Science Advances
AAAS

Supplementary Materials for

Model-based aversive learning in humans is supported by preferential task state reactivation

Toby Wise*, Yunzhe Liu, Fatima Chowdhury, Raymond J. Dolan

*Corresponding author. Email: t.wise@ucl.ac.uk

Published 28 July 2021, *Sci. Adv.* **7**, eabf9616 (2021)
DOI: 10.1126/sciadv.abf9616

This PDF file includes:

Supplementary materials

Model fitting

To verify that the G parameter in our model captured the kind of model-based choice we intended, we correlated it with an approximate behavioural index of generalisation, namely choice consistency between learning and generalisation trials. If subjects are indeed generalising based on learned value, their choices on generalisation trials should approximate those on learning trials immediately preceding these trials. In line with this, our generalisation parameter strongly correlated with the proportion of trials where generalisation choices were consistent with preceding non-generalisation choices for each subject ($r = 0.59$, $p = 0.001$, Figure 1F), indicating this parameter provides a valid index of generalisation. Finally, it is possible that subjects may use model-based inference due to poor general task knowledge, rather than absence of a model-based generalisation *per se*. To ascertain this was not the case, we examined the relationship between generalisation parameter values and the number of mistakes subjects made on generalisation trials (i.e. not selecting any option or making an incorrect state selection). If subjects do not generalise due to poor task knowledge, we would expect a negative correlation between these variables. These two variables were not correlated ($r = -.06$, $p = 0.75$, Figure 1F), indicating that low generalisation parameter values are not simply a reflection of poor task knowledge or execution.

Classifier accuracy

In Figure S1A, we show the timecourse of classification accuracy relative to stimulus onset during the localiser task for the 14-way classifier used in the sequenceness analyses. Notably, classification accuracy is substantially higher using a temporal embedding approach whereby data from timepoints +/-50ms either side of the timepoint of interest are provided to the classifier. Accuracy peaks at around 150ms, but is above chance until the end of a 0.8 second window tested. Generalisation to other time points was limited (Figure S1B) indicating that classifiers trained on each time point were detecting distinct stimulus characteristics from those trained on other timepoints. When investigating the spatial distribution of classifier weights in sensor space we observed that the strongest contributions in the

classifier came from occipital and temporal sensors across classifiers trained at 150ms, 370ms and 520ms timepoints (Figure S1C).

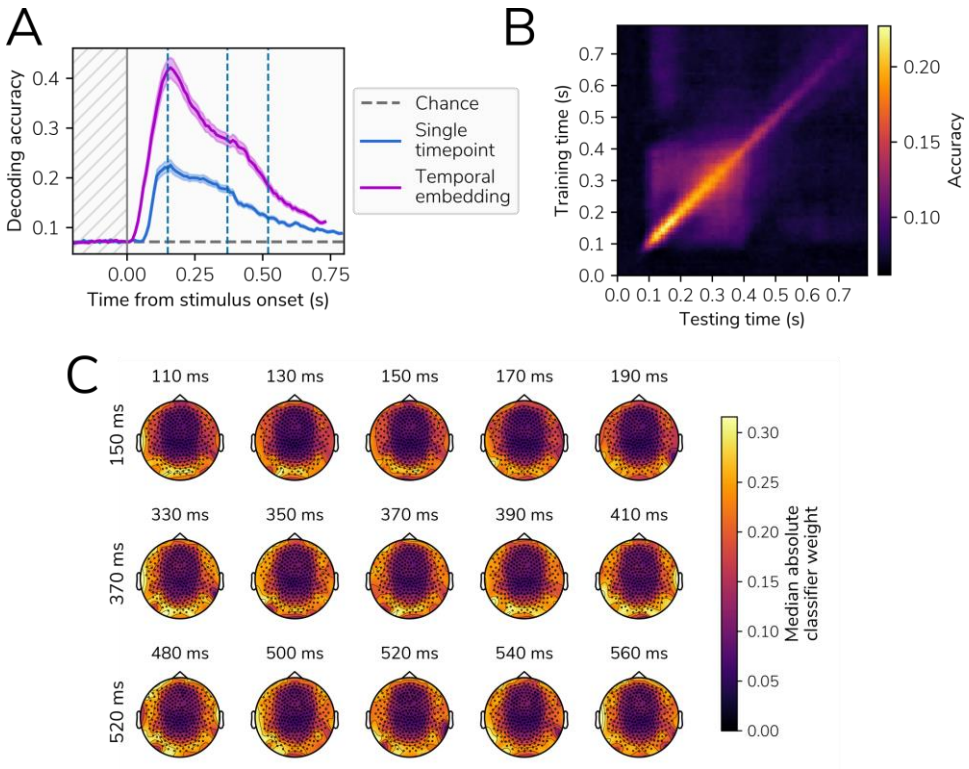

Figure S1. Classifier performance. A) Classification accuracy (for the 14-way classifier used in sequenceness analyses) from 0.2s prior to stimulus onset to 0.8s post stimulus onset in the localiser task. Accuracy is shown for classifiers trained on a single timepoint and trained using temporal embedding, whereby data from +/-50ms relative to the timepoint of interest are provided to the classifier. Chance is calculated as 1/14 (7.14%). B) Generalisation of stimulus classification across time points, showing how classifiers trained on a specific timepoint (Y axis), without temporal embedding, perform when tested on different time points (X axis). C) Approximate spatial distribution of classification weights for classifiers trained on the 150ms, 370ms, and 520ms time points. As classifiers were trained on reduced data following PCA, maps were obtained by projecting the weight of each principal component on to the original sensor space (272 channels) to approximate the weight of each individual sensor. The absolute value of these weights was then taken to indicate the contribution of each sensor (regardless of the weight being positive or negative), and the median value across stimuli and subjects was displayed.

To determine the accuracy of the 14-way classifiers used for determining state reactivation in the sequenceness analyses, we compare the accuracy of predictions in the localiser data, using 5-fold cross-validation. Confusion matrices, collapsed across subjects, are shown in Figure S1 for each time-point used for training. Classification accuracy was above chance (1/14 = 7.14%) for every stimulus.

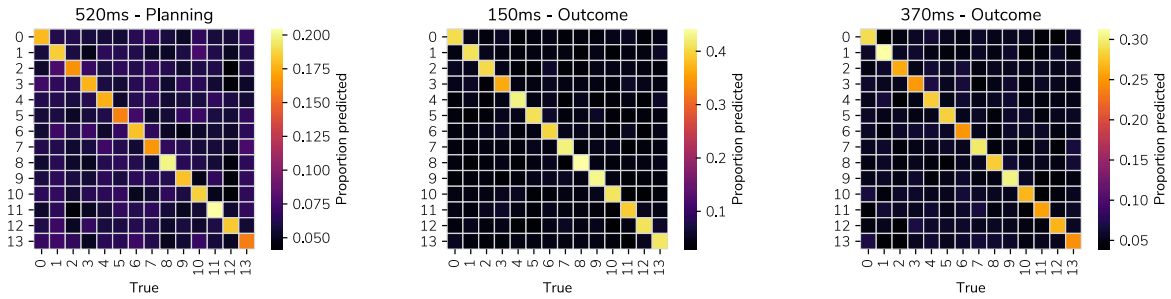

Figure S2. Confusion matrices representing classification accuracy across different training timepoints. Values show accuracy collapsed across subjects.

