

Sequential replay of non-spatial task states in the human hippocampus

Nicolas W. Schuck^{1,2,*} & Yael Niv²

¹Max Planck Research Group NeuroCode

Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany

²Princeton Neuroscience Institute and Department of Psychology

Princeton University, Washington Road, Princeton, NJ, 08544, USA

*Corresponding author contact:

Max Planck Research Group NeuroCode

Max Planck Institute for Human Development

Lentzeallee 94, 14195 Berlin, Germany

schuck@mpib-berlin.mpg.de

tel: +49 30 82406 649

Abstract

Neurophysiological research has found that previously experienced sequences of spatial events are reactivated in the hippocampus of rodents during wakeful rest. This phenomenon has become a cornerstone of modern theories of memory and decision making. Yet, whether hippocampal sequence reactivation at rest is of general importance also for humans and non-spatial tasks has remained unclear. Here, we investigated sequences of fMRI BOLD activation patterns in humans during wakeful rest following a sequential but non-spatial decision-making task. We found that pattern reactivations within the human hippocampus reflected the order of previous task state sequences, and that the extent of this offline reactivation was related to the on-task representation of task states in the orbitofrontal cortex. Permutation analyses and fMRI signal simulations confirmed that these results reflected underlying neural activity, and showed that our novel statistical analyses are, in principle, sensitive to sequential neural events occurring as fast as one hundred milliseconds apart. Our results support the importance of sequential reactivation in the human hippocampus for decision making, and establish the feasibility of investigating such rapid signals with fMRI, despite its substantial temporal limitations.

Highlights

- We provide fMRI evidence for sequential pattern reactivation in the human hippocampus
- Sequences of patterns reflect states from a sequential, non-spatial decision-making task
- Simulations show that our novel fMRI analysis is sensitive to fast sequences of sub-second neural events
- Results support the importance of sequential reactivation in the human hippocampus for decision making

26 Introduction

27 The hippocampus plays an important role in memory (1–3), and is known to represent
28 spatial as well as non-spatial information that is relevant to an animal’s current location
29 within a ‘map’ of the ongoing task (4–8). Recent research has suggested that hippocampal
30 memories are also used to inform spatial decision making and planing by reactivating neurally
31 encoded experiences that are relevant for the current task (9, 10). Specifically, studies in
32 rodents have shown that hippocampal representations of spatial locations are reactivated
33 sequentially during short on-task pauses, longer rest periods, and sleep (11–13). This
34 sequential reactivation, or replay, is related to better planning (12) and memory consolidation
35 (14), and suppression of replay-related short wave ripples impairs spatial memory (15).

36 While these findings have provided insights into the hippocampal computations underlying
37 spatial decision making in animals, what role replay plays in non-spatial decision making
38 tasks in humans has remained unclear. We instructed participants to perform a non-spatial
39 decision making task, and recorded functional magnetic resonance imaging (fMRI) activity
40 during resting periods before and after the task. We investigated whether sequences of fMRI
41 activation patterns during rest reflected hippocampal replay of task states. Evidencing such
42 replay, transitions between neural activity patterns were related to previously experienced
43 sequences of task states. Moreover, reactivation in the hippocampus during rest was asso-
44 ciated with better representation of the same task states in the orbitofrontal cortex during
45 decision making. In line with our previous work, these orbitofrontal on-task representations
46 were related to better performance of the task (16).

47 Our results demonstrate sequential reactivation of non-spatial decision-making states in
48 the human hippocampus and suggest that representations reflecting the structure of the
49 current task are supported by the interaction of hippocampal and prefrontal brain systems.
50 Our findings, together with a set of rigorous statistical tests and simulations, also establish
51 the utility of noninvasive fMRI to detect possibly rapid replay events, despite the low temporal
52 resolution of this method.

53 Results

54 Thirty three human participants performed a sequential decision-making task that required
55 integration of information from past trials into a mental representation of the current task
56 state (16, see Methods). Specifically, each stimulus consisted of overlapping images of a
57 face and a house and participants’ main task was to make age judgments (old or young)

58 about one of the images (Fig. 1A). The category to be judged (face or house) was instructed
59 before the first trial. Subsequent category switches were determined by the following rule:
60 if the age in the current trial was the same as the age in the previous trial, then the judged
61 category remained the same; on the other hand, if the age on the current trial was different
62 from the age on the previous trial, the participant had to switch to the other category from
63 the next trial onward (Fig. 1B). This created a ‘miniblock’ structure where each miniblock
64 involved judgment of one category. Miniblocks were at least two trials long (that is, no age
65 comparison was required on the first trial after a switch), and on average lasted for three
66 trials. These task rules resulted in a total of 16 task states, which were experienced in
67 a structured order (Fig. 1C). For example, the (Ho)Fy state, indicating an old house trial
68 followed by a young face trial, was only experienced after a change from young to old houses.
69 Participants performed the task with high accuracy (average error rate: 3.1 %, time-outs:
70 0.6%, reaction time: 969 ms), improving their performance throughout the course of the
71 experiment (see Fig. 1D, significant linear trends of task block for reaction times and errors,
72 both $ps < .001$, see Supplemental Information, SI, Fig. S4, for further details).

73 The experiment comprised two sessions during which participants engaged in the above
74 described decision-making task while undergoing fMRI. The first session included task in-
75 structions and four runs of task performance (388 trials, about 40 minutes duration). The
76 second session took place one to four days later and was identical to session 1, but without
77 instructions (Fig 1E). Resting state scans consisting of 5 minute periods of wakeful rest
78 without any explicit task or visual stimulation (100 volumes per resting state scan) were
79 administered for $N = 23$ participants (group 1) after session 1, before session 2 and after
80 session 2, resulting in a total of 300 wholebrain volumes acquired during rest. A second
81 group of participants ($N = 10$; group 2) underwent the same procedures as group 1,
82 plus one additional resting state scan at the beginning of session 1, before having had
83 any task experience or being exposed to task instructions. This resulted in a total of 400
84 wholebrain volumes acquired during rest. The analyses reported below focus on fMRI data
85 recorded during these resting scans. Resting state data acquired after participants had task
86 experience will from hereon referred to as the TASK rest condition, whereas resting state
87 data acquired before the task as the PRE rest condition. Data recorded while participants
88 received instructions will serve as another control and referred to as the INSTR condition.
89 To account for differences in the number of data points constituting the TASK vs control
90 conditions, we used a size matched TASK condition where appropriate. Notably, while none
91 of these conditions involves active experience of the sequential decision making task, they




Figure 1: Experimental task and performance. (A): On each trial, participants had to judge the age of either a face or a house shown overlaid in a compound stimulus. Trials began with the display of a fixation cross and the response mapping (1200ms), followed by the stimulus. Responses could be made at any time, and the stimulus stayed on screen for an average of 3300 ms. (B): The rule of the task required participants to switch between judging faces and houses whenever the age changed between two trials, see text. (C): The state space of the task. Each node represents one possible task state, and each arrow a possible transition. All transitions out of a state are equally probable, occurring with $p = 0.5$. Each state of the task is determined by the age and category of the previous and current trial, indicated by the acronyms (see legend). States are colored depending on whether they correspond to trials in which the age and category were repeated (orange), the age changed (green) or the category changed (purple). (D): Average error rates and reaction times across the two sessions. Bars: ± 1 S.E.. Grey dots represent individual subjects. (E): The experiment extended over two sessions, each of which included about 40 minutes task experience flanked by resting state scans. *:The pre-task resting state scan in Session 1 was performed only for a subgroup of our sample ($N = 10$; group 2).

92 differ in whether the task has been experienced before or not.

