down cueing was also counterbalanced across blocks. Each block was formed from four hierarchies: experimental Up, control Up, experimental Down, control Down (see Figure 4(A)). The minimum criterion number of blocks to include a participant in the analysis was 12, this means 288 cues were performed during the night.

The online detection of slow oscillations was based on the detection of the negative half-wave peaks of oscillations. The electrode used as reference for the on-line detection was F3. When the amplitude of the signal passed a threshold set at -80uV the auditory stimulus was delivered[41,42]. Inter-trials intervals were set to a minimum of 4 seconds. The slow oscillation detection, auditory stimulation and presentation of the trigger to the EEG recording was done via a custom-made Matlab-based toolbox (https://github.com/mnavarretem).

4.6 Statistics

Statistical behavioural results were calculated using robust statistical methods from the R package WRS2[54] to avoid any possible issues with normality and homoscedasticity assumptions. Repeated measures analysis of variance (RM-ANOVA) or simple ANOVA analysis was performed accordingly, keeping the trial information and adding individual differences into the analysis (subject ID's). One sample t-tests (Students or Wilcoxon signed-rank) were performed to test the difference over chance level (50 %) of each group, condition and session of interest. For the accuracy of the TMR protocol, the R package Circular was used[55]. For time-frequency analysis, we tested the difference between experimental and control sounds for the Up and Down conditions and its interaction with a cluster-based permutation test with dependent samples, cluster level alpha of 0.050 and 1000 random data partitions for the Monte Carlo p-value calculations. The statistical significance of Pearson correlations between classification accuracy and behaviour was using a bootstrap method implemented in R (boot package[56]).

4.7 EEG recordings

Sleep was recorded using standard polysomnography including EEG, electromyographic (EMG) and electrooculography (EOG). EEG was recorded using a 64-channel LiveAmp amplifier (Brain Vision). Electrode impedance was kept below 10KΩ and sampling rate was set to 500Hz. Initially referenced to Cpz electrode. In addition to the online identification of sleep stages, polysomnographic recordings were scored offline by 3 independent raters according to the ASSM criteria[57], all of them were blind to the periods when the sounds were replayed.

4.7.1 Pre-processing

EEG pre-processing was performed using FieldTrip[58] and custom Matlab functions. Data were low-pass and high-pass filtered (30Hz and 0.5Hz respectively). Eye and muscle related artefacts were removed using independent component analysis (ICA). Bad channels were interpolated and data epoched into 4 second segments, from -1s before stimulus onset to 3s afterwards. Finally, a visual inspection of the dataset was performed, and any residual artefact was manually removed. Afterwards, each epoch was categorised into Experimental Up, Control Up, Experimental Down, Control Down conditions.

4.8 Classification

Classification of single-trial data was performed using MVPA-light[59]. Classification was performed for each participant and each time point (-1 to 3s) using the ERP values (4 to 20Hz) of the 60 EEG channels as features. In order to corroborate our results, we compared the performance of two classifiers, a linear discriminant analysis (LDA) and a support vector machine with linear kernel (SVM). To avoid overfitting, we used a 5-fold cross-validation method with 2 repetitions and principal component analysis (PCA) as a dimensionality reduction technique (n=10, 20 and 30). The data within each fold was z-scored to avoid any bias. Additionally, we used two different metrics to evaluate the performance of each classifier: the traditional accuracy (ACC) and area under the curve (AUC). The former is defined as the number of correctly detected trials between the total number of trials, the later represents the trade-off between the true positive rate and false positive rate. Once the classifiers were calculated for each participant, we performed a between-subject cluster analysis[59] to determine at what time points the experimental and control sounds were statistically different

for each condition. All code for the analysis of this study is available at https://github.com/Contrerana.

Acknowledgements

We would like to thank Elena Schmidt, Ralph Andrews and Duarte Pereira for contributing to the data collection, Dominic Carr for helping with the sleep scoring and the NAPs lab for helping out in many ways.

Author contributions

PL conceived of the present idea and LS and PL designed the experiment. IK recruited the participants and IK, NM and LS collected the data. LS analysed the data. All authors discussed the results and contributed to the final manuscript.