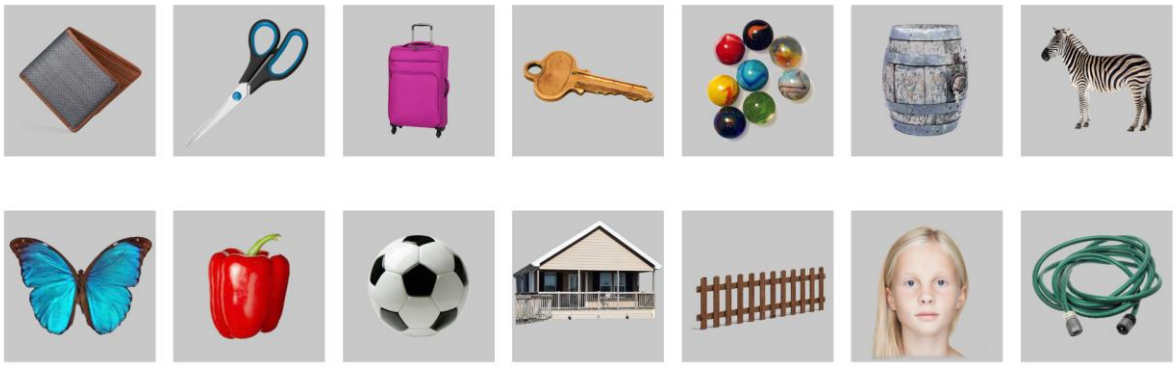

Figure S3. Images used to represent task states.

Associations between model-based planning and hippocampal theta

We tested whether the G parameter from our model, an index of engagement in model-based planning, was associated with theta power in the hippocampus, or with coupling between theta power and reactivation (as indexed by beta coefficients from our source-level regression models). We first addressed this by collapsing across the duration of the trial and examining the relationship with overall theta power and coupling across trials, but found no significant correlations with G (table S1). We next tested whether there were clusters within the trial where either measure was significantly associated with G. This revealed no significant clusters (Figure S4).

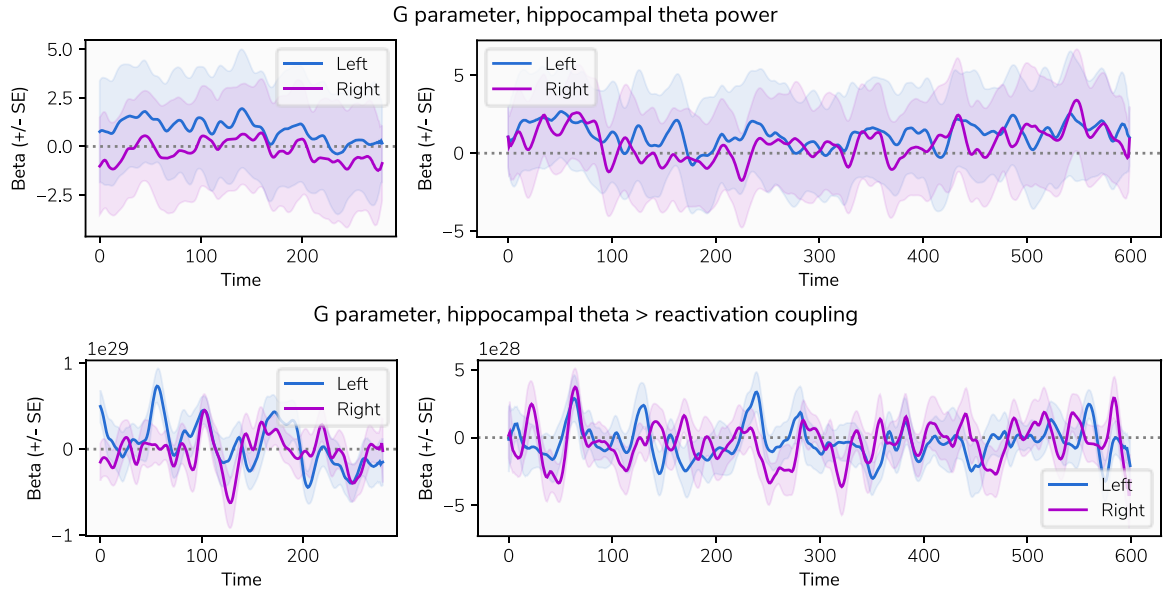

Figure S4. Neural associations with G parameter. Association between G parameter and theta power (top) and theta-reactivation coupling (bottom) within outcome (left) and planning (right) periods. Beta values on the y-axes represent the effect from a linear model predicting theta power or coupling from G value across subjects.

Table S1. Results of correlation analyses assessing relationships between the G parameter and theta power and theta power-state reactivation coupling, collapsing across the trial.

Measure	Task phase	Hemisphere	*β*	*P (uncorrected)*
Coupling	Outcome	L	0.21	0.29
		R	-0.08	0.69
	Planning	L	-0.22	0.27
		R	-0.26	0.18
Power	Outcome	L	0.07	0.73
		R	-0.02	0.92
	Planning	L	0.07	0.71
		R	0.04	0.83

Within-trial analyses of reactivation, replay and theta power.

Full results of our analyses investigating links between hippocampal activity, reactivation, and replay are shown in Table S2, with the strongest results in the theta band shown in Figure S5.