93 The main goal of our study was to investigate sequential reactivation, or replay, of
 94 task-related experiences in the human hippocampus during rest. To this end, we trained
 95 a multivariate pattern recognition algorithm (see Methods) to distinguish between the acti-
 96 vation patterns associated with each of the 16 task states in the data recorded during task
 97 performance (Fig. 2A,B). Leave-one-run-out cross-validated classification accuracy on the
 98 task data from the hippocampus (HPC) was significantly higher than chance (6.25%) and
 99 than classification obtained in a permutation test (11.6% vs 7.1%, $t_{32} = 8.9$, $p < .001$, Fig.
 100 2C), indicating that hippocampal activation patterns indeed reflected task states. We then
 101 applied the trained classifier to each volume of fMRI data acquired during the resting state
 102 scans. Although classification accuracy cannot be assessed for the resting scan data (due

103 to lack of ground truth), we could assess the quality of the classification using the mean
104 unsigned distance to the decision hyperplane, a proxy for classification certainty (17). This
105 distance was larger in the TASK condition compared to simulated spatiotemporally-matched
106 noise ('NOISE', $t_{32} = 12.9$, $p < .001$; for simulation details see Methods and SI) and the
107 PRE condition ($t_9 = 2.1$, $p = .031$, group 2 only, Fig. 2D). This suggests that fMRI patterns
108 recorded during resting-state scans following task experience could reflect reactivation of
109 task states, in line with previous findings (18–20).

110 The defining aspect of replay is that previously experienced states are reactivated *se-*
111 *quentially*. We therefore asked next whether it is theoretically possible to measure rapid
112 sequential replay events (on the order of few hundreds of milliseconds in humans (21)) using
113 fMRI, given its low temporal resolution. To this end, we simulated fMRI activity that
114 would result from fast replay events (see SI and below), and asked what order and state
115 information could be extracted from these spatially and temporally overlapping patterns.
116 The slow hemodynamic response measured in fMRI causes brief neural events to impact the
117 BOLD signal over several seconds. Although these same dynamics might obscure the details
118 of a replayed sequence, our simulations showed that two successive fMRI measurements can
119 still reflect two states from the same sequence, for instance the first and last element of a
120 multi-step replay event (see SI). Because replay events mainly reflect short sequences of states
121 (e.g., Ref. 13, their figure 3C), we can therefore expect that consecutively decoded states be
122 close in the task's state space (that is, separated by few intervening states in Fig. 1C), if they
123 indeed reflect sequential replay. We further asked whether we could expect to successfully
124 decode a pair of states from the same replay event, given the low accuracy of correctly
125 decoding task states during task performance. Our simulations showed that, because brain
126 activity recorded after a rapid replay event includes several superimposed states (Fig. S5B),
127 the likelihood of classifying *one out of several* replayed states in each resting state brain
128 volume is actually considerably higher than the overall decoding accuracy when classifying
129 a single event during task. The chance that analyzing two consecutive brain volumes results
130 in decoding one (ordered) set of two states out of several possible sets caused by the same
131 replay event may therefore be on the order of the overall decoding accuracy ($\sim 10\%$; see SI).

132 Having established that, in principle, we can detect sequential replay in fMRI data,
133 we next investigated whether the sequences of states we decoded in the TASK resting
134 data (Fig. 3A) reflected sequences experienced during the task. Note that testing for
135 sequentiality in the decoded data is not trivial given that the classifier was trained on task
136 data that was itself sequential. As a result, apparent 'sequentiality' can be found even in

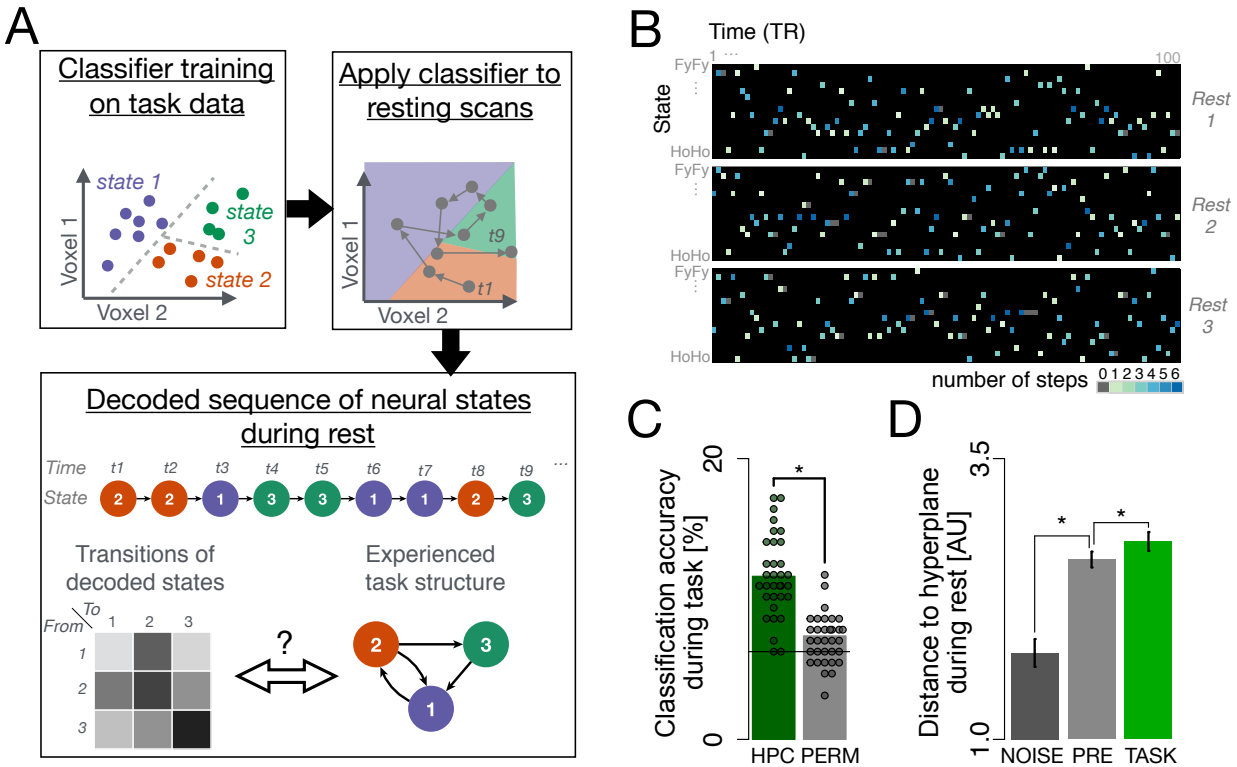


Figure 2: Sequential replay decoding analysis. (A): Illustration of analysis procedure. For simplicity, only two dimensions and three state classes are shown. We first trained a classifier to distinguish between the different task states in the hippocampal fMRI data acquired during the task. The trained classifier was then applied to each volume of fMRI data recorded during resting sessions (grey dots). This resulted in a sequence of predicted classifier labels that was then transformed into a ‘transition matrix’ T that summarized the frequency of decoding each pair of task states consecutively. The structure of the decoded sequences, as summarized by this matrix, was then compared to the sequential structure of the task (see text). Note that the real analysis involved 16-way classification of >1000 dimensional data, which was compared to the task state space shown in Fig. 1C. (B): Example data from one randomly selected participant. Each dark rectangle illustrates the sequence of classified states for the 100 volumes of fMRI data recorded in one resting state scan (depicted are three resting state scans acquired throughout the experiment; see Fig. 1E). Columns represent time, and rows states. Each color-filled cell represents the state classified at the respective time point, and color indicates the distance (in steps in the state space; Fig. 1C) from the state decoded in the previous timepoint (i.e., the previous TR). (C) Classification accuracy during task performance was significantly higher in hippocampal data (HPC) than in a permutation test (PERM). The solid line represents the theoretical chance baseline of $100/16=6.25$. (D): Average distance to the hyperplane for classified states during rest in the NOISE (dark grey, left bar), PRE (light grey, middle bar, $N=10$) and TASK conditions (green, rightmost bar, $N=33$). Larger distance indicates higher certainty in the classification of the state. Each dot indicates one participant, bars within-subject S.E.M., *: $p < .05$.