References

1. Dymond, S. & Llewellyn, S. Time, Sleep, and Stimulus Equivalence-Based Relational Memory. *Psychol. Rec.* **69**, 357–365 (2019).
2. Lerner, I. & Gluck, M. A. Sleep and the extraction of hidden regularities: A systematic review and the importance of temporal rules. *Sleep Medicine Reviews* (2019) doi:10.1016/j.smrv.2019.05.004.
3. BRYANT, P. E. & TRABASSO, T. Transitive Inferences and Memory in Young Children. *Nature* **232**, 456–458 (1971).
4. Grosenick, L., Clement, T. S. & Fernald, R. D. Fish can infer social rank by observation alone. *Nature* **445**, 429–432 (2007).
5. Lazareva, O. F., Paxton Gazes, R., Elkins, Z. & Hampton, R. Associative models fail to characterize transitive inference performance in rhesus monkeys (Macaca mulatta). *Learn. Behav.* (2020) doi:10.3758/s13420-020-00417-6.
6. Camarena, H. O., García-Leal, O., Burgos, J. E., Parrado, F. & Ávila-Chauvet, L. Transitive Inference Remains Despite Overtraining on Premise Pair C+D-. *Front. Psychol.* **9**, (2018).
7. Holyoak, K. J. & Lu, H. Emergence of relational reasoning. *Current Opinion in Behavioral Sciences* (2021) doi:10.1016/j.cobeha.2020.11.012.
8. Morgan, A. *Effect of Time and Sleep on the Transitive Inference Task*. (2017).
9. Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D. & Walker, M. P. *Human relational memory requires time and sleep*. www.pnas.org/cgi/content/full/ (2007).
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. Sanders, K. E. G. How 'sleeping on it' works: The cognitive mechanisms and neural signatures of sleep-facilitated problem solving. *Dissertation Abstracts International: Section B: The Sciences and Engineering* (2020).
12. Lutz, N. D., Diekelmann, S., Hinse-Stern, P., Born, J. & Rauss, K. Sleep Supports the Slow Abstraction of Gist from Visual Perceptual Memories. *Sci. Rep.* **7**, 42950 (2017).
13. Hennies, N., Ralph, M. A. L., Kempkes, M., Cousins, J. N. & Lewis, P. A. Sleep spindle density predicts

the effect of prior knowledge on memory consolidation. *J. Neurosci.* (2016) doi:10.1523/JNEUROSCI.3162-15.2016.
14. 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.* (2010) doi:10.1523/JNEUROSCI.3028-10.2010.
15. Rasch, B. & Born, J. About Sleep's Role in Memory. *Physiol. Rev.* **93**, 681–766 (2013).
16. 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.* (2020) doi:10.1037/bul0000223.
17. Nolan, B. REM, not incubation, improves creativity by priming associative networks. *Yearb. Neurol. Neurosurg.* **2010**, 191–192 (2010).
18. Fogel, S. M. & Smith, C. T. Learning-dependent changes in sleep spindles and Stage 2 sleep. *J. Sleep Res.* (2006) doi:10.1111/j.1365-2869.2006.00522.x.
19. Wilhelm, I. *et al.* Sleep Slow-Wave Activity Reveals Developmental Changes in Experience-Dependent Plasticity. *J. Neurosci.* **34**, 12568–12575 (2014).
20. Steriade, M., Nuñez, A. & Amzica, F. A novel slow (< 1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. *J. Neurosci.* **13**, 3252–65 (1993).
21. Steriade, M. Grouping of brain rhythms in corticothalamic systems. *Neuroscience* **137**, 1087–1106 (2006).
22. Göldi, M., van Poppel, E. A. M., Rasch, B. & Schreiner, T. Increased neuronal signatures of targeted memory reactivation during slow-wave up states. *Sci. Rep.* **9**, 2715 (2019).
23. Shimizu, R. E. *et al.* Closed-Loop Targeted Memory Reactivation during Sleep Improves Spatial Navigation. *Front. Hum. Neurosci.* **12**, (2018).
24. Vyazovskiy, V. V. *et al.* Cortical Firing and Sleep Homeostasis. *Neuron* (2009) doi:10.1016/j.neuron.2009.08.024.
25. Rakowska, M., Abdellahi, M. E. A., Bagrowska, Pa., Navarrete, M. & Lewis, P. A. Long term effects of cueing procedural memory reactivation during NREM sleep. *Neuroimage* **244**, 118573 (2021).
26. Hu, X. *et al.* Unlearning implicit social biases during sleep. *Science (80-.).* (2015) doi:10.1126/science.aaa3841.
27. Groch, S. *et al.* Targeted reactivation during sleep differentially affects negative memories in socially anxious and healthy children and adolescents. *J. Neurosci.* (2017) doi:10.1523/JNEUROSCI.1912-16.2017.
28. Abel, M. *et al.* Sleep reduces the testing effect—But not after corrective feedback and prolonged retention interval. *J. Exp. Psychol. Learn. Mem. Cogn.* **45**, 272–287 (2019).
29. Cousins, J. N., Wong, K. F., Raghunath, B. L., Look, C. & Chee, M. W. L. The long-term memory benefits of a daytime nap compared with cramming. *Sleep* **42**, (2019).
30. Schapiro, A. C. *et al.* Sleep Benefits Memory for Semantic Category Structure while Preserving Exemplar-Specific Information. *Sci. Rep.* (2017) doi:10.1038/s41598-017-12884-5.
31. Denis, D. *et al.* Sleep spindles preferentially consolidate weakly encoded memories. *J. Neurosci.* (2021) doi:10.1523/JNEUROSCI.0818-20.2021.
32. von Fersen, L., Wynne, C. D. L., Delius, J. D. & Staddon, J. E. R. Transitive Inference Formation in Pigeons. *J. Exp. Psychol. Anim. Behav. Process.* (1991) doi:10.1037/0097-7403.17.3.334.
33. Kuriyama, K., Stickgold, R. & Walker, M. P. Sleep-dependent learning and motor-skill complexity. *Learn. Mem.* (2004) doi:10.1101/lm.76304.
34. Sio, U. N., Monaghan, P. & Ormerod, T. Sleep on it, but only if it is difficult: Effects of sleep on problem solving. *Mem. Cogn.* (2013) doi:10.3758/s13421-012-0256-7.
35. Schreiner, T., Lehmann, M. & Rasch, B. Auditory feedback blocks memory benefits of cueing during sleep. *Nat. Commun.* **6**, 8729 (2015).
36. Schreiner, T. & Rasch, B. Boosting vocabulary learning by verbal cueing during sleep. *Cereb. Cortex* (2015) doi:10.1093/cercor/bhu139.
37. Cairney, S. A., Guttesen, A. á. V., El Marj, N. & Staresina, B. P. Memory Consolidation Is Linked to Spindle-Mediated Information Processing during Sleep. *Curr. Biol.* (2018) doi:10.1016/j.cub.2018.01.087.
38. Abdellahi, M. E. A., Koopman, A. C. M., Treder, M. S. & Lewis, P. A. *Targeting targeted memory reactivation: characteristics 2 of cued reactivation in sleep*. (2021) doi:https://doi.org/10.1101/2021.12.09.471945.
39. Murphy, M., Stickgold, R., Parr, M. E., Callahan, C. & Wamsley, E. J. Recurrence of task-related electroencephalographic activity during post-training quiet rest and sleep. *Sci. Rep.* (2018) doi:10.1038/s41598-018-23590-1.
40. Frank, M. J., Rudy, J. W., Levy, W. B. & O'Reilly, R. C. When logic fails: Implicit transitive inference in