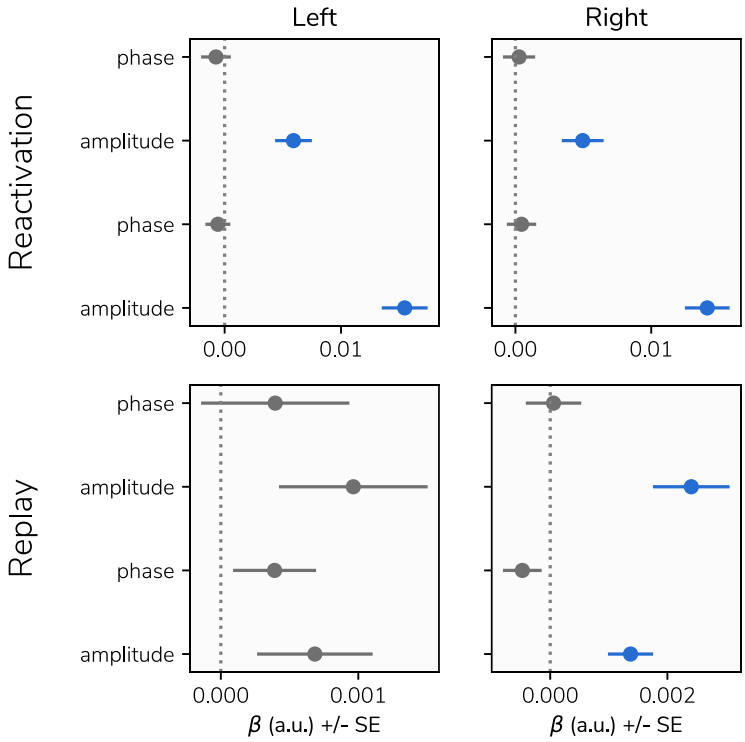

Figure S5. Results of autoregressive models predicting theta amplitude from reactivation and replay intensity within-trial. Significant effects, corrected for multiple comparisons using false discovery rate correction, and highlighted in blue.

Table S2. Results of all autoregressive models used to assess relationships between reactivation, replay, and hippocampal signals. P values are corrected for multiple comparisons across all models using FDR correction, and significant effects are highlighted in bold.

Measure	Task phase	Hippocampus metric	Frequency band	Hemi	β	t	p
Reactivation	Outcome	Amplitude	High gamma	L	0.01	0.04	0.967
				R	0.31	2.24	0.107
			Low gamma	L	**3.39**	**3.41**	**0.009**
				R	**4.52**	**4.50**	**0.001**
		Amplitude	Theta	L	**15.47**	**7.85**	**0.000**
				R	**14.12**	**8.58**	**0.000**
		Phase	Theta	L	-0.58	-0.54	0.790
				R	0.45	0.42	0.871
	Planning	Amplitude	High gamma	L	-0.10	-0.73	0.686
				R	-0.44	-2.18	0.111
		Amplitude	Low gamma	L	-0.16	-0.18	0.967
				R	1.31	1.49	0.294
		Amplitude	Theta	L	**5.92**	**3.75**	**0.007**
				R	**4.95**	**3.21**	**0.013**
		Phase	Theta	L	-0.75	-0.58	0.786
				R	0.27	0.23	0.967
Replay	Outcome	Amplitude	High gamma	L	0.00	-0.07	0.967
				R	0.00	0.06	0.967
			Low gamma	L	-0.38	-2.58	0.055
				R	-0.28	-1.75	0.211
		Amplitude	Theta	L	0.68	1.63	0.246
				R	**1.37**	**3.56**	**0.008**
		Phase	Theta	L	0.39	1.30	0.366
				R	-0.48	-1.44	0.301
	Planning	Amplitude	High gamma	L	0.03	0.85	0.667
				R	0.01	0.35	0.897
		Amplitude	Low gamma	L	-0.46	-2.10	0.120
				R	0.21	0.82	0.667
		Amplitude	Theta	L	0.96	1.78	0.211
				R	**2.41**	**3.68**	**0.007**
		Phase	Theta	L	0.40	0.73	0.686
				R	0.06	0.12	0.967

Power spectral density of reactivation and replay timecourses

We conducted further exploratory analyses investigating whether the state reactivation and replay timecourses themselves exhibited particularly strong rhythmicity across frequency bands. To achieve this, we calculated the power spectral density across frequencies from 0.5 to 50hz using Welch's method. Full power spectra are shown in Figure S6A). We used an established method to decompose the spectra into an aperiodic component (representing 1/f noise), and periodic signals (61). This revealed the presence of a strong alpha component and a weaker beta component across both replay and reactivation timeseries. During the outcome phase, replay timeseries also showed evidence of theta rhythmicity, but this was not observed in the planning phase. Reactivation timeseries also showed evidence of a component in the gamma range. Estimates of power in each band are shown in Figure S6B

