Transcranial Ultrasound Stimulation in Anterior Cingulate Cortex Impairs Information Sampling and Learning in Loss Contexts

Kianoush Banaie Boroujeni[1*], Michelle K Sigona[3,4], Robert Louie Treuting[4], Thomas J. Manuel[3,4], Charles F. Caskey[2,3,4], Thilo Womelsdorf[1,4,*]

[1]*Department of Psychology, Vanderbilt University, Nashville, TN 37240.*
[2]*Vanderbilt University Medical Center Department of Radiology and Radiological Sciences, Nashville, TN, USA.*
[3]*Vanderbilt University Institute of Imaging Science, Nashville, TN, USA.*
[4]*Department of Biomedical Engineering, Vanderbilt University, Nashville, TN 37240.*

Corresponding Authors: Kianoush Banaie Boroujeni
(kianoush.banaie.boroujeni@vanderbilt.com) and Thilo Womelsdorf
(thilo.womelsdorf@vanderbilt.edu)

Keywords: Striatum, focused ultrasound, cognitive flexibility, nonhuman primate, loss aversion, motivation, effort control,

Neural activity in anterior cingulate cortex and the anterior striatum predicts which visual objects are sampled and how likely objects are paired with positive or aversive outcomes. We causally tested whether these neural signals contribute to behavioral flexibility. Disrupting with transcranial ultrasound the ACC, but not striatum, prolonged information sampling when attentional demands were high, impaired flexible learning, and reduced the ability to avoid losses. These results support a role of the ACC in guiding attention and information sampling to overcome motivational conflict during adaptive behaviors.

Anterior cingulate cortex (ACC) and the anterior striatum (aSTR) are involved in flexible learning behavior, but it is unclear how they causally contribute to behavioral flexibility. One hallmark of flexibility is the ability of subjects to allocate attention efficiently to those visual objects that reduce uncertainty and enhance the likelihood to receive reward. This function might be supported by neural circuits that rapidly learn which stimuli have positive or negative expected values and control directing gaze and attention to those stimuli[1,2]. These neural circuits have been found in

ACC[3–7] and in the striatum[8,9], suggesting that an overarching functional contribution of the ACC - striatum network to behavioral flexibility could be the guidance of information sampling (or attention) to optimize behavioral outcomes[1].

This viewpoint, however, needs to be reconciled with evidence for a motivational role of the same neural circuits suggesting they control how much effort is exerted to overcome task challenges[10]. ACC and aSTR activate when subjects face enhanced task demands and exert effort to succeed with a task[11]. Such effort-control might be supported by neural populations that calculate the benefits of exerting more effort in a task, as well as the costs that would be incurred for not putting effort into a task[12]. Both, ACC and aSTR contain neurons signaling these motivational variables in separate neural populations encoding prospective gains and the anticipated losses and aversive consequences associated with specific visual stimuli[6,13].

To understand the relative importance of the ACC and aSTR for guiding information sampling and controlling motivational effort we adopted a causal transcranial ultrasound approach to disrupt local neural processing in ACC or aSTR while monkeys performed a learning task that independently varied (*i*) the attentional demands for information sampling from (*ii*) motivational demands modulated by varying prospective gains and losses. The task required subjects to learn through trial-and-error which unique object feature predicted reward by choosing in individual trials one of three multidimensional visual objects (**Fig. 1a**,**b**). The rewarded feature remained constant for blocks of ~55 trials and then changed unannounced. Learning blocks had either low attentional demands when the objects varied in only one dimension (e.g. different shapes), or high attentional demands when objects varied in two or three dimensions (e.g. varying shapes, surface