humans. *Mem. Cognit.* **33**, 742–750 (2005).
41. Ngo, H. V. V., Martinetz, T., Born, J. & Mölle, M. Auditory closed-loop stimulation of the sleep slow oscillation enhances memory. *Neuron* (2013) doi:10.1016/j.neuron.2013.03.006.
42. Ngo, H. V. V. *et al.* Driving sleep slow oscillations by auditory closed-loop stimulation—A self-limiting process. *J. Neurosci.* (2015) doi:10.1523/JNEUROSCI.3133-14.2015.
43. 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.* (2016) doi:10.1523/JNEUROSCI.3175-15.2016.
44. Cordi, M. J. & Rasch, B. How robust are sleep-mediated memory benefits? *Curr. Opin. Neurobiol.* **67**, 1–7 (2021).
45. Wagner, U., Hallschmid, M., Rasch, B. & Born, J. Brief Sleep After Learning Keeps Emotional Memories Alive for Years. *Biol. Psychiatry* (2006) doi:10.1016/j.biopsych.2006.03.061.
46. Jensen, G. *et al.* Category learning in a transitive inference paradigm. *Mem. Cogn.* (2021) doi:10.3758/s13421-020-01136-z.
47. Frank, M. J., Rudy, J. W., Levy, W. B. & O'Reilly, R. C. When logic fails: Implicit transitive inference in humans. *Mem. Cogn.* (2005) doi:10.3758/BF03195340.
48. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 433–436 (1997).
49. Lundqvist, D. & Flykt, A. The Averaged Karolinska Directed Emotional Faces - AKDEF. *CD ROM from Dep. Clin. Neurosci. Psychol. Sect. Karolinska Institutet* (1998).
50. Horst, J. S. & Hout, M. C. The Novel Object and Unusual Name (NOUN) Database: A collection of novel images for use in experimental research. *Behav. Res. Methods* **48**, 1393–1409 (2016).
51. Åkerstedt, T. & Gillberg, M. Subjective and objective sleepiness in the active individual. *Int. J. Neurosci.* (1990) doi:10.3109/00207459008994241.
52. Hoddes, E., Dement, W. & Zarcone, V. The development and use of the Stanford sleepiness scale (SSS). *Psychophysiology* (1972).
53. von Hecker, U., Klauer, K. C. & Aßfalg, A. A robust anchoring effect in linear ordering. *Q. J. Exp. Psychol.* (2019) doi:10.1177/1747021819855234.
54. Mair, P. & Wilcox, R. Robust statistical methods in R using the WRS2 package. *Behav. Res. Methods* (2020) doi:10.3758/s13428-019-01246-w.
55. Jammalamadaka, S. Rao and SenGupta, A. *Topics in Circular Statistics*. *World Scientific* (2001).
56. Buckland, S. T., Davison, A. C. & Hinkley, D. V. Bootstrap Methods and Their Application. *Biometrics* **54**, 795 (1998).
57. Berry, R. B. *et al. The AASM Manual for the Scoring of Sleep and Associated Events. Rules, terminology and technical specifications.* (American Academy of Sleep Medicine, 2015).
58. Oostenveld, R., Fries, P., Maris, E. & Schoffelen, J. M. FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Comput. Intell. Neurosci.* (2011) doi:10.1155/2011/156869.
59. Treder, M. S. MVPA-Light: A Classification and Regression Toolbox for Multi-Dimensional Data. *Front. Neurosci.* **14**, (2020).