137 random noise—although clearly those data do not reflect sequential replay. We therefore
 138 conducted a series of carefully controlled assessments of the levels of sequentiality in our
 139 data. Indeed, several forms of sequentiality predicted by replay were evident in our data
 140 when compared to a series of carefully matched controls. First, we predicted that replay
 141 would be reflected in a decreased number of steps that separate two consecutively decoded

142 states, as indicated by the above mentioned simulations. In line with this idea, the number
143 of steps between state transitions decoded in the TASK resting condition was smaller, on
144 average, than the distance between states in the INSTR condition ($t_{32} = 2.4, p = .01$),
145 smaller than the distance found in the PRE condition ($t_9 = 2.3, p = 0.02$, group 2 only) and
146 smaller compared to a permutation test in which classified states were randomly reordered
147 to control for overall state frequency (PERM condition: $t_{32} = 4.6, p < .001$; Fig. 3B,C).
148 Second, because replay events are separated by long pauses (21), and sequentiality should be
149 present only following the replay events, we expected the occurrence of short-distance state
150 pairs to be clustered in time. Indeed, short-distance state pairs (less than 3 steps apart)
151 were not only more frequent than expected, but were also more likely to occur in clusters
152 in the TASK rest condition compared to the INSTR ($t_{32} = 1.7, p = .046$), PRE ($t_9 = 1.9,$
153 $p = .044$, group 2 only), and PERM controls ($t_{32} = 4.5, p < .001$, Fig. 3D). Third, we
154 confirmed that neither the high prevalence of one particular step size nor sustained state
155 activation would distort our conclusions regarding sequenceness in the TASK condition.
156 To this end, we tested whether the frequency of decoded state transitions was linearly
157 related to the distance between them in task space while also excluding state repetitions
158 from the analysis. Specifically, we tested whether the empirical frequency of decoding each
159 pair of task states consecutively (the ‘transition probability’ for each pair of decoded states,
160 summarized in matrix T ; Fig. 3A) was negatively correlated with the distance D between
161 states during the task (where D_{ij} corresponds to the minimum number of steps necessary to
162 get from state i to state j ; Fig. 3E). This was indeed the case, with an average correlation
163 between D and T of $r = -.16$ ($t_{32} = -17.7, p < .001$, t-test of individual correlations across
164 participants, Fig. 3F). While we also found a correlation when the order to decoded states
165 was permuted (PERM condition, $r = -.08, p < .001$), reflecting an effect of overall state
166 frequency, this correlation was substantially smaller than in the TASK data ($\Delta r = -.07,$
167 $t_{32} = -5.8, p < .001$). Likewise, applying the trained classifier to matched fMRI noise
168 (see Methods) also showed that temporal contingencies between states in the training data
169 for the classifiers lead to spurious correlations (NOISE condition, $r = -.08, p < .001$),
170 but these were also significantly smaller than the correlation found in the TASK rest data
171 ($\Delta r = -.08, t_{32} = -5.6, p < .001$, Fig. 3G). Importantly, our hypothesis that sequential
172 reactivation of task-state representations during rest was caused by task experience was also
173 supported by a significantly lower correlation between D and T in the TASK condition
174 as compared to the INSTR data ($t_{32} = -12.1, p < .001$, when compared only subset of
175 TASK condition matched in number of TRs), as well as the PRE resting scan ($t_9 = -7.9,$

176 $p < .001$, group 2, $p = .059$ when compared to only first resting scan in TASK condition,
177 Fig. 3H). Finally, we also assessed the effect of the sequential structure in the training
178 data on our results in an additional control analysis in which we systematically excluded
179 sets of state pairs from classifier training (see SI, Fig S2), to test if, as a result, these pairs
180 would show a lower frequency in the resting data. The excluded transitions were observed
181 as often as the included transitions ($t_{32} = 0.3$, $p = 0.73$), in line with our conclusion that
182 the transition frequencies observed during rest reflected sequential reactivation above and
183 beyond any sequential structure in the classifier.

184 In order to investigate the effects of task experience on pair-decoding frequency data T
185 while simultaneously (a) excluding state repetitions, (b) controlling for the above-mentioned
186 effect of temporal contingencies in the classifier training and (c) incorporating the different
187 sources of between- and within-participant variability, we performed a logistic mixed-effects
188 model that included including nuisance covariates (see Methods). We will henceforth call the
189 effect estimate (beta weight) of the distances D on the data T in this model ‘sequenceness,’
190 and the nuisance effect ‘randomness.’ Comparing models containing the randomness regres-
191 sor with vs. without an additional sequenceness regressor to explain frequency of transitions
192 showed no difference in model fit in the PRE rest condition (Aikake Information Criterion,
193 AIC, 3651.5 vs 3651.4, $\chi^2_1 = 1.9$, $p = .17$). In the TASK rest condition, in contrast, adding
194 the sequenceness regressor improved model fit (AIC 3642.1 vs 3645.4, $\chi^2_1 = 5.2$, $p = .02$,
195 group 2 only and considering only the first TASK resting scan from the first session to
196 equate power). Modelling both conditions within one model also showed improved fit when
197 the interaction of condition factor with sequenceness and randomness was included (AIC
198 3660.2 vs 3674.1, $p < .001$). Figure 4A/B shows the sequenceness and randomness effects in
199 the TASK compared to the PRE condition. Comparing the INSTR to the TASK condition in
200 all participants showed the same pattern of effects: No effect of the sequenceness regressor
201 was found in the INSTR condition (AIC 10046 vs 10047, $p = .27$), but in the TASK rest
202 condition (AIC 10130 vs. 10146, $p < .001$, TASK data matched in size to equate power),
203 see Fig.s. 4C/D. A combined model indicated no interaction between condition and the
204 pattern transitions however (10142 vs 10130, $p > .1$). Note that the lack of sequenceness
205 before task experience shows that our modelling approach analysis successfully controlled for
206 bias effects due to the temporal contingencies between states in the classifier training data.
207 Analyzing data from all participants (groups 1 & 2) and all TASK resting-state scans with
208 this model showed that the inclusion of a state distance factor led to significantly better
209 model fits even after controlling for the randomness (bias) effect as above (AIC 10789 vs

