nature human behaviour

SUPPLEMENTARY INFORMATION

https://doi.org/10.1038/s41562-018-0324-5

In the format provided by the authors and unedited.

Multiple motor memories are learned to control different points on a tool

James B. Heald[1]*, James N. Ingram[1], J. Randall Flanagan[2] and Daniel M. Wolpert[1]

[1]Computational and Biological Learning Laboratory, Department of Engineering, University of Cambridge, Cambridge, UK. [2]Center for Neuroscience Studies and Department of Psychology, Queen's University, Kingston, ON, Canada. *e-mail: jbh40@cam.ac.uk

Multiple motor memories are learned to control different points on a tool

Supplementary Information

James B. Heald[1*], James N. Ingram[1], J. Randall Flanagan[2] & Daniel M. Wolpert[1]

[1] Computational and Biological Learning Lab, Department of Engineering,
University of Cambridge, Cambridge, CB2 1PZ, United Kingdom
[2] Center for Neuroscience Studies and Department of Psychology,
Queen's University, Kingston, ON, Canada

*Corresponding author: jbh40@cam.ac.uk

Supplementary Figures

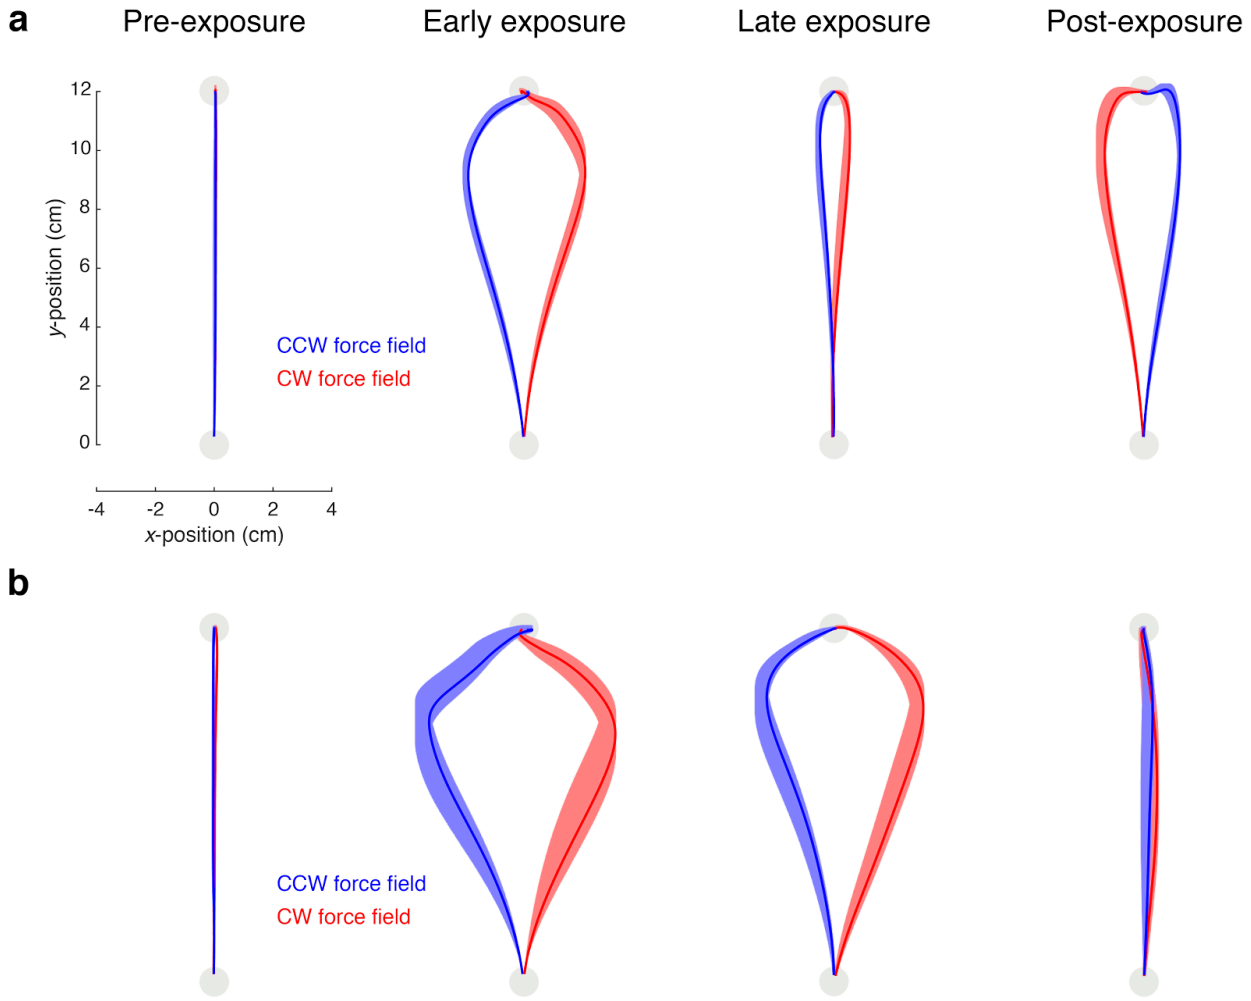

Supplementary Figure 1 | Hand paths at different stages of Experiment 1. a,b, In early exposure, the force fields produced substantial deviations of the hand path from a straight line. For the two-control-points groups (n = 10) **(a)**, hand paths were significantly straighter by late exposure, and clear after-effects were seen in early post-exposure. For the single-control-point group (n = 8) **(b)**, hand paths were not significantly straighter by late exposure, and no after-effects were seen. Traces show mean ± s.e.m. across participants hand paths over two blocks.