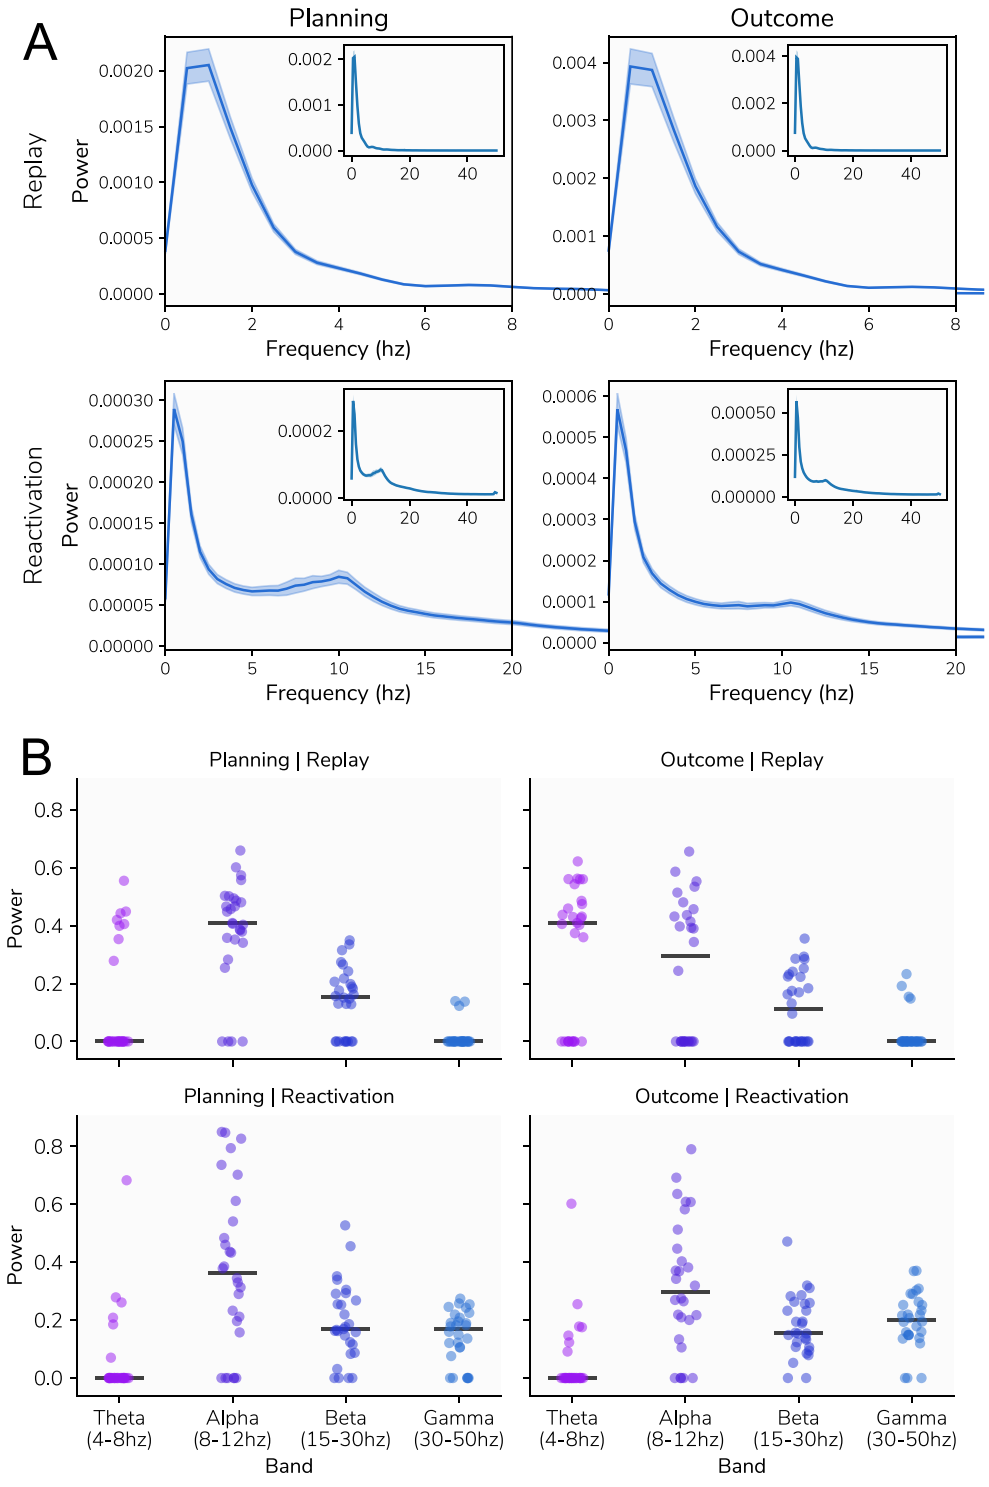

Figure S6. Analyses of power spectra within replay timeseries. A) Power spectra for the replay (top) and state reactivation (bottom) time courses. Full spectra are shown in the inset figures, while the main figures focus on the ranges exhibiting prominent peaks (0-8hz for replay, 0-20hz for reactivation). B) Estimates of power within each band, accounting for 1/f noise. Each point represents the power estimate for a subject and power estimates of zero indicate that no periodic component in this band was detected. Bars represent the median across subjects within each band.

Prediction accuracy in the planning phase

Although we observed a significant association between prediction accuracy and the G parameter from our behavioural model in our temporal generalisation analyses during the planning phase, we found no significant predictive effects (i.e. successfully predicting the chosen option using a classifier trained to distinguish between end states). Predictive accuracy maps are shown in Figure S7A. We also observed no significant effect when predicting predictive accuracy from the G parameter during learning trials, as shown in Figure S7B.

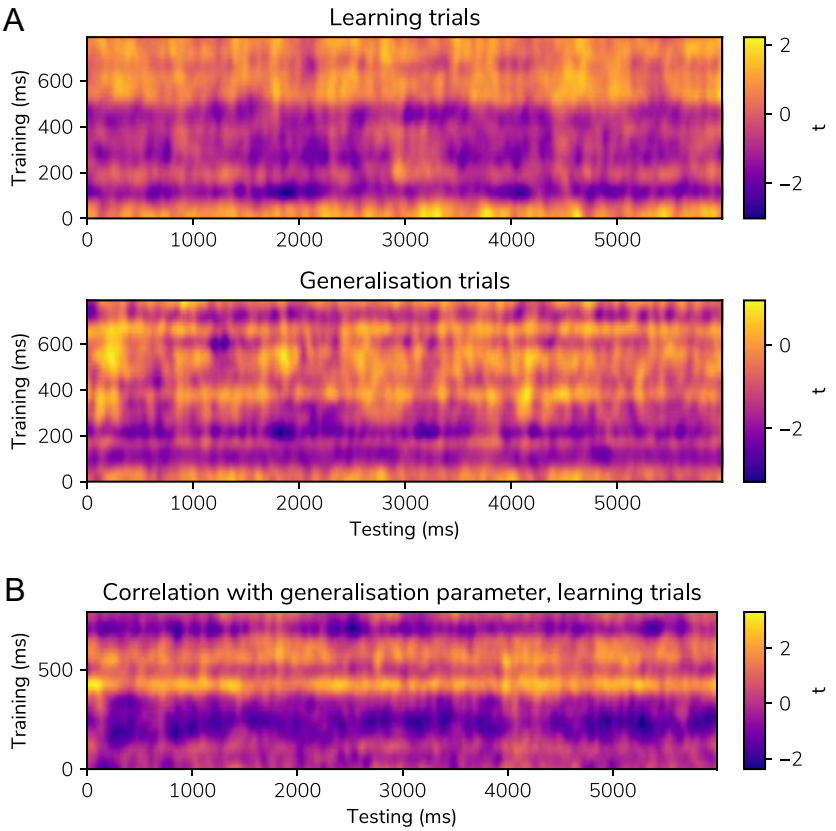

Figure S7. Temporal generalisation analyses showing no significant effects. A) Predictive accuracy in the planning phase, predicting a subsequently chosen option based on a classifier trained to distinguish between end states. Values represent t statistics from a one-sample t-test across subjects against a chance value of 0.5, with higher values indicating reactivation of the terminal state for the subsequently chosen path and lower values representing reactivation of the subsequently unchosen terminal state. No significant clusters were found. B) Correlation between G parameter and predictive accuracy on learning trials. All effects are non-significant.

Phase-amplitude coupling

We also performed additional exploratory analyses to examine the role of theta-gamma phase amplitude coupling in state reactivation, given the known role of theta-gamma coupling in learning and memory formation (62, 63). To assess coupling between theta phase and gamma amplitude, we extracted unfiltered timeseries from the bilateral hippocampus. We then used the approach developed by Tort et al (64), as implemented in pactools (https://github.com/pactools/pactools) to measure the degree of coupling between theta (4-8hz) phase and gamma (30-200hz) amplitude across the entire trial for planning and outcome phases independently. As our primary question pertained to the effect of state reactivation, we performed this analysis separately for trials showing high and low levels of state reactivation, using the same reactivation strength metric as for our analyses investigating source-level prediction of reactivation strength. A median split was then used to classify trials as high or low reactivation. This resulted in a theta frequency X gamma frequency matrix of coupling values for each phase and hemisphere, for each subject. These were then averaged across subject to produce a group-level matrix. To determine statistical significance, we repeated this process 100 times on data with permuted labels to provide an empirical null distribution for the difference in coupling between high and low replay trials. We set the significance threshold at .0125 (.05 / 4 tests) and determined effects to be significant if they exceeded the 99.375^{th} percentile, or did not exceed the 0.625^{th} percentile, across the entire 3D null matrix (permutations X theta phase X gamma phase). This conservative approach served to protect against the risk of false positives due to multiple comparisons.