patterns, and arm types) (**Fig. 1c**). In addition, motivational effort varied by defining learning blocks either as pure *gain contexts,* or as a mixed *gain-loss contexts*. In the gain-only context, subjects received three tokens for correct choices that they later cashed out for fluid reward, while in the mixed gain/loss-context subjects received two tokens for correct choices and lost already attained token assets when choosing objects with non-rewarded features (**Fig. 1a-c**). The loss of tokens triggers a vigilance response[14] and imposes a motivational conflict[15]. We used transcranial focused ultrasound (TUS) to disrupt bilaterally ACC or anterior striatum in separate learning sessions adopting the same TUS protocol as in[16,17]. The protocol imposes a ~6mm wide / 40mm tall sonication region that has been shown previously to alter behavior in foraging tasks[18], to reduce functional connectivity of the sonicated area in macaques[16], and in in-vitro preparations to modulate neuronal excitability to external inputs[19]. We provide detailed acoustic simulation of the ultrasound pressure dispersion around the target, the anatomical sub-millimeter targeting precision of TUS and the real-time monitored ultrasound power in **Extended Data Fig. 3**. We bilaterally sonicated or sham-sonicated the ACC or the aSTR in individual sessions when monkeys had completed the first six learning blocks. Following the sonication procedure, monkeys resumed the task and proceeded on average for 23.6 (±4 SE) learning blocks (monkey W: 20.5±4, monkey I: 26.5±4) (**Fig. 1d**).

Across learning blocks, monkeys reached the learning criterion (≥80% correct choices over 12 trials) systematically later in blocks with high attentional load (LMEs, p<0.001; **Fig. 1e**, **Extended Data Fig. 1a-c**). Both monkeys also showed longer foveation durations onto the objects prior to making a choice when the attentional load was high (LMEs, p<0.001; **Fig. 1f**; **Extended Data Fig. 1d-e**). These longer foveation durations index more information sampling at higher attentional

load. Monkeys also learned slower and showed reduced plateau performance (**Extended Data Fig. 2a-c**) as well as increased explorative sampling in blocks with gains and losses (gain/loss contexts) compared to blocks with only gains (gain-only contexts) (**Fig. 1f**, **Extended Data Fig. 2d-e**; LMEs, p<0.001).

TUS in ACC but not sham-TUS in ACC or TUS or sham-TUS in the striatum (**Fig. 2a**) changed this behavioral pattern. ACC-TUS selectively slowed learning in the gain/loss contexts compared to the gain-only contexts (LME's, t=2.67, p=0.007)) (**Fig. 2b**; **Extended Data Fig. 4a,b, Supplementary Table 1**). ACC-TUS increased the number of trials needed to reach 80% learning criterion to 14.7±0.8 (monkey W/I: 14.3±1.2 / 15±1.2) relative to pre-TUS baseline (trials to criterion: 10.7±1.4; monkey W/I: 9.9±2.3 / 11.4±1.8) (Wilcoxon test, p=0.049; **Fig. 2b**; **Extended Data Fig. 4c,d**). The slower learning with ACC-TUS in the gain/loss context was significantly different to other TUS conditions (Kruskal-Wallis test, p=0.003), ACC-sham (pairwise Wilcoxon test, FDR multiple comparison corrected for dependent samples, p=0.019), and to the conditions with TUS-striatum and TUS-sham (pairwise Wilcoxon test, FDR multiple comparison corrected for dependent samples, p=0.019, p=0.003) (**Fig. 2b**). This effect interacted with attentional load. The slower learning after ACC-TUS in the gain/loss condition was stronger when attentional load was intermediate (2 distracting features) and high (3 distracting features) (random permutation, p<0.05; LME's, t=-2.8, p=0.004) (**Fig. 3c,d**; **Extended Data Fig. 5a, Supplementary Table 2**). TUS did not affect learning in the gain-only contexts even at high attentional load (Kruskal-Wallis test, p=0.933) (**Fig. 3c,d**; **Extended Data Fig. 5b**).

The slower learning in the gain/loss context after ACC-TUS was accompanied by prolonged explorations of objects prior to making a choice compared to pre-TUS baseline (Wilcoxon test, p=0.016), and compared to aSTR-TUS (Kruskal-Wallis test, p=0.036; pairwise Wilcoxon test, FDR multiple comparison corrected for dependent samples, p=0.03) (**Fig. 3a, Extended Data Fig. 6a,c**). Explorative information sampling in the gain-only context did not vary between TUS conditions (Kruskal-Wallis test, p=0.55) (**Fig. 3b, Extended Data Fig. 6b,d**). Once subjects reached the learning criterion in a learning context, they could exploit the learned feature rule until the learning contexts changed after ~30-55 trials. During this period, they showed high (>85%) plateau performance which was significantly lower in the gain/loss (91%±0.07) than the gain-only contexts (92%±0.05; LME's, t-stat= -3.95, p=0.001) (**Fig. 2c**). ACC-TUS exacerbated this performance drop in the gain/loss block, leading to significantly lower plateau accuracy than in the pre-TUS baseline and other TUS conditions in gain-loss learning context (LME's, t-stat= -2.05, p=0.04; **Fig. 3c**; **Extended Data Fig. 6e**), but not in the gain-only learning context (Kruskal-Wallis test, p=0.8) (**Fig. 3d**; **Extended Data Fig. 6f**).