SM1 Sound-Image association

Participants first learned the sound-image association to 90% criteria to increase chances that the overnight stimulation would trigger the associated image (see Figure 1(A)). They were then tested on these associations but still received feedback to reinforce the learning. In the morning and during session 3 (2-weeks follow-up) they performed the same test but without feedback. A RM-ANOVA was performed to assess any difference in performance across sessions. Accuracy remained almost constantly at ceiling level for the three sessions (M:0.96 SE:0.004, M:0.97 SE:0.004, M:0.96 SE:0.004 respectively) with non-significant differences among them (smallest p=0.55, refer to Figure SM1).

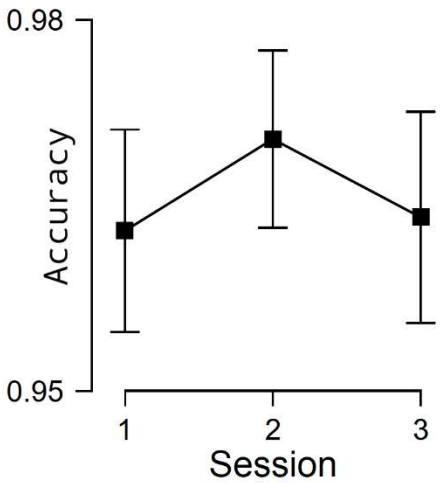

Figure SM1.1: Sound-Image association performance. Accuracy for each session (evening, next morning and 2-weeks follow up). Horizontal bars represent 95 % confident intervals. No statistically significant differences were found.

SM2 Questionnaire

Ten questions were asked at the end of the experiment dived into 4 pages as shown below:

Page 1:

1. Did you have the impression that some of the pairs of patterns were easier to choose between than others? Yes/No
2. Did you think any of the patterns were ALWAYS correct (no matter what the other pattern was)? Yes/No
3. Did you think that any of the patterns were ALWAYS incorrect (no matter what the other pattern was)? Yes/No
4. Did you have any tricks to memorising the individual patterns or the pairs of patterns? Yes/No
   If yes, explain briefly:

Page 2:

5. Did you give names to the patterns? Yes/No

   If yes, can you explain briefly:

6. Did you have the impression that there was some kind of logical rule or order? Yes/No

   If yes, can you explain briefly:

Page 3:

7. In the test phase, did you notice any new combination of patterns taken from those you saw before in the training phase? Yes/No

8. How did you make your choice in the cases? (for example, you guessed, went with instinct). Please explain.

Page 4:

9. Did you think that there was a hierarchy among the patterns seen in training? That is, did you think they could be ranked from "bigger" to "smaller" or from "best" to "worst? Yes/No

10. If you answered Yes, can you write the hierarchy down? (you can use the "names" you have given to the images or any other trick you have used to remember them).

SM3 Premise Pairs

Five pairs of premise-pairs were presented to the participants in each one of the three testing sessions: A-B, B-C, C-D, D-E, E-F. Performance for each individual pair, Condition and Session illustrated in Figure SM3.