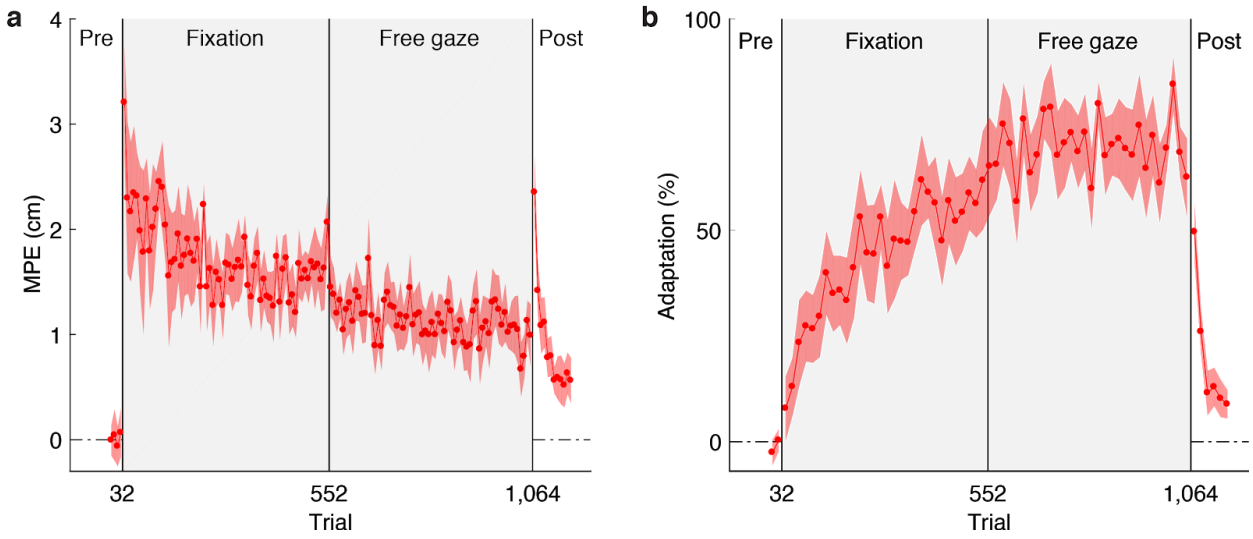

Supplementary Figure 2 | Adaptation is not gaze dependent. In the fixation stage of the exposure phase, participants (n = 8) were required to fixate a cross in the midline of the workspace. In the free-gaze stage of the exposure phase, no fixation cross was displayed and gaze was unconstrained. **a,b,** Mean ± s.e.m. **(a)** MPE and **(b)** adaptation across all participants. Each MPE data point represents the mean of a block of eight trials. To assess learning, we compared adaptation in the final two blocks of the pre-exposure phase with the final two blocks of the fixation stage. We found a significant increase in adaptation (two-tailed paired t-test, 60.2 ± 8.4%, t_7 = 7.70, P = 1 x 10^{-4}), which reached 59.0 ± 7.9% (mean ± s.e.m.) of full compensation by the final two blocks of the fixation stage. We contrasted the adaptation reached by blocks 23-24 of the exposure phase with the two-control-points and opposing-field groups who experienced the same task with free gaze. This revealed no significant difference (between-subjects ANOVA, $F_{2,23}$ = 2.07, P = 0.149), indicating that fixation did not detriment learning. We allowed the fixation group to continue exposure after the fixation phase but with free gaze. There was no significant additional learning after 616 further trials (two-tailed paired t-test, 6.4 ± 10.4%, t_7 = 0.67, P = 0.526).