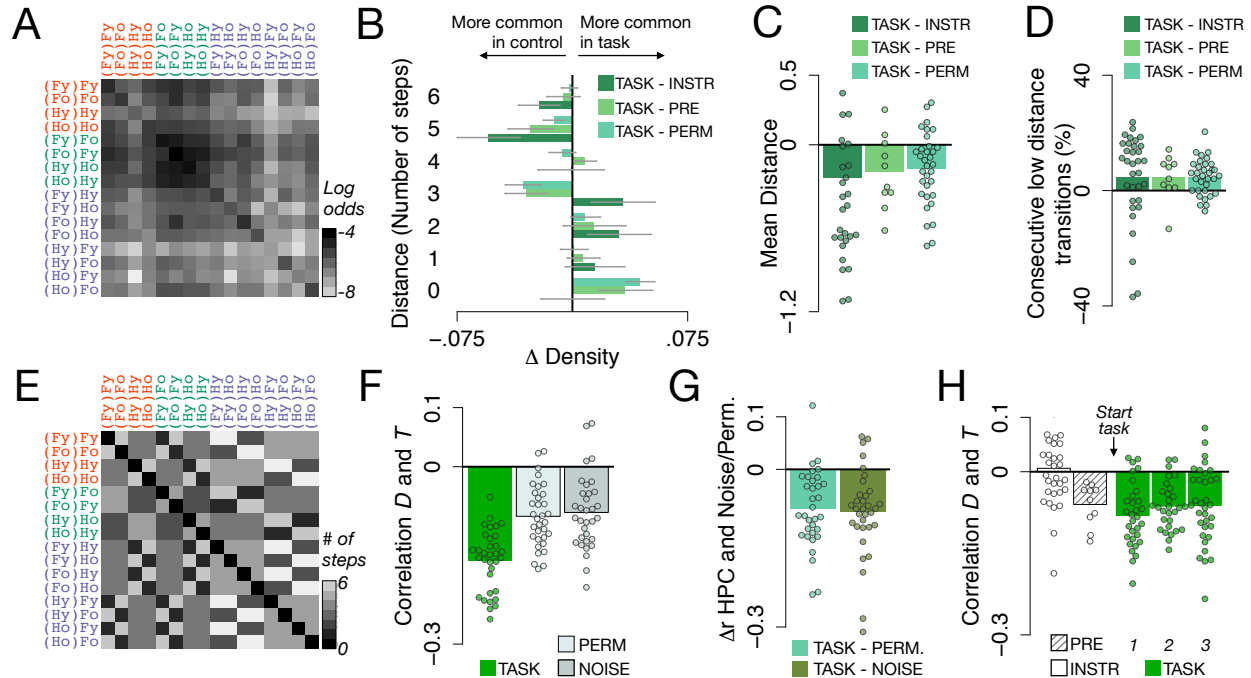


Figure 3: Hippocampal state transitions during rest are related to state distances in the task. (A): The matrix T , expressing the log odds of transitions between all states in the sequence of classification labels in the hippocampal TASK rest data, averaged across all participants. Y-axis: first state, x-axis: second state, in each consecutively decoded pair. Darker colors reflect a higher probability of observing a pair in the data. (B): Relative distributions of number of steps separating two consecutively decoded states. A distance of 1 corresponds to a decoded state transition as experienced in the task, 2 corresponds to a transition with one item missing in between as compared to the task, etc. Barplots show the difference in relative frequency (Δ Density) with which each transition type was observed in the TASK condition compared to INSTR and PRE control conditions and a permutation test (PERM), see legend. (C): The average distance in state space of two consecutively decoded states was significantly lower in the TASK data as compared to the INSTR, PRE and PERM controls (all $ps < .05$, t-test comparing difference to 0). (D): Low-distance transitions (fewer than 3 steps) occurred in succession significantly more frequently in the TASK data compared to all controls (all $ps < .05$). (E): The matrix D , indicating the minimum number of steps between each pair of states in the task, i.e. the state distances. Lighter colors reflect a higher number of steps between states. (F): Average correlations between the state distance matrix D and the corresponding decoded transition matrix T in the TASK condition (green bar, left), as compared to a permutation test (light grey, middle) or when the same classifier was applied to spatio-temporally matched noise (NOISE; dark gray bar, right). (G): Within-participant differences between correlations in TASK versus the PERM and NOISE controls (all $ps < .05$) (H): The correlation between D and T in the PRE and INSTR phases and each of the TASK resting state sessions. Dots reflect correlations/differences of individual participants, bars S.E.M.

210 10780, $\chi_1^2 = 11.0, p < .001$), supporting the conclusion that previously experienced sequences
 211 of task states are replayed in the human hippocampus during rest periods. These results
 212 were unaffected by the choice of distance metric, see SI. No comparable pattern of results
 213 emerged when data from the orbitofrontal cortex, a brain area known to contain task-state
 214 information during decision making (16, 22), was analyzed (see SI).

215 We next tested whether the sequenceness found in the TASK condition could be explained

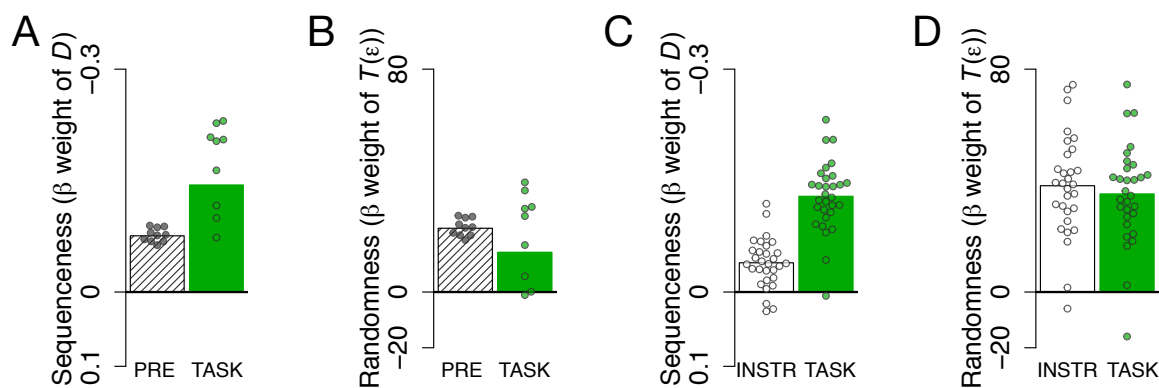


Figure 4: Effect of state distance (sequenceness) on transition frequency in hippocampal data is specific to TASK rest conditions. Bars indicate strength of fixed effects in mixed effects model (see text). Dots indicate individual random effects. Note that variability of dots in this case cannot be used to infer significant differences. (A): Effect of sequenceness regressor on resting data from the PRE and TASK conditions. Model comparisons based on AIC showed that the sequenceness regressor led to better model fit in the TASK but not the PRE condition. (B): Effect of randomness across the PRE and TASK conditions. The randomness regressor captures the sequentiality in the data due to a classifier bias, see text. (C): Sequenceness in the INSTR and TASK conditions, as in panel A. Adding the sequenceness regressor led to better model fit only in the TASK condition. (D): Randomness in the INSTR and TASK conditions as in (B).

216 by backwards replay, or replay of partial states such as stimuli, instead of forward replay
 217 of complete state information. To this end, we defined alternative transition functions
 218 corresponding to the above hypotheses, and tested the power of these transition functions to
 219 explain the sequences of states during rest. We used one-step transition functions instead of
 220 state distances to avoid statistical disadvantages of alternative models that have very evenly
 221 distributed distance (high entropy) functions. As in our original analysis, all transitions
 222 functions were based on the sequence of trials experienced in the task. The alternative
 223 transition functions were created by assuming that only partial aspects of each trial are
 224 represented, for instance by computing the experienced transitions between attended stimuli
 225 without contextualisation by the event in the previous trial. As the classifier was trained
 226 to distinguish all 16 possible states, we assumed that different states corresponding for
 227 instance to the same stimulus would be fully aliased. This approach allowed to calculate the
 228 likelihood that the observed sequences of states were generated by (a) replay of trials, (b)
 229 replay of states containing only information about the current attention, (c) replay of states
 230 containing information about the current and past information and (d) backward replay.
 231 Fig. 5 A-D shows the transition functions used in these analyses. Model comparison of the
 232 same mixed effects models as above showed that the transition function assuming full state
 233 representations (Fig. 5E) led to a better model fit compared to all four alternative models
 234 (AIC: 20808, 20808, 20806, 20796, for the 4 alternative models, respectively; AIC of true

235 model: 20782, see Fig. 5F).