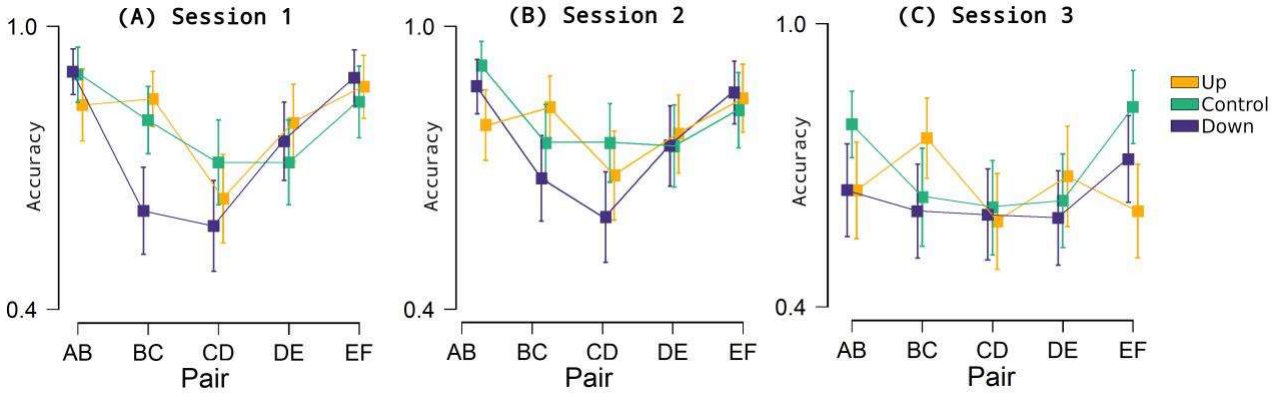

Figure SM3.1: Serial position effect (SPE)[32] for premise pairs (horizontal axis) in each session and each Condition (Up: yellow, Control: green and Down in purple).

Table SM3.1: Number of times each stimulus (F: Faces, O: objects, S: scenes) is used for each one of the three TMR conditions: Up, Down, Control.

Stimuli	Up	Down	Control
F	6	8	6
O	8	5	9
S	8	7	5

Table SM3.2: Premise pairs performance for each session (columns) and condition (rows). Accuracy is indicated as: number of trials, mean accuracy and standard error of the mean respectively. (*) Indicates statistically significant difference from chance level (50%).

	Session 1			Session 2			Session 3		
	n	Mean	SE	n	Mean	SE	n	Mean	SE
Down	390	0.747*	0.018	390	0.747*	0.019	340	0.629*	0.022
Control	390	0.792*	0.018	390	0.796*	0.018	340	0.696*	0.021
Up	390	0.796*	0.018	390	0.782*	0.018	340	0.653*	0.022

SM4 Inference Pairs

Table SM4.1: Inference pair performance for Overnight experiment. First column represents the averaged pair performance, the other two the accuracy divided into the 1st and 2nd degree of separation. Accuracy is depicted for each session and condition of interest and indicated as: number of trials, mean accuracy and standard error of the mean respectively. (*) Indicates statistically significant difference from chance level (50%).

	Averaged			1st Degree			2nd Degree		
	n	Mean	SE	n	Mean	SE	n	Mean	SE
Session 2									
Down	234	0.464	0.028	156	0.455	0.034	78	0.481	0.049
Control	234	0.560*	0.027	156	0.574*	0.032	78	0.532	0.049
Up	234	0.628*	0.028	156	0.587*	0.035	78	0.712*	0.045
Session 3									
Down	204	0.586*	0.029	136	0.577*	0.036	68	0.603*	0.050
Control	2048	0.559*	0.029	136	0.555	0.036	68	0.566	0.050
Up	204	0.583*	0.031	136	0.518	0.038	68	0.713*	0.046