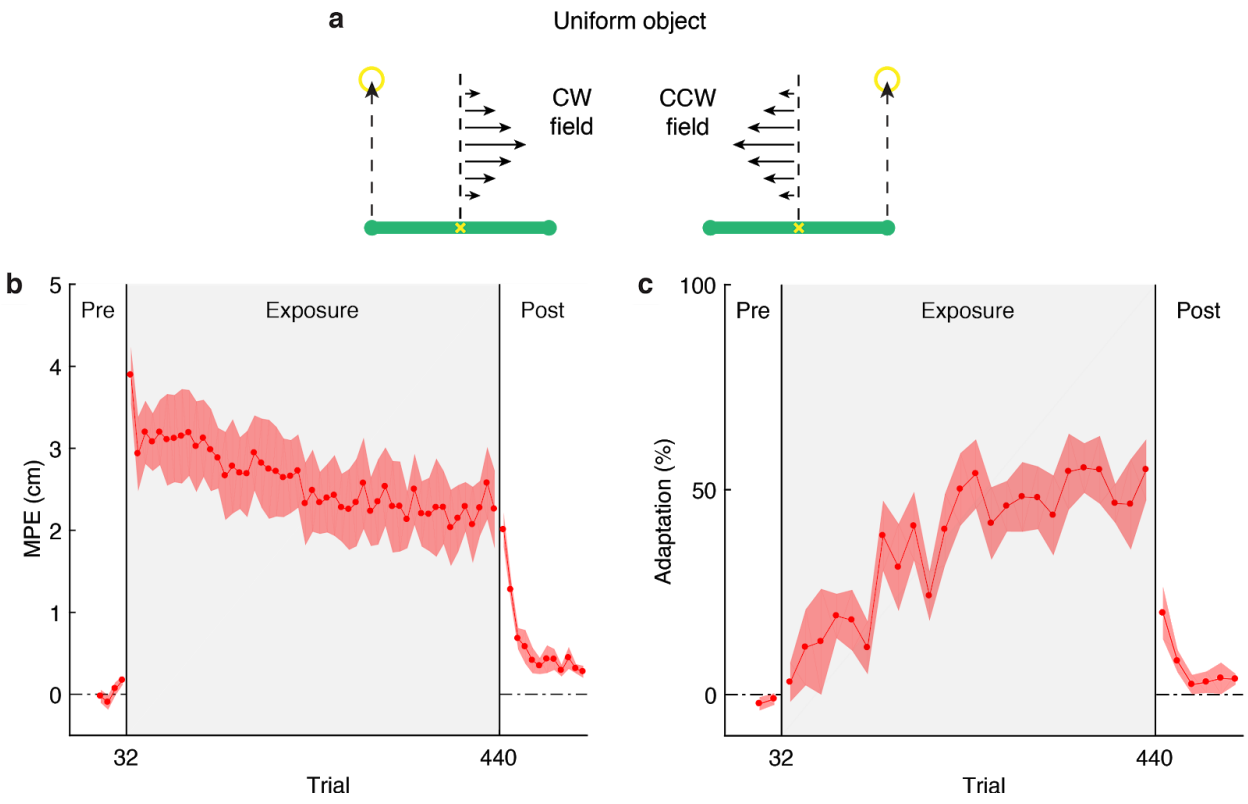

Supplementary Figure 3 | Adaptation does not require control points to be distinct from the rest of the object. **a,** Participants (n = 8) controlled a uniform object composed of two green disks connected by a green bar. Because the disks and bar were the same color, the object did not contain any distinct visual elements. **b,c,** Mean ± s.e.m. **(b)** MPE and **(c)** adaptation across all participants. Each MPE data point represents the mean of a block of eight trials. To assess learning, we compared adaptation in the final two blocks of the pre-exposure phase with the final two blocks of the exposure phase. We found a significant increase in adaptation (two-tailed paired t-test, 52.3 ± 9.3%, t_7 = 6.08, P = 5 x 10^{-4}), which reached 50.7 ± 8.5% (mean ± s.e.m.) of full compensation by the final two blocks of the exposure phase. We contrasted the adaptation reached by this stage with the two-control-point and opposing-field groups who experienced the same task with control points displayed as green disks on a red rectangle. There was no significant difference (between-subjects ANOVA, $F_{2,23}$ = 2.11, P = 0.144), indicating that the uniform object did not detriment learning.