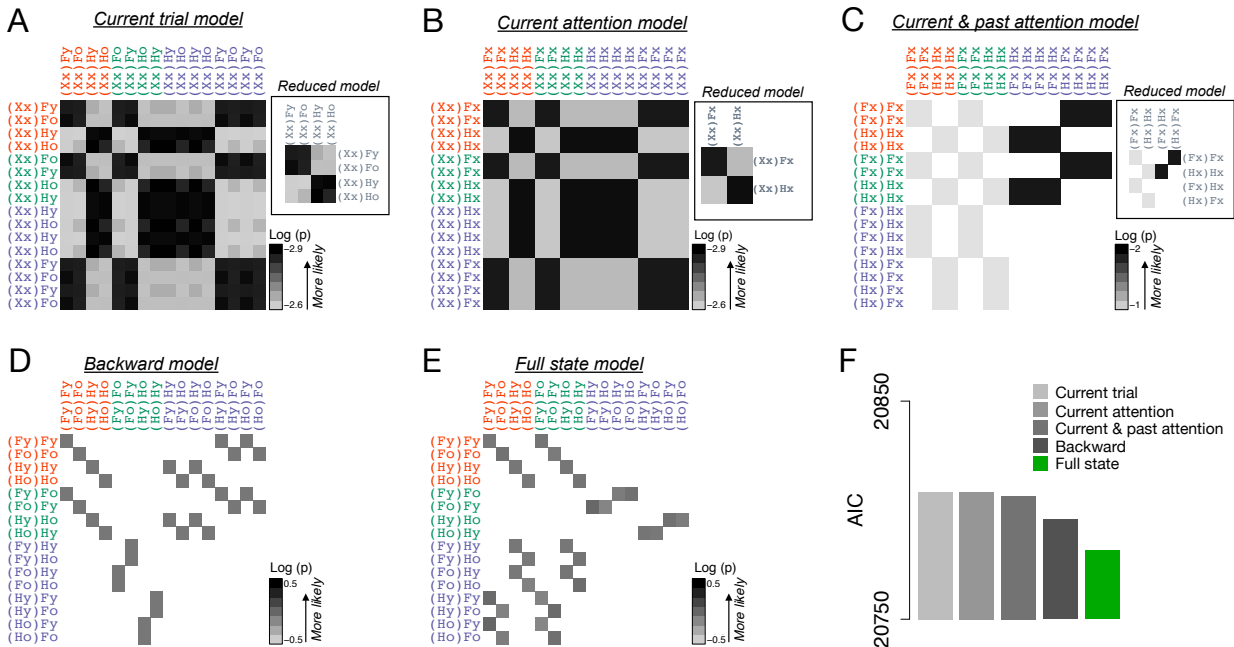


Figure 5: Alternative state transition functions have less power to explain hippocampal state sequences during rest. (A-E) Alternative state transitions. Rows indicate origin states and columns indicate receiving states for a given transition. Color shading indicates log likelihood of the corresponding transition in each model, see legend. Empty (white) cells indicate that a transition is not possible. 'Reduced model' panels in A-C show the transition function when aliased states are collapsed. (F) Akaike information criterion when data from the TASK rest condition was modelled using the transition functions shown in A-E.

236 To test whether the observed sequenceness in hippocampal fMRI data could have been
 237 caused by fast sequences of neural events in principle, we then simulated fMRI signals
 238 generated by sequences of neural events of different speeds, and asked at which speed the
 239 above analyses can uncover the underlying sequential structure. In these simulations, each
 240 neural event triggered a hemodynamic response in a distributed pattern of voxels (see SI; Fig.
 241 S3). The simulations confirmed that following replay, decoding of replayed state identities is
 242 possible over multiple TRs, even when fast replay speeds caused the involved fMRI patterns
 243 to be highly overlapping (Fig. S5). Importantly, when signal-to-noise ratio was adjusted to
 244 yield state-decoding levels that were matched to our data (12.1% accuracy in simulations,
 245 vs. 11.6% in the data), significant correlations between the consecutively decoded state pair
 246 frequencies T and the corresponding distances D were found even at replay speeds of about
 247 14 items per second (i.e. inter-event intervals of 60-80ms; $r = -0.018$; permutation test:
 248 $r = -0.003$; t-test of sequence vs permutation results: $t_{199} = -4.42$, $p < .001$, corrected for

249 multiple comparisons; corresponding test for events separated by faster events at 40-60 ms:
250 $p = .18$; $p < .05$ for all slower sequences; Fig. S6). This suggests that our findings in the
251 resting-state data may reflect fast sequential replay in the human hippocampus.

252 In combination, these analyses show that sequences of hippocampal fMRI activity pat-
253 terns during rest were systematically related to previous task experience. Interestingly, we
254 found no such effect when we included backward distances between states instead of the
255 forward distance in the model. This indicates that the sequences of hippocampal activity
256 patterns became directionally structured in correspondence to participants' task experience.

257 Finally, we investigated the functional significance of hippocampal replay of abstract
258 task states. One idea is that replay helps to form, or further solidify, a representation of
259 the transitions between states of the task (23–25). We therefore tested for a relationship
260 between sequential state pattern reactivation during rest and better representation of states
261 during the task, as measured through cross-validated state decoding accuracy in fMRI data
262 recorded during task performance. We did not find any evidence of a relationship between
263 hippocampal sequenceness at rest and decoding of states during task performance ($r = -.05$,
264 $p > .05$). However, we did find a significant correlation between hippocampal sequenceness
265 at rest and orbitofrontal state representations during the task ($r = -.47$, $p = .005$). Notably,
266 in previous work we have shown that improved state decoding in the orbitofrontal cortex is
267 associated with better decision making in this task (16, see). This finding therefore suggests
268 a role for hippocampal replay in supporting the integrity of task-relevant orbitofrontal
269 state representations during decision making. We also tested for a relationship between
270 hippocampal replay at rest and behavioral measures of task performance, but did not find
271 any evidence for a direct relationship between sequenceness and reaction times, error rates,
272 or the change in these measures across runs (all $ps > .10$). However, in line with our previous
273 work, we did find a relationship between the change in orbitofrontal decoding accuracy during
274 the task and improvements in task performance. That is, runwise decoding within the first
275 session was correlated with runwise error rates ($\chi^2_1 = 9.1$, $p = .003$, using the same decoding
276 measure as used before, see ref. 16). This was not the case for on-task decoding in the
277 hippocampus ($p = .87$, interaction with ROI: $\chi^2_1 = 5.2$, $p = .023$). This result suggests
278 that the hippocampus supports the offline formation or maintenance of a 'cognitive map' of
279 the task, while the orbitofrontal cortex is deployed to represent such a map during decision
280 making (16, 26).

281 Discussion

282 Our findings demonstrate that fMRI patterns recorded from the human hippocampus dur-
283 ing rest reflect sequential replay of non-spatial task states previously experienced in a
284 decision-making task. Previous studies have relied on sustained elevated fMRI activity in the
285 hippocampus or sensory cortex as evidence for replay (18–20, 27), investigated wholebrain
286 MEG signals (28), or studied EEG sleep spindles and memory improvements that are thought
287 to index replay activity (29–33), but were not able to directly demonstrate sequential replay
288 in the human hippocampus. Our study represents an important extension of these findings
289 by providing evidence of sequential offline reactivation of non-spatial decision-making states
290 in the human hippocampus.

291 Evidence of sequentiality and localization of replay in the human hippocampus is in direct
292 correspondence with animal studies in which replay has been shown to be sequential and
293 specific to hippocampal place cells (e.g. 34). Importantly, unlike the majority of previous
294 investigations in animals, the here reported sequences of activation patterns signify the
295 replay of non-spatial, abstract task states. Our results therefore add to a growing literature
296 proposing a significant role for ‘cognitive maps’ in the hippocampus in non-spatial decision
297 making (3, 8, 26, 35).

298 Our findings are in line with the idea that the human hippocampus samples previous task
299 experiences in order to improve the current decision-making policy, a mechanism that has
300 been shown to have unique computational benefits for achieving fast and yet flexible decision
301 making (23–25). Dating back to Tolman (36), this idea requires a neural mechanism that
302 elaborates on and updates abstract state representations of the current task, regardless of the
303 task modality. Several studies have suggested that the hippocampus and adjacent structures
304 support a broad range of relational cognitive maps (35), as evidenced by hippocampal
305 encoding of not only spatial relations but also temporal (37, 38), social (7), conceptual
306 (6) or general contingency relations (39). Here, we showed that the human hippocampus
307 not only represents these abstract task states, but also performs sequential offline replay of
308 these states during rest.