Results of this analysis are shown in Figure S8. This revealed evidence of stronger coupling between low theta (4-6hz) phase and gamma amplitude across multiple frequency bands (40-50hz, 60-75hz, 140-150hz, 190-200hz) during the planning phase when the subsequently chosen option was reactivated to a greater extent. There was also evidence of stronger phase-amplitude coupling between low theta (4-5hz) and high gamma (110-170hz) in the outcome phase for trials when the previously chosen option was reactivated more strongly, but these effects were weaker than those seen during the planning phase. However, we note that these analyses should be interpreted with caution given the limited time windows and trial numbers available for the calculation of phase-amplitude coupling.

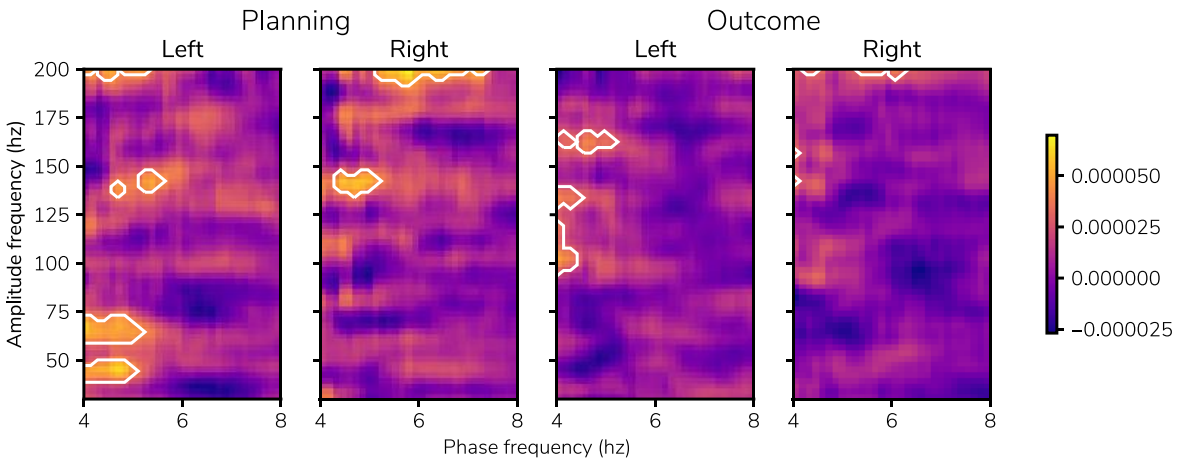

Figure S8. Results of phase-amplitude coupling analysis. These figures show the difference in theta-gamma coupling in the hippocampus between high and low reactivation trials. Brighter colours represent greater coupling in trials with high levels of reactivation, and highlighted areas represent significant effects.

Validation of hierarchical latent Gaussian process regression

To validate the statistical approach used for testing time-varying sequenceness, we evaluated its performance on simulated data. The method itself is described in full in the Methods section. We generate data according to a regression model where the regressors are represented by time-varying functions. Thus, the sequenceness value on each trial is represented by the value of that trial on a given variable (this may be trial number, for example), multiplied by its regression coefficient. As we expect these regression coefficients to vary over the course of the trial (for example the effect of trial number may be greatest at the start of the trial) with a degree of autocorrelation, representing the regressors as functions of time provides a simple model of the data-generating process. We used 100 simulated trials, and 100 timepoints.

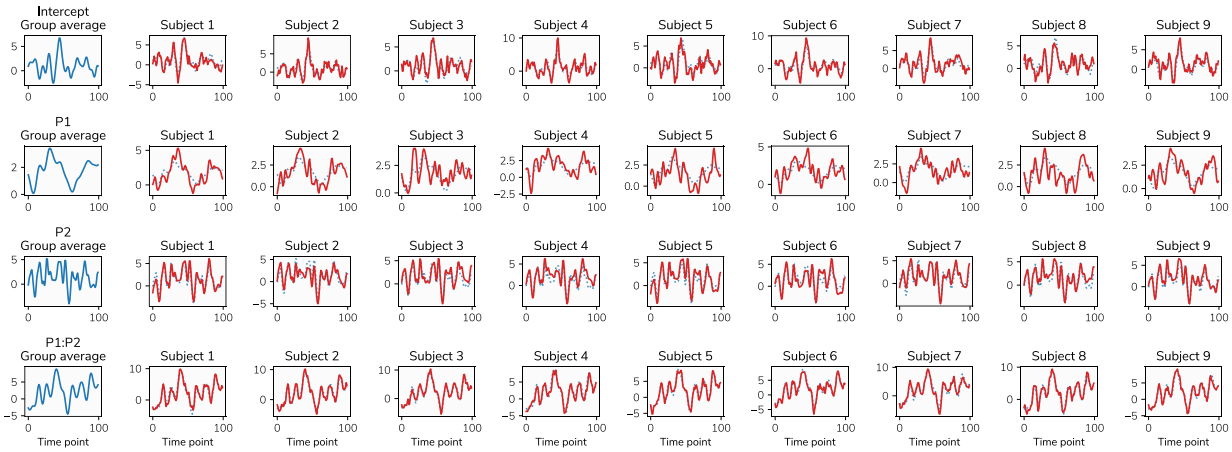

Figure S9. Functions used to generate simulated data. The blue lines on the left of the figure are the group-level functions. The red lines on the right side are the subject-level functions, with the group-level function shown in the dotted blue line.

For the purpose of validation, we simulate data from four known functions determined by Gaussian processes with pre-determined covariance functions, with the first representing an intercept, two representing separate regressors, and the final representing the interaction between the previous two regressors. We structure the data in a hierarchical manner, with these functions representing the group-level process. The covariance function of each Gaussian process has its own length scale parameter (values of 3, 5, 2, 4), and its own variance (2, 0.8, 2, 3). The mean function of these Gaussian processes is set to a constant (with values of 1.5, 2, 0.3, 0.7), representing an effect that does not depend on time.

We then determine 20 subject-level functions (assuming 20 simulated subjects) for each regressor, which are represented by the sum of the group level function and a separate, subject level function.