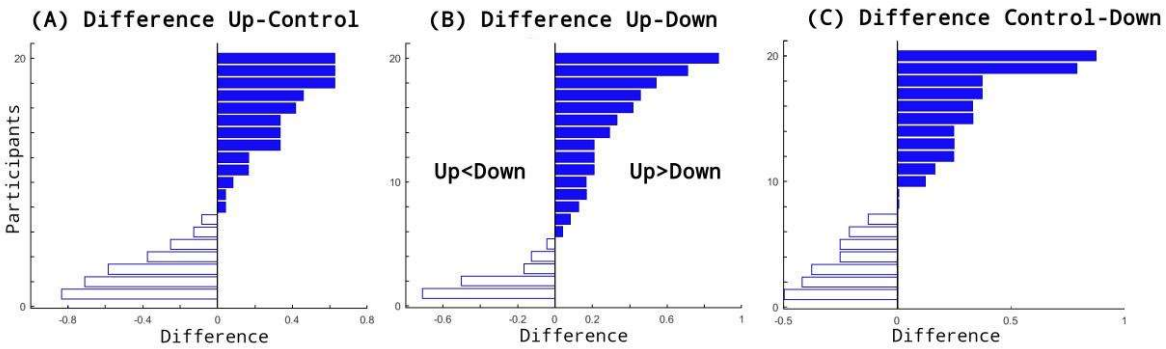

Figure SM4.1: Consistency of inference pair results among participants measured as the difference between the average performance of condition 1 minus condition 2 for the Second Session only. Blue bars represent positive differences, hence condition 1 > condition 2, and white bars negative differences. (A)Difference Up minus Control conditions, where 65% of participants (blue bars) performed better in the Up condition. (B) Difference Up minus Down conditions, where 75% of participants performed better in the Up condition. (C) Difference Control minus Down conditions, where 65% of participants performed equal or better in the Control condition.

Table SM4.2 Inference pair statistical results for RM-ANOVA with Session, Distance and Condition as factors. Statistically significant results are highlighted in bold letters.

	F	*p*
Session	1.969	0.170
Condition	11.81	**0.003**
Degree	5.804	**0.017**
Session*Condition	6.225	**0.046**
Condition*Degree	0.356	0.551
Session*Degree	8.559	**0.015**
Session*Condition*Degree	0.729	0.696

Table SM4.3 Inference pair statistical results per each condition (Down, Up and Control) separately. Statistically significant results are highlighted in bold letters.

	Down		Up		Control	
	F	*p*	F	*p*	F	*p*
Degree	0.44	0.5040	9.61	**0.0023**	0.064	0.806
Session	10.72	**0.0013**	1.16	0.2821	0.151	0.697
Interaction	0.16	0.6868	0.74	0.3824	0.598	0.440
Post-hoc						
Session 2						
Degree	0.68	0.4109	4.39	**0.0378**	0.389	0.5335
Session 3						
Degree	0.10	0.7419	9.94	**0.0020**	0.025	0.8735
Degree 1						
Session	6.90	**0.0094**	1.35	0.2457	0.055	0.8145
Degree 2						
Session	2.22	0.1397	0.02	0.8757	0.247	0.6199

Table SM4.4 Inference pairs statistical results per each session by degree of separation. D: Down, U: Up and C: control. Statistically significant results are highlighted in bold letters.

	Session 2		Session 3	
	$\hat{\Psi}$	p	$\hat{\Psi}$	p
Degree 1				
C vs D	0.191	**0.0303**	-0.042	0.6093
C vs U	-0.026	0.7369	0.054	0.6093
D vs U	-0.218	0.0230	0.097	0.6093
Degree 2				
C vs D	0.052	0.6390	-0.059	0.6083
C vs U	-0.281	**0.0053**	-0.214	0.0844
D vs U	-0.333	**0.0012**	-0.154	0.2149

SM5 EEG results

We performed time-frequency analysis independently for the Up and Down conditions, examining the differences between experimental and control sounds and the interaction effect in the broad-band from 4 to 20Hz. One negative cluster was found for the Up stimulation (p=0.0180) but no significant clusters, positive or negative (p=0.13 and p=0.36 respectively) for the Down case, or interaction (p=0.094).