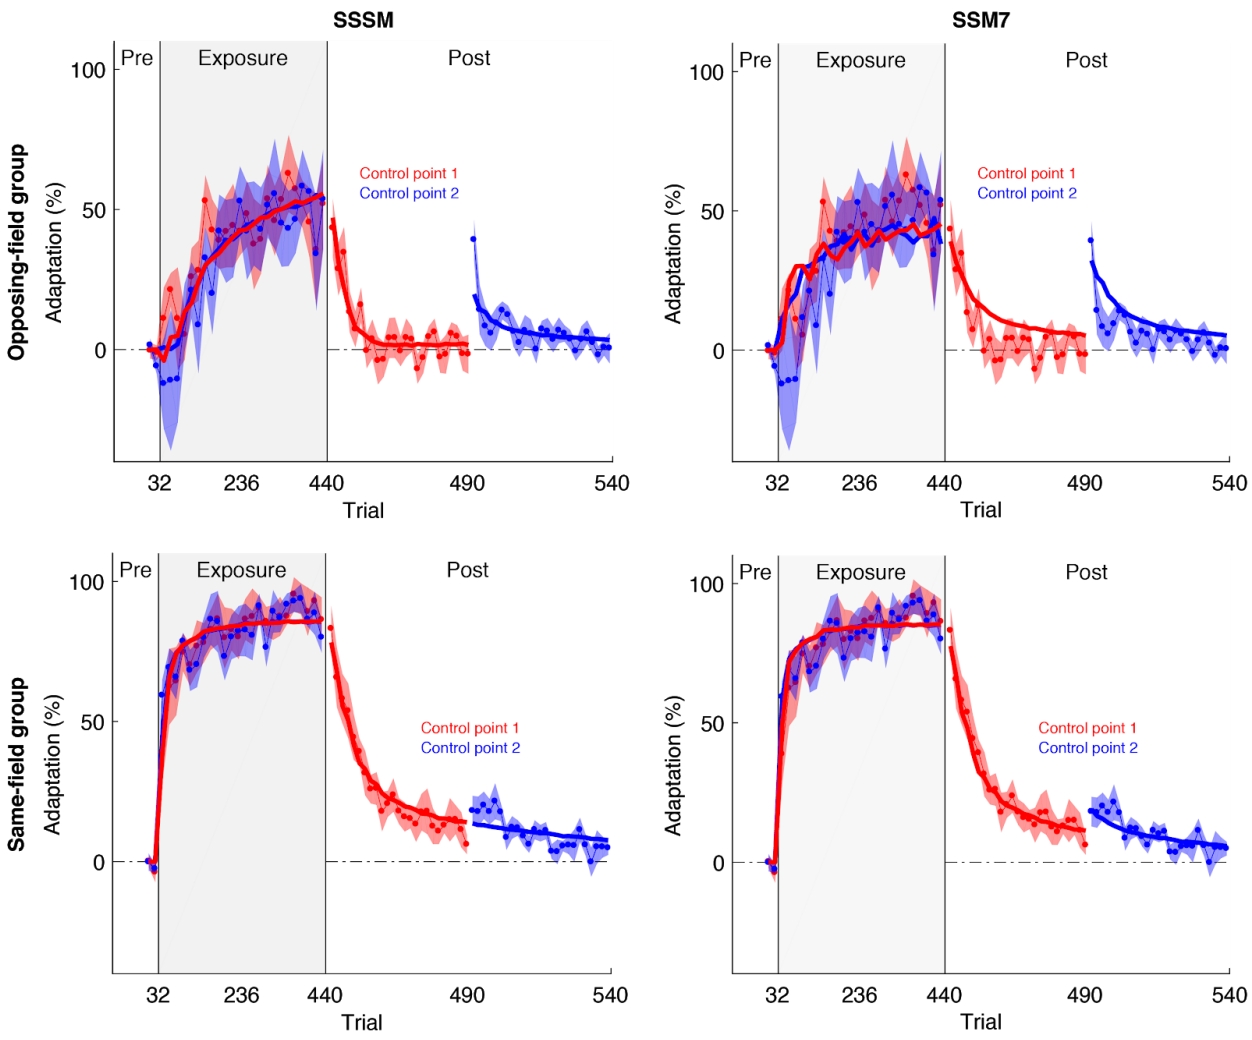

Supplementary Figure 4 | Model fits to adaptation data from Experiment 2. Solid lines show the mean fits across participants for the switching state-space model (SSSM) and state-space model 7 (SSM7), which had the next best BIC. Experimental data is shown as in Fig. 3b,d for the opposing-field (n = 8) and same-field (n = 8) groups. The superiority of the SSSM is most apparent in the fit to data from the opposing-field group.