Each subject-level function is drawn from another Gaussian process, again with separate length scale (3, 5, 2, 4) and variance (1, 0.8, 0.9, 0.8) parameters for the covariance function. Thus, each subject's function is offset from the group-level function by a subject-specific function. The length scale parameters represent the degree of correlation between adjacent timepoints, while the variance parameters have the effect of representing the range of the function. A greater subject-level variance relative to the group-level variance has the effect of increasing the subject-level variability and reducing the influence of the group-level function. Examples of these functions are shown in Figure S2. Finally, to represent error, we add noise determined by a Gaussian distribution with zero mean and a standard deviation of 20. We also repeated this process using simulated data with no effect at the group level; that is, subjects had their own effects but these were offset from a group effect that was zero across all time points.

Fitting the models of the form used for the sequenceness analyses to this simulated dataset allowed us to recover the true group-level functions used to generate the data (for the purposes of these analyses, we are interested only in the group-level effect). As in the analyses reported in the main results section, we used the 99.9% highest posterior density intervals (HPDI) of the posterior distributions over these functions to obtain a conservative metric through which to determine whether the recovered interval includes the true function. Calculating the number of timepoints where the HPDI includes the true function value provides a measure of the false positive rate. The results of this test are shown in Figure S3.

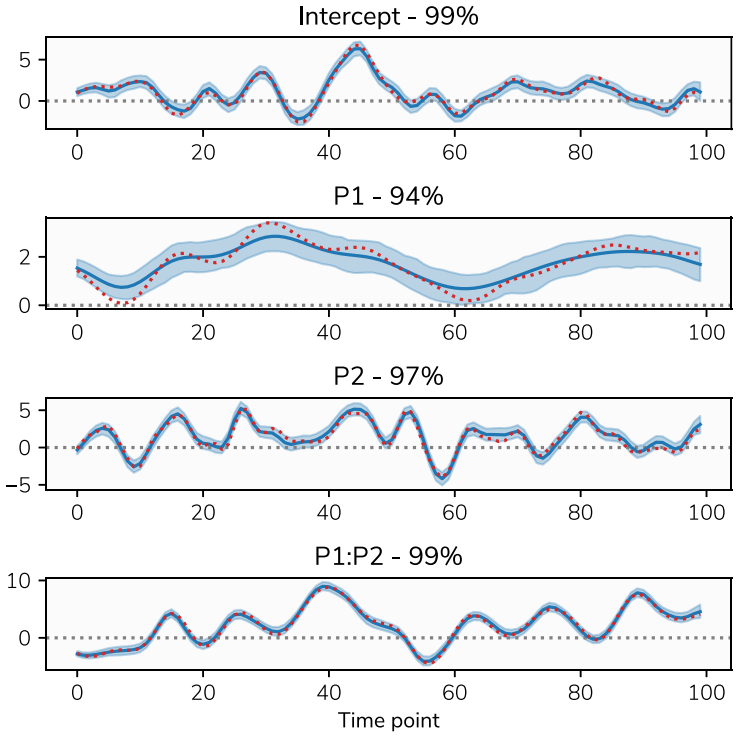

Figure S10. Means and 99.9% HDPIs of the recovered group-level function, shown in blue. The red dotted line represents the true group-level function.

As shown in Figure S2, the true function was included in the HPDI in 97.25% of time points across the four regressors, indicating that while the HDPI is not perfectly calibrated with respect to the false positive rate, using the conservative 99.9% HPDI provides acceptable control of the false positive rate. As shown in Figure S4, repeating this with a null effect at the group level (i.e. an effect of zero across all time points), the 99.9% HPDI includes zero at 100% of time points.

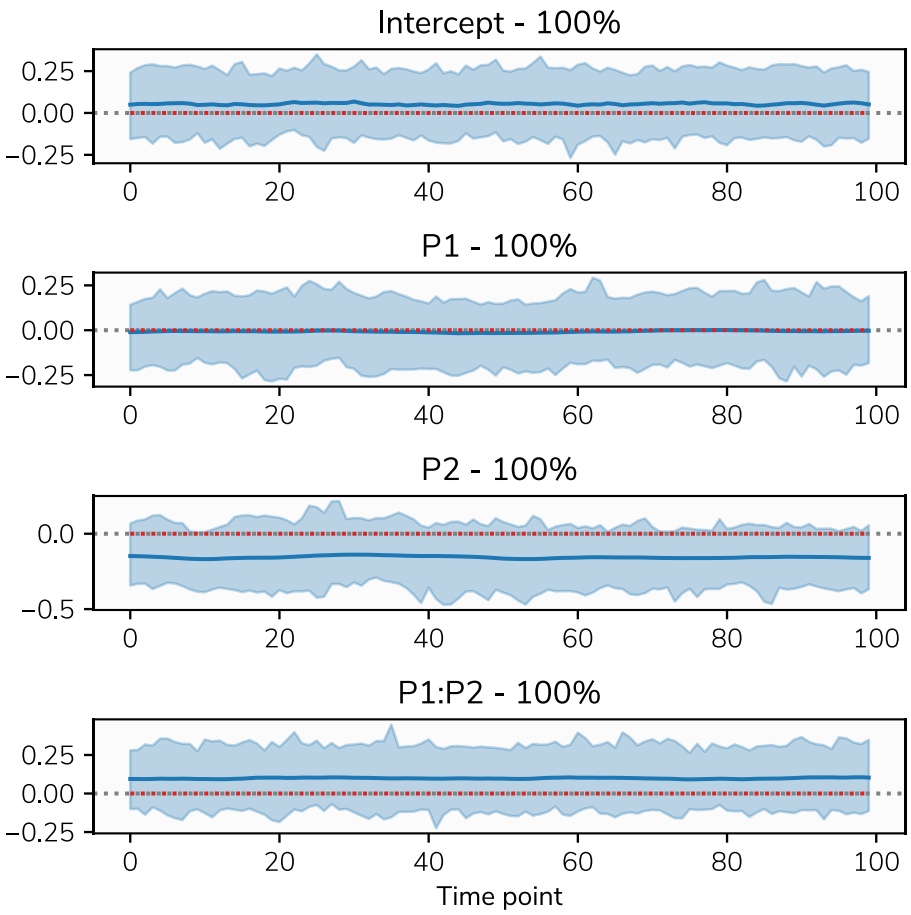

Figure S11. Means and 99.9% HDPIs of the recovered group-level function (showing no effect across the trial), shown in blue. The red dotted line represents the true group-level function (zero across the entire trial).

REFERENCES AND NOTES

1. D. L. Schacter, R. G. Benoit, K. K. Szpunar, Episodic future thinking: Mechanisms and functions. *Curr. Opin. Behav. Sci.* **17**, 41–50 (2017).

2. R. E. Ambrose, B. E. Pfeiffer, D. J. Foster, Reverse replay of hippocampal place cells is uniquely modulated by changing reward. *Neuron* **91**, 1124–1136 (2016).

3. H. F. Ólafsdóttir, D. Bush, C. Barry, The role of hippocampal replay in memory and planning. *Curr. Biol.* **28**, R37–R50 (2018).