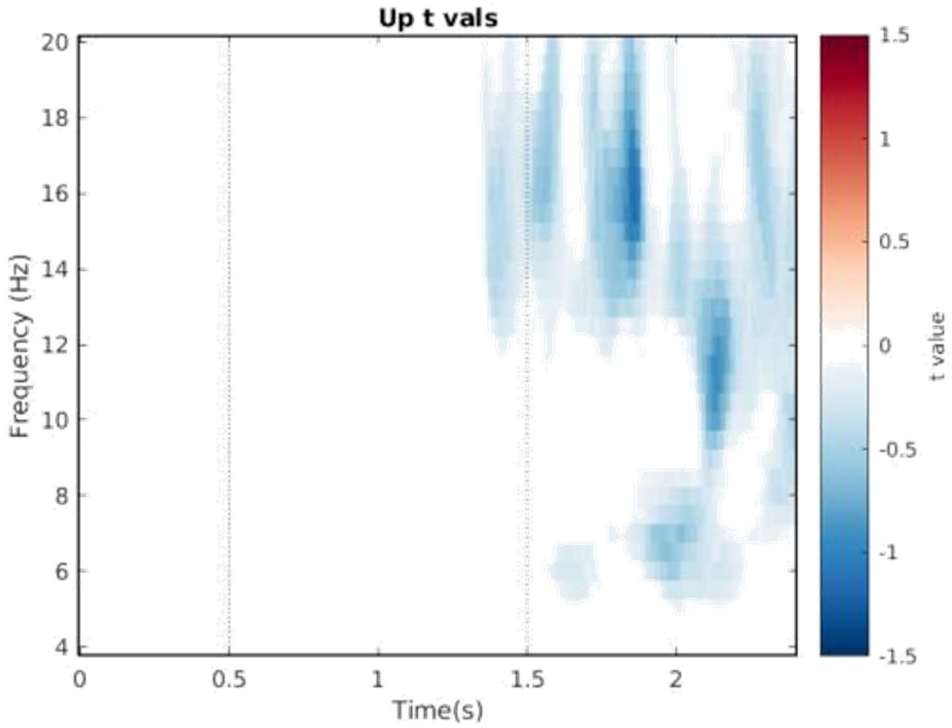

Figure SM5.1: Time frequency representation of the t-values, after thresholding (p<0.05) from the cluster analysis performed on the Up stimulation comparing experimental versus control sounds.

SM6 Awareness

Confidence ratings

To ensure that participants were not explicitly aware of the hierarchy, hence allowing different strategies rather than sleep or TMR enhancement to influence performance[8,10] they were asked how confident there were about their answers after each selection was made in the Late test of Session 2 and 3. The 5-point scale was set from -2 (guessing) to +2 (certain) and was shown for each type of pair presented. Following the same reasoning as in[10], we analysed the confidence ratings of Premise vs. Inference pairs for each condition (Up, Control and Down) (see Figure SM6). If participants were completely aware of the hierarchy, their confidence ratings should be equal for both type of pairs. However, Premise pair confidence was clearly significantly higher than Inference pair confidence for each one of the conditions (all $p<0.0001$), indicating that the participant's ability of make transitive inferences is not due to explicit awareness.

We repeated this analysis for Session 3, where non-significant differences were found between premise and inference pairs' confidence intervals, (lowest $p=0.060$). This is not surprising if we have in account the decremented performance on Premise pairs on this session compared with the two previous ones. Even though they can still perform over chance level, 2-weeks is a long period of time hence participants lost their confidence for premise pairs.

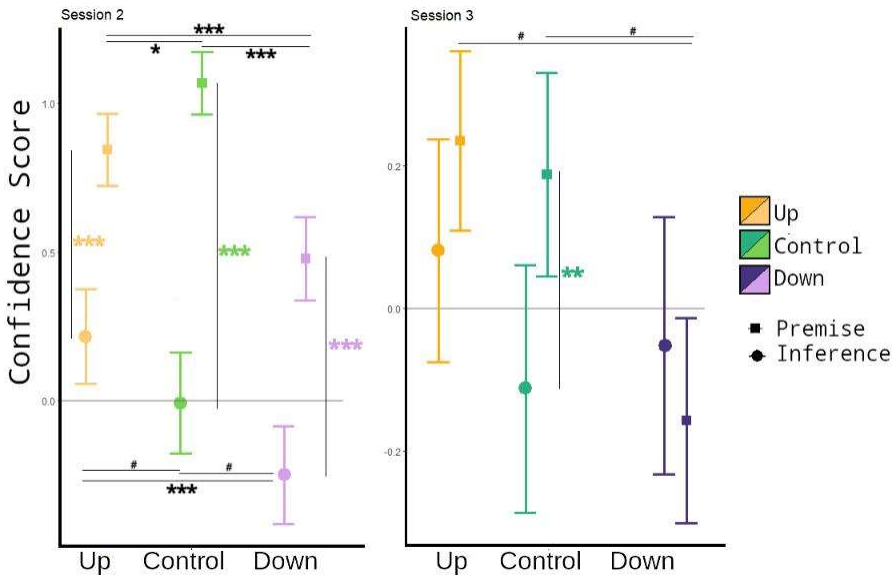

Figure SM6: Confidence score for Session 2 (left) and Session 3 (right) divided into Premise (squares) and Inference (circles) pairs for each condition (Up, Control and Down). Five-point-scale range was set from -2 (completely guessing) to +2 (absolute sure). * $p<0.05$, ** $p<0.01$, *** $p<0.001$, # $p<0.09$.