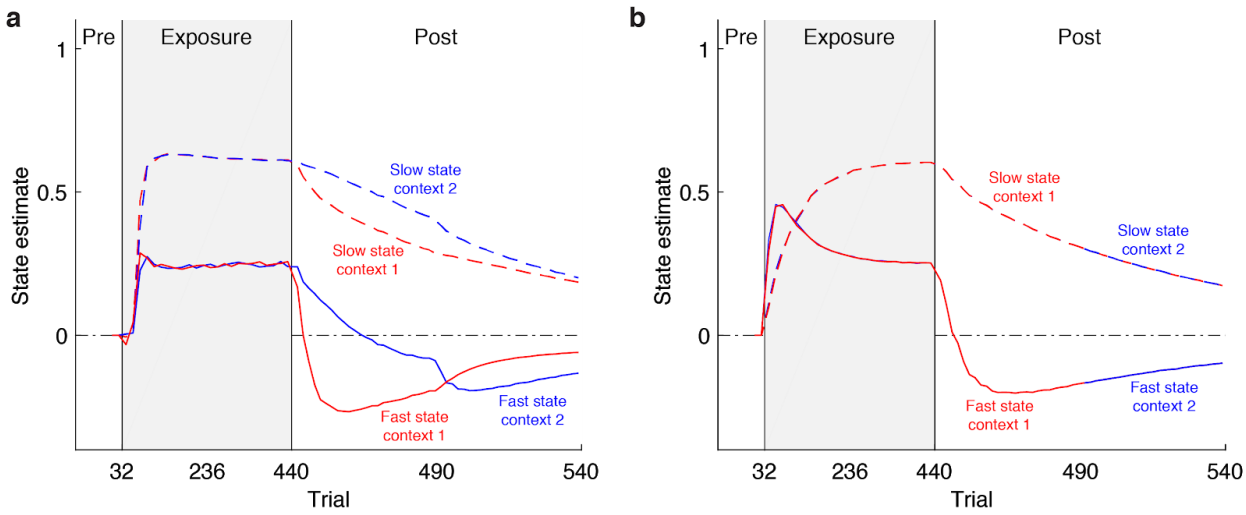

Supplementary Figure 5 | Slow and fast states of the switching state-space model. a, Mean state estimates across participants for the group that experienced opposing force fields at each control point. The slow and fast states of each context learn to represent different perturbations and, due to the learned association between contexts and cues, de-adapt at different rates during washout. For simplicity, the signs of the slow and fast states for context 2 have been inverted. **b,** For the group that experienced the same force field at each control point, the slow and fast states of each context represent the same perturbation and de-adapt at the same rate during washout.

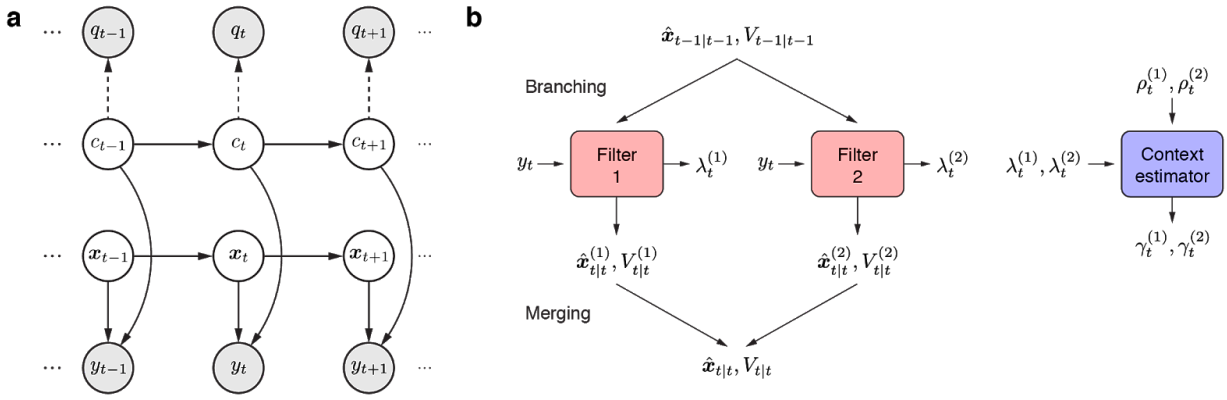

Supplementary Figure 6 | Motor adaptation as online state and parameter estimation in a switching state-space model. a, The graphical model. The discrete context, c_t, and the continuous perturbations, $\boldsymbol{x}_t$, evolve according to Markovian dynamics. The context determines both which observable cue, q_t, is emitted and how each perturbation contributes to the continuous observation y_t. Observed and latent variables are represented by gray and white nodes, respectively. The dashed arrows indicate dependencies that are learned online. **b,** The GPB1 algorithm. Given a single estimate of the state on trial $t-1$, each Kalman filter predicts the state on trial t and then updates its prediction when y_t is observed. The updated estimates of each filter are then merged into a single estimate. The probability of each context, which depends on both the contextual cue and the prediction error of each filter, determines how to merge the estimates in an optimal manner.