4. C.-T. Wu, D. Haggerty, C. Kemere, D. Ji, Hippocampal awake replay in fear memory retrieval. *Nat. Neurosci.* **20**, 571–580 (2017).

5. C. Gagne, P. Dayan, S. J. Bishop, When planning to survive goes wrong: Predicting the future and replaying the past in anxiety and PTSD. *Curr. Opin. Behav. Sci.* **24**, 89–95 (2018).

6. A. S. Heller, R. C. Bagot, Is hippocampal replay a mechanism for anxiety and depression? *JAMA Psychiat.* **77**, 431–432 (2020).

7. T. D. Borkovec, W. J. Ray, J. Stober, Worry: A cognitive phenomenon intimately linked to affective, physiological, and interpersonal behavioral processes. *Cogn. Ther. Res.* **22**, 561–576 (1998).

8. B. B. Doll, K. D. Duncan, D. A. Simon, D. Shohamy, N. D. Daw, Model-based choices involve prospective neural activity. *Nat. Neurosci.* **18**, 767–772 (2015).

9. G. E. Wimmer, C. Büchel, Learning of distant state predictions by the orbitofrontal cortex in humans. *Nat. Commun.* **10**, 2554 (2019).

10. Z. Kurth-Nelson, G. Barnes, D. Sejdinovic, R. Dolan, P. Dayan, Temporal structure in associative retrieval. *eLife* **4**, e04919 (2015).

11. G. E. Wimmer, D. Shohamy, Preference by association: How memory mechanisms in the hippocampus bias decisions. *Science* **338**, 270–273 (2012).

12. I. Momennejad, A. R. Otto, N. D. Daw, K. A. Norman, Offline replay supports planning in human reinforcement learning. *eLife* **7**, e32548 (2018).

13. E. Eldar, G. Lièvre, P. Dayan, R. J. Dolan, The roles of online and offline replay in planning. *eLife* **9**, e56911 (2020).

14. Z. Kurth-Nelson, M. Economides, R. J. Dolan, P. Dayan, Fast sequences of non-spatial state representations in humans. *Neuron* **91**, 194–204 (2016).

15. Y. Liu, R. J. Dolan, Z. Kurth-Nelson, T. E. J. Behrens, Human replay spontaneously reorganizes experience. *Cell* **178**, 640–652.e14 (2019).

16. Y. Liu, M. G. Mattar, T. E. J. Behrens, N. D. Daw, R. J. Dolan, Experience replay is associated with efficient nonlocal learning. *Science* **372**, eabf1357 (2021).

17. N. W. Schuck, Y. Niv, Sequential replay of nonspatial task states in the human hippocampus. *Science* **364**, eaaw5181 (2019).

18. G. E. Wimmer, Y. Liu, N. Vehar, T. E. J. Behrens, R. J. Dolan, Episodic memory retrieval success is associated with rapid replay of episode content. *Nat. Neurosci.* **23**, 1025–1033 (2020).

19. G. Castegnetti, A. Tzovara, S. Khemka, F. Melinščak, G. R. Barnes, R. J. Dolan, D. R. Bach, Representation of probabilistic outcomes during risky decision-making. *Nat. Commun.* **11**, 2419 (2020).

20. D. R. Bach, P. Dayan, Algorithms for survival: a comparative perspective on emotions. *Nat. Rev. Neurosci.* **18**, 311–319 (2017).

21. D. Mobbs, D. B. Headley, W. Ding, P. Dayan, Space, time, and fear: survival computations along defensive circuits. *Trends Cogn. Sci.* **24**, 228–241 (2020).

22. S. Zorowitz, I. Momennejad, N. D. Daw, Anxiety, avoidance, and sequential evaluation. *Comput. Psychiat.* **4**, 1–17 (2020).

23. R. A. Rescorla, A. R. Wagner, A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement, in *Classical Conditioning II: Current Research and Theory*, A. H. Black, W. F. Prokasy, Eds. (Appleton-Century-Crofts, 1972), vol. 2, pp. 64–99.

24. T. Wise, R. J. Dolan, Associations between aversive learning processes and transdiagnostic psychiatric symptoms in a general population sample. *Nat. Commun.* **11**, 4179 (2020).

25. T. Wise, J. Michely, P. Dayan, R. J. Dolan, A computational account of threat-related attentional bias. *PLoS Comput. Biol.* **15**, e1007341 (2019).

26. S. Watanabe, Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory. *J. Mach. Learn. Res.* **11**, 3571–3594 (2010).

27. E. Eldar, T. U. Hauser, P. Dayan, R. J. Dolan, Striatal structure and function predict individual biases in learning to avoid pain. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 4812–4817 (2016).

28. Y. Liu, R. J. Dolan, C. Higgins, H. Penagos, M. W. Woolrich, H. Freyja Ólafsdóttir, C. Barry, Z. Kurth-Nelson, T. E. Behrens, Temporally delayed linear modelling (TDLM) measures replay in both animals and humans. *eLife* **10**, e66917 (2021).

29. E. Eldar, G. J. Bae, Z. Kurth-Nelson, P. Dayan, R. J. Dolan, Magnetoencephalography decoding reveals structural differences within integrative decision processes. *Nat. Hum. Behav.* **2**, 670–681 (2018).

30. D. J. Foster, M. A. Wilson, Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680 (2006), 683.

31. M. G. Mattar, N. D. Daw, Prioritized memory access explains planning and hippocampal replay. *Nat. Neurosci.* **21**, 1609–1617 (2018).

32. K. Z. Tanaka, A. Pevzner, A. B. Hamidi, Y. Nakazawa, J. Graham, B. J. Wiltgen, Cortical representations are reinstated by the hippocampus during memory retrieval. *Neuron* **84**, 347–354 (2014).

33. A. R. Backus, J.-M. Schoffelen, S. Szebényi, S. Hanslmayr, C. F. Doeller, Hippocampal-prefrontal theta oscillations support memory integration. *Curr. Biol.* **26**, 450–457 (2016).

34. R. Kaplan, D. Bush, M. Bonnefond, P. A. Bandettini, G. R. Barnes, C. F. Doeller, N. Burgess, Medial prefrontal theta phase coupling during spatial memory retrieval. *Hippocampus* **24**, 656–665 (2014).

35. S. Khemka, G. Barnes, R. J. Dolan, D. R. Bach, Dissecting the function of hippocampal oscillations in a human anxiety model. *J. Neurosci.* **37**, 6869–6876 (2017).

36. B. D. Van Veen, W. Van Drongelen, M. Yuchtman, A. Suzuki, Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Trans. Biomed. Eng.* **44**, 867–880 (1997).

37. J. O'Neill, T. J. Senior, K. Allen, J. R. Huxter, J. Csicsvari, Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nat. Neurosci.* **11**, 209–215 (2008).