309 Our results imply a relationship between hippocampal replay and the representation
310 of decision-relevant task states that are thought to reside in the orbitofrontal cortex (16,
311 22, 40–42). The relationship between ‘offline’ hippocampal sequenceness and the fidelity of
312 ‘online’ orbitofrontal task-state representations raises the possibility that the hippocampus
313 supports the maintenance and consolidation of state transitions that characterize the task
314 and are employed during decision making (38). Given our findings and recent evidence

315 implying hippocampal place and entorhinal grid cells in signaling non-spatial task-relevant
316 stimulus properties (6, 8), a crucial question for future studies will be to further specify how
317 flexible, task specific representations in the hippocampus interact with task representations in
318 other brain regions (26). Of particular interest will be investigations asking whether neural
319 populations in the hippocampus and entorhinal cortex share a common neural code for
320 abstract task states with orbitofrontal (16) and medial prefrontal regions (43), as indicated
321 by recent studies (6, 44, 45). Together with our findings, this research promises to shed light
322 on the neural representations and computations underlying memory and decision making.

323 **Experimental Procedures**

324 **Participants**

325 Thirty nine participants were selected according to standard fMRI screening criteria (right
326 handedness, 18–35 years of age, normal or corrected-to-normal vision and no contraindication
327 for fMRI) from the Princeton University community, and were compensated with \$20 per
328 hour plus up to \$5 performance-related bonus. Six participants were excluded from analysis
329 due to either technical errors (3 participants), violation of performance criteria standards (2
330 subjects with over 3 times the average error rate in the last two blocks of the experiment)
331 or incomplete data (1 participant). The final sample consisted of 33 participants (22 female,
332 mean age 23.4 years).

333 **Stimuli**

334 Stimuli consisted of spatially superimposed images of a face and a house (see (16); face
335 images from <http://faces.mpdl.mpg.de/faces> described in (46), see Fig. 1). Faces and
336 houses could be classified as either young or old, e.g. a stimulus could show an old face
337 image blended with a contemporary (i.e., young) house image. Four classes of stimuli were
338 possible: (1) two old or (2) two young face and house pictures, (3) a young face with an old
339 house or (4) vice versa.

340 **Task**

341 The task was identical to Schuck et al. 2016 and will be described only briefly. Each trial
342 began with the display of the mapping of a left and right key to a young and old response
343 (changing randomly trialwise) below a fixation cross for 1.2s (range: 0.5–3.5s). Then, a

344 compound face-house stimulus was shown for 3.3s (range: 2.75-5s; Fig. 1) and participants
345 had to make an age judgment about one of the two image categories. Participants knew
346 which category of the stimulus they had to judge by applying the following rules: 1. before
347 the first trial of each run, the category to judge was displayed on the screen; 2. Once the age
348 of the relevant category changed (e.g., from young to old), the judged category changed on
349 the next trial. 3. No age comparison was necessary on the first trial after a category change,
350 i.e. each category was judged for at least two trials in a row before a switch. The average
351 trial duration was 4.5s (range: 3.25-8.5s), all timings were randomly drawn from a truncated
352 exponential distribution and the response deadline was 2.75s. The category instruction cue
353 at the beginning of a run was displayed for 4s. Erroneous or time-out responses led to
354 feedback (written above stimulus for 0.7s) and trial repetition. If an error trial involved an
355 age change (and thus would require a category switch on the next trial), participants had to
356 repeat the trial before the error as well as the error trial, giving them the chance to observe
357 the age change. Otherwise, they had to repeat the trial on which they made the error.

358 Design

359 Participants underwent two fMRI sessions. The first session began with the display of
360 written instructions while participants underwent a functional scan (group 1), or a 5 minute
361 resting-state scan followed by instructions (group 2). The instructions explained the rules of
362 the task and contained a training phase in which simple age judgments had to be made on
363 (non-overlapping) face and house images. The images shown in this period were later used in
364 the task, thus familiarizing participants with the age judgment aspect of the task as well as
365 the stimuli. The instructions furthermore involved an annotated walk-through of four trials
366 of the real task (i.e., with overlapping images and the requirement to switch attention after
367 an age change). Following the instructions, participants performed 4 runs of the task (97
368 trials per run, 388 total). Each run lasted about 7-10 minutes and participants were given
369 the chance to rest briefly between runs. A 5 minute fieldmap scan was done between runs
370 2 and 3, resulting in a longer break for participants. After run 4, participants underwent a
371 resting state scan as well as a structural scan. Lights were turned off during resting-state
372 scans and participants were instructed to stay awake for the entire duration of the scan (5
373 minutes, 100 TRs). The second session was identical for all participants and involved the
374 following scans: resting state, 2 task runs, fieldmap, 2 task runs, resting state and structural
375 scan. Thus, overall, participants performed 8 task runs and 3 (group 1) or 4 (group 2)
376 resting-state scans. In all other regards, the task design involved the same characteristics as

377 detailed in Schuck et al. (2016).

378 Behavioral Analyses

379 Behavioral analyses were done using mixed effects models implemented in the package lme4
380 (47) in R (48). The model included fixed effects for Block, Condition, Category and intercept.
381 Participants were considered a random effect on the intercept and the slopes of all fixed
382 effects. The reported p -values correspond to Wald chi-square (χ^2) tests as implemented in
383 R. Reaction time (RT) analyses were done on error-free trials only and reflect the median
384 RT within each factor cell.

385 fMRI Scanning Protocol

386 Magnetic-resonance images were acquired using a 3-Tesla Siemens Prisma MRI scanner
387 (Siemens, Erlangen, Germany) located at the at the Princeton Neuroscience Institute. A
388 T2*- weighted echo-planar imaging (EPI) pulse sequence was used for functional imaging
389 (2×2 mm in plane resolution, TR = 3000 ms, TE = 27 ms, slice thickness = 2 mm, 53
390 slices, 96×96 matrix (FOV = 192 mm), iPAT factor: 3, flip angle = 80° , A→P phase
391 encoding direction). Slice orientation was tilted 30° backwards relative to the anterior
392 – posterior commissure axis to improve acquisition of data from the orbitofrontal cortex
393 (Deichmann2003). Field maps for distortion correction were acquired using the same
394 resolution (TE1 = 3.99ms) and a MPRAGE pulse sequence was used to acquire T1-weighted
395 images (voxel size = 0.9^3 mm). The experiment began 20 seconds after acquisition of the
396 first volume of each run to avoid partial saturation effects.

397 fMRI Data Preprocessing

398 FMRI data preprocessing was done using SPM8 (<http://www.fil.ion.ucl.ac.uk/spm>)
399 and involved fieldmap correction, realignment, and co-registration to the segmented struc-
400 tural images. The task data used to train the classifier were further submitted to a mass-
401 univariate general linear model that involved run-wise regressors for each state (see below),
402 nuisance regressors that reflected participant movement (6 regressors) and run-wise inter-
403 cepts. The resulting voxelwise parameter estimates were z-scored and spatially smoothed (4
404 mm FWHM). The resulting activation maps were used as the training set for a support-vector
405 machine with a radial basis function (RBF) kernel that was trained to predict the task state
406 from which a particular activation pattern came from (Chang2011). Like the activation

407 maps used for classifier training, the resting-state data were z-scored and smoothed (4mm
408 FWHM). Anatomical ROIs were created using SPM's `wfupick` toolbox. The hippocampus
409 (HC) was defined as the left and right hippocampus AAL labels. The orbitofrontal cortex
410 was defined as in (16). Behavioral analyses and computations within the assumed graphical
411 model of state space (see below) were done using R (48).

412 **fMRI Classification Analysis**

413 The support vector machines were trained on 8 maps of parameter estimates (“betas”) for
414 each of the 16 states (one map for each state and run) restricted to the anatomical mask
415 of the hippocampus (back-transformed into each subject’s individual brain space) or the
416 orbitofrontal cortex. Classification accuracy was assessed with a leave-one-run-out cross-
417 validation scheme in which data from 7 runs were used for training and the held-out run was
418 used for testing (Fig. 2). The resting-state analysis used a classifier trained on all available
419 task data (8 runs). This classifier was applied to each volume of the resting-state data
420 and the most highly classified state was considered as the output of the classifier for that
421 volume. The resulting sequence of predictions was the main focus of our analyses (see below).
422 We obtained the distance to the hyperplane by dividing the decision value by the norm of
423 the weight vector w , as specified in the libSVM webpage (<http://www.csie.ntu.edu.tw/~cjlin/libsvm/faq.html#f4151>). For each volume, we then calculated the average of the
425 distances of all pairwise comparisons of the predicted class against all other classes, to obtain
426 a proxy of how certain the classifier is in its prediction. Student t tests pertaining to decoding
427 results were one-tailed, given the *a priori* expectation of larger-than-chance decoding in the
428 hippocampus.