Supplementary Tables

Supplementary Table 1 | Variants of the context-dependent state-space model

Model	DOF	Retention vector ***a***	Learning-rate vector ***b***	Context vector ***c***	Number of effective states
SSM1[1]	3	$\begin{bmatrix} a_s & 0 & a_s & 0 \end{bmatrix}^{\mathrm{T}}$	$\begin{bmatrix} b_s & 0 & b_s & 0 \end{bmatrix}^{\mathrm{T}}$	$\boldsymbol{c}^{(1)} = \begin{bmatrix} 1 & 0 & c_s & 0 \end{bmatrix}^{\mathrm{T}}$ $\boldsymbol{c}^{(2)} = \begin{bmatrix} c_s & 0 & 1 & 0 \end{bmatrix}^{\mathrm{T}}$	2
SSM2[2]	3	$\begin{bmatrix} a & a & a & 0 \end{bmatrix}^{\mathrm{T}}$	$\begin{bmatrix} b_{\mathrm{s}} & b_{\mathrm{f}} & b_{\mathrm{s}} & 0 \end{bmatrix}^{\mathrm{T}}$	$\boldsymbol{c}^{(1)} = \begin{bmatrix} 1 & 1 & 0 & 0 \end{bmatrix}^{\mathrm{T}}$ $\boldsymbol{c}^{(2)} = \begin{bmatrix} 0 & 1 & 1 & 0 \end{bmatrix}^{\mathrm{T}}$	3
SSM3[1]	4	$\begin{bmatrix} a_{\mathrm{s}} & a_{\mathrm{f}} & a_{\mathrm{s}} & 0 \end{bmatrix}^{\mathrm{T}}$	$\begin{bmatrix} b_{\mathrm{s}} & b_{\mathrm{f}} & b_{\mathrm{s}} & 0 \end{bmatrix}^{\mathrm{T}}$	$\boldsymbol{c}^{(1)} = \begin{bmatrix} 1 & 1 & 0 & 0 \end{bmatrix}^{\mathrm{T}}$ $\boldsymbol{c}^{(2)} = \begin{bmatrix} 0 & 1 & 1 & 0 \end{bmatrix}^{\mathrm{T}}$	3
SSM4[1]	5	$\begin{bmatrix} a_{\mathrm{s}} & a_{\mathrm{f}} & a_{\mathrm{s}} & 0 \end{bmatrix}^{\mathrm{T}}$	$\begin{bmatrix} b_{\mathrm{s}} & b_{\mathrm{f}} & b_{\mathrm{s}} & 0 \end{bmatrix}^{\mathrm{T}}$	$\boldsymbol{c}^{(1)} = \begin{bmatrix} 1 & 1 & c_{\mathrm{s}} & 0 \end{bmatrix}^{\mathrm{T}}$ $\boldsymbol{c}^{(2)} = \begin{bmatrix} c_{\mathrm{s}} & 1 & 1 & 0 \end{bmatrix}^{\mathrm{T}}$	3
SSM5[1]	5	$\begin{bmatrix} a_{\mathrm{s}} & a_{\mathrm{f}} & a_{\mathrm{s}} & a_{\mathrm{f}} \end{bmatrix}^{\mathrm{T}}$	$\begin{bmatrix} b_{\mathrm{s}} & b_{\mathrm{f}} & b_{\mathrm{s}} & b_{\mathrm{f}} \end{bmatrix}^{\mathrm{T}}$	$\boldsymbol{c}^{(1)} = \begin{bmatrix} 1 & 1 & 0 & c_{\mathrm{f}} \end{bmatrix}^{\mathrm{T}}$ $\boldsymbol{c}^{(2)} = \begin{bmatrix} 0 & c_{\mathrm{f}} & 1 & 1 \end{bmatrix}^{\mathrm{T}}$	4
SSM6[1]	5	$\begin{bmatrix} a_{\mathrm{s}} & a_{\mathrm{f}} & a_{\mathrm{s}} & a_{\mathrm{f}} \end{bmatrix}^{\mathrm{T}}$	$\begin{bmatrix} b_{\mathrm{s}} & b_{\mathrm{f}} & b_{\mathrm{s}} & b_{\mathrm{f}} \end{bmatrix}^{\mathrm{T}}$	$\boldsymbol{c}^{(1)} = \begin{bmatrix} 1 & 1 & c_{\mathrm{s}} & 0 \end{bmatrix}^{\mathrm{T}}$ $\boldsymbol{c}^{(2)} = \begin{bmatrix} c_{\mathrm{s}} & 0 & 1 & 1 \end{bmatrix}^{\mathrm{T}}$	4
SSM7[1]	6	$\begin{bmatrix} a_{\mathrm{s}} & a_{\mathrm{f}} & a_{\mathrm{s}} & a_{\mathrm{f}} \end{bmatrix}^{\mathrm{T}}$	$\begin{bmatrix} b_{\mathrm{s}} & b_{\mathrm{f}} & b_{\mathrm{s}} & b_{\mathrm{f}} \end{bmatrix}^{\mathrm{T}}$	$\boldsymbol{c}^{(1)} = \begin{bmatrix} 1 & 1 & c_{\mathrm{s}} & c_{\mathrm{f}} \end{bmatrix}^{\mathrm{T}}$ $\boldsymbol{c}^{(2)} = \begin{bmatrix} c_{\mathrm{s}} & c_{\mathrm{f}} & 1 & 1 \end{bmatrix}^{\mathrm{T}}$	4