Awareness questionnaire

Participants' awareness reports are subjective and not very reliable since a confident rating does not necessarily imply awareness of the hierarchy. We therefore introduced a more direct measure of awareness. This was a questionnaire on a scale of 0 to 4; where 3 indicates a high degree of awareness[8,47]. Only one participant scored 4 points; the rest of the participants scored less than 2.25 points indicating no explicit knowledge of the hierarchy. The performance of this participant on the inference pairs was at ceiling for session 2 (Up:1, Down:0.95, Control:0.95) and session 3 (Up:1, Down:1, Control:1).

SM7 Classification

We used two classification algorithms, SVM and LDA, with two different performance metrics, AUC and ACC to distinguish between experimental and control sounds. Cluster statistics resulted in a consistent positive cluster for the Up (see table SM7.1) condition but not significant clusters for the Down condition (see table SM7.2). Results of the SVM classifier with AUC as performance metrics for both conditions are shown in Figure SM7.

Additionally, we performed correlations between the classification performance and the behavioural results for the Up condition taking both the mean and the peak of the significant classifier cluster. Results for the mean performance within the cluster can be seen in table SM7.3 and results for the peak performance within the cluster in table SM7.4.

Table SM7.1: Significant cluster across different algorithms for Up condition

Classifier	Performance	Cluster starts (ms)	Cluster ends (ms)	*p* value
LDA	ACC	1216	1322	0.0060
LDA	AUC	1200	1326	0.0120
SVM	ACC	1200	1306	0.0050
SVM	AUC	1204	1298	0.0370

Table SM7.2: Cluster statistic results across different algorithms for Down condition

Classifier	Performance	lowest p value
LDA	ACC	0.095
LDA	AUC	0.143
SVM	ACC	0.106
SVM	AUC	0.074

Table SM7.3: Correlations results between classifier performance (mean within the significant cluster) and behavioural accuracy of the second degree-inference pairs for the Up condition.

Classifier	Rho	*p* original	*p* corrected
LDA ACC	-0.59	0.006	0.028
LDA AUC	-0.59	0.006	0.031
SVM ACC	-0.64	0.002	0.027
SVM AUC	-0.65	0.002	0.022

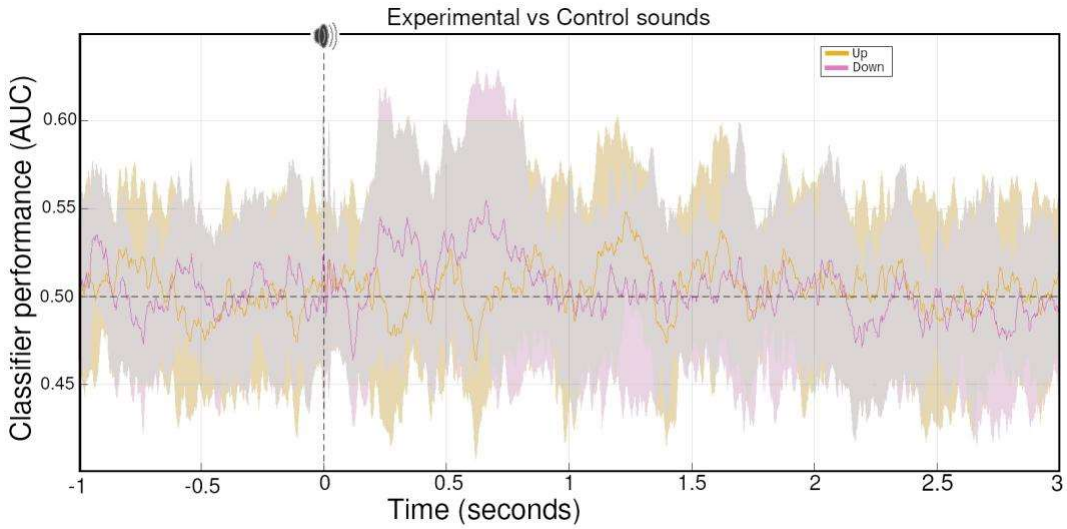

Figure SM7.1: Grand average classifier results for Up (yellowish) and Down (purple) conditions using a SVM with AUC as performance metric. Shadow areas corresponding to the stand deviation across participants.

Table SM7.4: Correlations results between classifier performance (peak of the significant cluster) and behavioural accuracy of the second degree-inference pairs for the Up condition.

Classifier	Rho	p original	p corrected
LDA AUC	-0.49	0.039	0.063
LDA ACC	-0.48	0.039	0.065
SVM ACC	-0.57	0.012	0.042
SVM AUC	-0.56	0.012	0.047