429 **Sequenceness Analysis**

430 The main question of the sequenceness analysis was whether the state transitions decoded
431 from resting state scans, T , were related to the distance between states experienced during
432 the task, D . To this end, we analyzed the neural state transitions T with logistic mixed-effects
433 models, using the `lme4` (47) package in R (48). Because the slow hemodynamic response
434 function leads to encoding of sequential structure in activity patterns (i.e., there is high
435 similarity between temporally adjacent patterns), a classifier trained on sequential task data
436 can be biased to decode states in a similar sequence to the training data, even if the test data
437 are random (i.e., the sequenceness comes from the training, not the test set). We therefore

438 applied the trained classifier to matched fMRI noise (see below) and used the resulting
439 spurious ‘state transitions,’ $T(\epsilon)$, as a covariate that would account of the spurious base
440 rate of transitions that is due to the classifier rather than the data. Applying these models
441 to control conditions consistently yielded non-significant effects of sequenceness, showing
442 that this analysis appropriately controls for the above mentioned spurious structure that is
443 observable for instance in the significant correlations between D and T in the noise data
444 (Fig. 3F). Specifically, our model included the following fixed effects: (1) the distance
445 between states, D , which was the regressor of interest, and as regressors of no interest
446 (2) the transition probabilities obtained in the above mentioned noise simulations, $T(\epsilon)$,
447 (3) an orthogonal quadratic polynomial of $T(\epsilon)$ that was included in order to account
448 for as much noise-related variance as possible, and (4) an intercept. Models of change
449 in sequenceness across PRE, INSTR and TASK conditions (Fig. 4) additionally involved
450 interaction terms of condition with the distance D and condition with the noise transitions
451 $T(\epsilon)$. Participant identity was included as a random factor to account for between subject
452 variability. To capture state-related variability (state frequency effects affect the distribution
453 of state transitions), state identity s_j of a transition from state i to state j was used as an
454 additional random effect nested within subject. Participant and state were random grouping
455 factors for all fixed effects with exception of the quadratic expansion of $T(\epsilon)$, where including
456 these random factors caused problems in fitting the logistic regression models.

Formally, the model followed the general assumption that the number of transitions Y is drawn from a binomial distribution of n draws and probability T :

$$Y_{ijk} \sim B(n_k, T_{ijk})$$

457 where n_k corresponds to the number of measurements for subject k , and i and j index
458 the outgoing and incoming states of a given transition. The logit transformed probabilities
459 T (shown in Fig. 2D; logit is the canonical link function for binomial models) were then
460 modeled in a mixed effects regression model with the above mentioned fixed and random
461 effects structure:

$$\begin{aligned} \text{logit}(T_{ijk}) = & \beta_0 + D_{ij}\beta_1 + T(\epsilon)_{ij}\beta_2 + T(\epsilon)_{ij}^2\beta_3 + \\ & \gamma_{0k} + D_{ij}\gamma_{1k} + T(\epsilon)_{ij}\gamma_{2k} + \\ & \zeta_{0kj} + D_{ij}\zeta_{1kj} + T(\epsilon)_{ij}\zeta_{2kj} + \epsilon_{ijk} \end{aligned}$$

462 In the text, we describe the fixed effect of D , β_1 in the models, as ‘sequenceness,’ and the
463 fixed effect of $T(\epsilon)$, β_2 , as ‘randomness’ (Fig. 4B,C). The subject-specific random effects of

464 D , γ_{1k} , were used as individual sequenceness indicators in the correlations in Fig. 4F,G. The
 465 state and subject specific random effects are indicated by ζ . Correlations between random
 466 effects were estimated. Model comparisons were conducted using likelihood-ratio tests by
 467 comparing base models including the noise transitions $T(\epsilon)$ with versus without the fixed
 468 effect regressor of distance (sequenceness), or without the condition interaction terms to the
 469 full models that included these terms. The random effects structure was kept constant across
 470 these comparisons.

471 T-tests pertaining to sequenceness results (number of steps, etc.) are one-tailed, given
 472 our *a priori* expectation of larger sequenceness in the hippocampus compared to the various
 473 controls.

474 Alternative Transition Functions

475 Alternative transition functions were computed directly from the true transition functions T .
 476 These alternatives were based on the assumption that the hippocampus has access to only
 477 partial state information, and hence correspond to transition functions defined over subsets
 478 of states. We define the set of all states \mathcal{S} :

$$\mathcal{S} = \{(Fy)Fy, (Fy)Fo, (Fy)Hy, (Fy)Ho, (Fo)Fy, (Fo)Fo, (Fo)Hy, (Fo)Ho, \\ (Hy)Fy, (Hy)Fo, (Hy)Hy, (Hy)Ho, (Ho)Fy, (Ho)Fo, (Ho)Hy, (Ho)Ho\}$$

479 To compute the transition function of the current trial model, for instance, we define
 480 that \mathcal{S}_{Fy}^{trial} is the subset of states that indicate that Fy was the current trial:

$$\mathcal{S}_{Fy}^{trial} = \{(Fy)Fy, (Fo)Fy, (Hy)Fy, (Ho)Fy\} \subseteq \mathcal{S}$$

481 \mathcal{S}_{Fo}^{trial} , \mathcal{S}_{Hy}^{trial} , \mathcal{S}_{Ho}^{trial} are the corresponding subsets of states with Fo, Hy and Ho as current
 482 trials. The transition function is then computed such that if a transition between states s_i
 483 and s_j exists, a transition between s_i and all states of the trial matched subset to which s_j
 484 belongs, $\mathcal{S}_{s_j}^{trial}$, exists with equal probability:

$$T_{s_i, s_j}^{trial} = \frac{1}{N} \sum_{\{s_x: s_x \in \mathcal{S}_{s_j}^{trial}\}} T_{s_i, s_x}$$

485 Following this procedure, we defined subsets of states that have the same current atten-
486 tion, and subsets of states that have the same current and past attention, and then computed
487 the transition functions as described above. The transition functions of the different models
488 are shown in Figure 5A-C. The reverse replay transition function was simply the transpose
489 of T .

490 Synthetic fMRI Data and Noise Simulations

491 In order to estimate to what extent training the classifiers on sequential data influenced
492 the sequenceness of their predictions, we simulated, for each participant, individually spatio-
493 temporally matched fMRI noise, and applied the classifiers to these data. For each partici-
494 pant and resting state session, we first extracted fMRI data from the hippocampus and the
495 orbitofrontal cortex. As in the classification analyses, we applied linear detrending to each
496 voxel. We then estimated the average standard deviation of the voxels within these regions, as
497 well as the average autocorrelation using an AR(1) model in R. Next, we used the neuRosim
498 toolbox in R (49) to simulate fMRI noise with the same standard deviation and temporal
499 autocorrelation as the real data. Finally, we used AFNI's 3dFWHMx and 3dBlurToFWHM
500 functions to first estimate the spatial smoothness of the real data, and then smooth the
501 simulated noise until it has the same effective smoothness. For each existing resting-state
502 run, matched noise data with the same number of TRs and voxels were generated. Figure
503 S1 (SI) shows the temporal and spatial smoothness of the real and simulated data separately
504 for each run. In all cases, the properties of the simulated data did not differ from the real
505 data, paired t-tests, all $ps > .05$.