We examined 7 context-dependent state-space models, which varied in their number of parameters (DOF). The models could vary in the retention parameters, learning rates and coupling. The setting of the parameters determine the effective number of states in each model. Note that models 2-4 do not map the states on to contexts 1 and 2 in equation (4) as one of the states is shared by both contexts (i.e., context independent). Furthermore, model 2 does not necessarily map the states on to the fast and slow states in equation (4) as the context-independent state can be either fast or slow.

Supplementary Table 2 | Comparison of model fits for the state-space (SSM) and switching state-space (SSSM) models

State-space models	R^2	ΔBIC	DOF	a_s	b_s	c_s	a_f	b_f	c_f
SSM1	0.92	135.3	3	0.9858	0.0730	0.4578	-	-	-
SSM2	0.92	135.3	3	0.9858	0.0215	[0]	0.9858	0.0668	[1]
SSM3	0.92	140.6	4	0.9858	0.0215	[0]	0.9858	0.0668	[1]
SSM4	0.93	117.3	5	0.9905	0.0445	0.4040	0.9229	0.0922	[1]
SSM5	0.92	145.9	5	0.9862	0.0007	[0]	0.9858	0.0721	0.4636
SSM6	0.93	129.6	5	0.9895	0.0502	0.5234	0.8860	0.0553	[0]
SSM7	0.94	112.4	6	0.9932	0.0295	0.4660	0.9405	0.0888	0.4340

Switching state-space model				a_s	q_s	a_f	q_f	r	α
SSSM	0.96	0	6	0.9946	0.0003	0.9404	0.0068	0.0792	0.6067

Numbers in square brackets are fixed and not fit to the data. Note that although the R^2 values are similar these are independent of the number of data points, whereas the BIC scales with the number of data points. Therefore, the small improvement in R^2 for a model can translate into a large difference in likelihood and hence BIC. In SSM2, the fast state corresponds to the context-independent state shared by both contexts (see Supplementary Table 1 for details).

Supplementary References

1. Kim, S., Oh, Y. & Schweighofer, N. PLoS One 10, e0142963 (2015).
2. Nozaki, D. & Scott, S.H. Exp. Brain Res. 194, 451–463 (2009).