38. G. R. Sutherland, B. McNaughton, Memory trace reactivation in hippocampal and neocortical neuronal ensembles. *Curr. Opin. Neurobiol.* **10**, 180–186 (2000).

39. K. Diba, G. Buzsáki, Forward and reverse hippocampal place-cell sequences during ripples. *Nat. Neurosci.* **10**, 1241–1242 (2007).

40. I. Momennejad, Learning structures: predictive representations, replay, and generalization. *Curr. Opin. Behav. Sci.* **32**, 155–166 (2020).

41. A. O. Constantinescu, J. X. O'Reilly, T. E. J. Behrens, Organizing conceptual knowledge in humans with a gridlike code. *Science* **352**, 1464–1468 (2016).

42. S. A. Park, D. S. Miller, H. Nili, C. Ranganath, E. D. Boorman, Map making: constructing, combining, and inferring on abstract cognitive maps. *Neuron* **107**, 1226–1238.e8 (2020).

43. A. K. Barbey, F. Krueger, J. Grafman, Structured event complexes in the medial prefrontal cortex support counterfactual representations for future planning. *Philos. Trans. R. Soc. B* **364**, 1291–1300 (2009).

44. N. Camille, G. Coricelli, J. Sallet, P. Pradat-Diehl, J.-R. Duhamel, A. Sirigu, The involvement of the orbitofrontal cortex in the experience of regret. *Science* **304**, 1167–1170 (2004).

45. G. Coricelli, H. D. Critchley, M. Joffily, J. P. O'Doherty, A. Sirigu, R. J. Dolan, Regret and its avoidance: A neuroimaging study of choice behavior. *Nat. Neurosci.* **8**, 1255–1262 (2005).

46. E. A. Boeke, J. M. Moscarello, J. E. LeDoux, E. A. Phelps, C. A. Hartley, Active avoidance: Neural mechanisms and attenuation of pavlovian conditioned responding. *J. Neurosci.* **37**, 4808–4818 (2017).

47. S. Qi, L. Cross, T. Wise, X. Sui, J. O'Doherty, D. Mobbs, The role of the medial prefrontal cortex in spatial margin of safety calculations. *bioRxiv*, 2020.06.05.137075 (2020).

48. G. E. Wimmer, C. Büchel, Reactivation of reward-related patterns from single past episodes supports memory-based decision making. *J. Neurosci.* **36**, 2868–2880 (2016).

49. N. L. Kocovski, N. S. Endler, N. A. Rector, G. L. Flett, Ruminative coping and post-event processing in social anxiety. *Behav. Res. Ther.* **43**, 971–984 (2005).

50. N. J. Roese, J. M. Olson, in *Advances in Experimental Social Psychology*, M. P. Zanna, Ed. (Academic Press, 1997), vol. 29, pp. 1–59.

51. W. C. Drevets, J. L. Price, J. R. Simpson, R. D. Todd, T. Reich, M. Vannier, M. E. Raichle, Subgenual prefrontal cortex abnormalities in mood disorders. *Nature* **386**, 824–827 (1997).

52. Green S, Lambon Ralph MA, Moll J, Deakin JW, Zahn R, Guilt-selective functional disconnection of anterior temporal and subgenual cortices in major depressive disorder. *Arch. Gen. Psychiatry* **69**, 1014–1021 (2012).

53. T. Wise, A. J. Cleare, A. Herane, A. H. Young, D. Arnone, Diagnostic and therapeutic utility of neuroimaging in depression: an overview. *Neuropsychiatr. Dis. Treat.* **10**, 1509–1522 (2014).

54. T. Wise, J. Radua, E. Via, N. Cardoner, O. Abe, T. M. Adams, F. Amico, Y. Cheng, J. H. Cole, C. de A. M. Périco, D. P. Dickstein, T. F. D. Farrow, T. Frodl, G. Wagner, I. H. Gotlib, O. Gruber, B. J. Ham, D. E. Job, M. J. Kempton, M. J. Kim, P. C. M. P. Koolschijn, G. S. Malhi, D. Mataix-Cols, A. M. McIntosh, A. C. Nugent, J. T. O'Brien, S. Pezzoli, M. L. Phillips, P. S. Sachdev, G. Salvadore, S. Selvaraj, A. C. Stanfield, A. J. Thomas, M. J. van Tol, N. J. A. van der Wee, D. J. Veltman, A. H. Young, C. H. Fu, A. J. Cleare, D. Arnone, Common and distinct patterns of grey-matter volume alteration in major depression and bipolar disorder: Evidence from voxel-based meta-analysis. *Mol. Psychiatry* **22**, 1455 (2017).

55. A. Gramfort, M. Luessi, E. Larson, D. A. Engemann, D. Strohmeier, C. Brodbeck, R. Goj, M. Jas, T. Brooks, L. Parkkonen, M. Hämäläinen, MEG and EEG data analysis with MNE-Python. *Front. Neurosci.* **7**, 267 (2013).

56. F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, É. Duchesnay, Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).

57. J. Bergstra, Y. Bengio, Random search for hyper-parameter optimization. *J. Mach. Learn. Res.* **13**, 281–305 (2012).

58. J. Hensman, N. D. Lawrence, M. Rattray, Hierarchical Bayesian modelling of gene expression time series across irregularly sampled replicates and clusters. *BMC Bioinformatics* **14**, 252 (2013).

59. N. Twomey, H. Chen, T. Diethe, P. Flach, An application of hierarchical Gaussian processes to the detection of anomalies in star light curves. *Neurocomputing* **342**, 152–163 (2019).

60. J. Salvatier, T. V. Wiecki, C. Fonnesbeck, Probabilistic programming in Python using PyMC3. *PeerJ Comput. Sci.* **2**, e55 (2016).

61. T. Donoghue, M. Haller, E. J. Peterson, P. Varma, P. Sebastian, R. Gao, T. Noto, A. H. Lara, J. D. Wallis, R. T. Knight, A. Shestyuk, B. Voytek, Parameterizing neural power spectra into periodic and aperiodic components. *Nat. Neurosci.* **23**, 1655–1665 (2020).

62. B. Lega, J. Burke, J. Jacobs, M. J. Kahana, Slow-theta-to-gamma phase–amplitude coupling in human hippocampus supports the formation of new episodic memories. *Cereb. Cortex* **26**, 268–278 (2016).

63. A. B. L. Tort, R. W. Komorowski, J. R. Manns, N. J. Kopell, H. Eichenbaum, Theta–gamma coupling increases during the learning of item–context associations. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 20942–20947 (2009).

64. A. B. L. Tort, R. Komorowski, H. Eichenbaum, N. Kopell, Measuring phase-amplitude coupling between neuronal oscillations of different frequencies. *J. Neurophysiol.* **104**, 1195–1210 (2010).