506 Finally, we applied each participant's classifier to the matched noise data. The classifier
507 was identical to the classifier that was used in estimating the sequences of states from
508 the real data. The resulting sequence of predicted states reflects the bias of the classifier
509 to make sequential predictions because of pattern overlap in the training set, even when
510 applied to noise, as well as any tendency of the classifier to decode certain states more
511 often than others. We therefore used the sequence of states from this analysis to construct
512 the nuisance covariate for the mixed effects models, i.e. the noise 'transition matrix,' $T(\epsilon)$,
513 and to perform the appropriate comparisons in the correlation analysis. These comparisons
514 between sequenceness in real data versus simulated noise in the correlation and mixed effect
515 analyses indicated that the noise sequenceness $T(\epsilon)$ indeed explained a significant amount
516 of sequential variability of the decoded states (see Figs. 3F,G, 4B, D), and thus served as a
517 powerful control. Together with the permutation tests (Fig. 3B-D, 3F,G), the comparisons

518 across brain regions (Fig. 4E) and the within-participant comparisons between the PRE,
519 INSTR and TASK conditions (3B-D, H and 4A-D), our approach represents a stringent
520 control of potential biases.

521 Author Contributions

522 NWS, and YN designed research. NWS conducted research. NWS and YN analyzed and
523 interpreted the data and wrote the manuscript.

524 Acknowledgments

525 This research was funded by NIH grant R01DA042065 and funding from John Templeton
526 Foundation. The opinions expressed in this publication are those of the authors and do not
527 necessarily reflect the views of the John Templeton Foundation. We thank Nathaniel Daw,
528 Christian Doeller, Eran Eldar and David Tank for helpful comments on this research and/or
529 previous versions of the manuscript.

530 References

- 531 1. W. Scoville, B. Milner, Journal of Neurology, Neurosurgery & Psychiatry **20**, 11–21 (1957).
- 532 2. L. R. Squire, Psychological Review **99**, 195–231 (1992).
- 533 3. H Eichenbaum, P Dudchenko, E Wood, M Shapiro, H Tanila, Neuron **23**, 209–26 (1999).
- 534 4. J. O’Keefe, L. Nadel, The hippocampus as a cognitive map (Oxford University Press,
535 Oxford, 1978).
- 536 5. S. McKenzie et al., Neuron **83**, 202–215 (2014).
- 537 6. A. O. Constantinescu, J. X. O’Reilly, T. E. J. Behrens, Science **352**, 1464–1468 (2016).
- 538 7. R. M. Tavares et al., Neuron **87**, 231–243 (2015).
- 539 8. D. Aronov, R. Nevers, D. W. Tank, Nature **543**, 719–722 (2017).
- 540 9. D. J. Foster, J. J. Knierim, Current Opinion in Neurobiology **22**, 294–300 (2012).
- 541 10. A. M. Wikenheiser, A. D. Redish, Current Opinion in Neurobiology **32**, 8–15 (2014).
- 542 11. W. E. Skaggs, B. L. McNaughton, Science **271**, 1870–1873 (1996).
- 543 12. A. Johnson, D. Redish, Journal of Neuroscience **27**, 12176–89 (2007).
- 544 13. M. P. Karlsson, L. M. Frank, Nature Neuroscience **12**, 913–918 (2009).

- 545 14. M. F. Carr, S. P. Jadhav, L. M. Frank, Nature neuroscience **14**, 147–53 (2011).
- 546 15. G. Girardeau, K. Benchenane, S. I. Wiener, G. Buzsáki, M. B. Zugaro, Nature neuroscience
547 **12**, 1222–1223 (2009).
- 548 16. N. W. Schuck, M. B. Cai, R. C. Wilson, Y. Niv, Neuron **91**, 1402–1412 (2016).
- 549 17. S. Tong, E. Chang, presented at the Proceedings of the ninth ACM international conference
550 on Multimedia, vol. October, pp. 107–118, DOI: 10.1145/500156.500159.
- 551 18. P. Peigneux et al., **44**, 535–545 (2004).
- 552 19. L. Deuker et al., Journal of Neuroscience **33**, 19373–19383 (2013).
- 553 20. B. P. Staresina, A. Alink, N. Kriegeskorte, R. N. Henson,
554 Proceedings of the National Academy of Sciences of the United States of America **110**,
555 21159–64 (2013).
- 556 21. N. Axmacher, C. E. Elger, J. Fell, Brain **131**, 1806–1817 (2008).
- 557 22. R. C. Wilson, Y. K. Takahashi, G. Schoenbaum, Y. Niv, Neuron **81**, 267–79 (2014).
- 558 23. R. S. Sutton, presented at the Proceedings of the 7th International Conference on Machine
559 Learning, pp. 216–224, DOI: 10.1.1.51.7362.
- 560 24. H. van Seijen, R. S. Sutton,
561 Proceedings of the 32nd International Conference on Machine Learning **37** (2015).
- 562 25. E. M. Russek, I. Momennejad, M. M. Botvinick, S. J. Gershman, bioRxiv **083857**, DOI:
563 10.1101/083857 (2016).
- 564 26. R. Kaplan, N. W. Schuck, C. F. Doeller, Trends in Neurosciences **40**, 256–259 (2017).
- 565 27. T. O. Bergmann, M. Mölle, J. Diedrichs, J. Born, H. R. Siebner, NeuroImage **59**, 2733–2742
566 (2012).
- 567 28. Z. Kurth-Nelson, M. Economides, R. J. Dolan, P. Dayan, Neuron **91**, 194–204 (2016).
- 568 29. R. Cox, W. F. Hofman, M. de Boer, L. M. Talamini, NeuroImage **99**, 103–110 (2014).
- 569 30. B. Rasch, C. Buchel, S. Gais, J. Born, Science **315**, 1426–1429 (2007).
- 570 31. J. W. Antony, E. W. Gobel, J. K. O’Hare, P. J. Reber, K. A. Paller, Nature Neuroscience
571 **15**, 1114–1116 (2012).
- 572 32. A. G. Siapas, M. A. Wilson, Neuron **21**, 1123–1128 (1998).
- 573 33. D. Ji, M. A. Wilson, Nature neuroscience **10**, 100–107 (2007).
- 574 34. A. K. Lee, M. A. Wilson, Neuron **36**, 1183–1194 (2002).
- 575 35. D. Schiller et al., Journal of Neuroscience **35**, 13904–13911 (2015).

- 576 36. E. C. Tolman, Psychological Review **55**, 189–208 (1948).
- 577 37. N. J. Fortin, K. L. Agster, H. B. Eichenbaum, Nature neuroscience **5**, 458–62 (2002).
- 578 38. A. C. Schapiro, L. V. Kustner, N. B. Turk-Browne, Current Biology **22**, 1622–1627 (2012).
- 579 39. K. L. Stachenfeld, M. M. Botvinick, S. J. Gershman, Nature Neuroscience **20**, 1643–1653
580 (2017).
- 581 40. L. A. Bradfield, A. Dezfouli, M. van Holstein, B. Chieng, B. W. Balleine, Neuron, 1–13
582 (2015).
- 583 41. J. D. Howard, T. Kahnt, Journal of Neuroscience **37**, 3473–16 (2017).
- 584 42. R. Nogueira et al., Nature Communications **8**, 14823 (2017).
- 585 43. N. W. Schuck et al., Neuron **86**, 331–340 (2015).
- 586 44. C. F. Doeller, C. Barry, N. Burgess, Nature **463**, 657–61 (2010).
- 587 45. J Jacobs et al., Nat Neurosci **16**, 1188–1190 (2013).
- 588 46. N. C. Ebner, M. Riediger, U. Lindenberger, Behavior Research Methods **42**, 351–362 (2010).
- 589 47. D. Bates, M. Mächler, B. Bolker, S. Walker, Journal of Statistical Software **67**, 51 (2015).
- 590 48. R Core Team, R: A Language and Environment for Statistical Computing, Vienna, Austria,
591 2015.
- 592 49. M. Welvaert, J. Durnez, B. Moerkerke, G. Verdoolaege, Y. Rosseel,
593 Journal of Statistical Software **44**, 1–18 (2011).