We validated that the behavioral impairments emerged shortly after the sonication and lasted until the end of the ≤120 min. long session (**Extended Data Fig. 7**). We also confirmed that the observed behavioral effects were not only evident at the block level (**Fig. 1-2**), but were evident also when averaging the learning performance per session and applying session-level statistics (**Extended Data Fig. 8**, **Supplementary Table 1**).

So far, we found that ACC-TUS slowed learning speed, increased exploration durations, and reduced plateau accuracy in the loss context. These results could be due to motivational

difficulties when adjusting to the loss of an already attained token. A first analysis showed that loss experience after a single choice leads to overall poorer performance in subsequent trials when TUS was directed to ACC but not after sham or aSTR-TUS (random permutation test, p<0.05, **Fig. 3e, Extended Data Fig. 9a-d**). This overall performance decrement was dependent on the recent history of losses. ACC-TUS reduced performance specifically when subjects had lost 3 or 2 tokens in the preceding four trials, but not when their net token gain in the past four trials was ≥0 tokens (random permutation test, p<0.05, **Fig. 3f**). This dependence of the ACC-TUS effect on the recent gross token income was evident in both monkeys (**Extended Data Fig. 9e**) and could be a main source that led to slower learning.

In summary, with sonicated ACC, learning was slowed down, information sampling was longer, and performance was poorer when the task required overcoming motivational challenges in the gains-loss context in conjunction with the cognitive demand to overcome high uncertainty in the high attentional load condition. The slower learning reflects suboptimal behavioral adjustment that reduced the overall reward intake when the ACC was disrupted, which is consistent with reports that more widespread lesions in ACC causes perseveration[20], reduced inhibition of prevalent motor programs[21], and disrupts integrating reward history[22].

Our findings document the ACC is causally supporting efficient information sampling when the feature sampling space is large (at high attentional load conditions), and when the benefits over cost ratio for performance outcomes is low (in the gains/loss contexts). This conclusion supports views that an overarching functional contribution of the ACC is to support information sampling for optimizing behavioral outcomes during adaptive behavior[2]. The results suggest, however, that

this guidance is not limited to situations with high cognitive uncertainty, but include situations with reduced payoffs in return for effort control. This conclusion resonates with reduced effort control and attentional control in patients with depression[23] and schizophrenia[24] and suggests that enhancing ACC functioning could enhance information sampling and effort control and thereby improve behavioral flexibility.

Acknowledgments

This work was supported by the National Institute of Mental Health of the National Institutes of Health under Award Number R01MH123687 and 1UF1NS107666. The authors acknowledge Huiwen Luo for assistance with calibrating optically tracked devices using MRI and Adrienne Hawkes for technical assistance with voltage monitoring. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Author Contributions

K.B.B., C.F.C. and T.W. conceived the experiments. K.B.B. and R.L.T. performed the experiment. M.K.S. and T.M. contributed the image-guided procedures and computer simulations. K.B.B. reconstructed the brain images. K.B.B. and M.K.S. wrote the code to run the experiments. K.B.B. analyzed and visualized the data. All authors contributed writing the paper.

Competing Interests

The authors declare no competing interests.

Data and code availability

All data supporting this study and its findings, as well as custom MATLAB code generated for analyses, are available from the corresponding author upon reasonable request.

Figures and Figure Legends

Fig. 1. Task paradigm and TUS protocol. a, A trial started with central gaze fixation and appearance of 3 objects. Monkeys could then explore objects and choose one object by fixating it for 700 ms. A correct choice triggered visual feedback (a yellow halo of the chosen object) and the appearance of green circles (tokens for reward) above the chosen object. The tokens were then animated and traveled to a token bar on the top of the screen. Incorrect choices triggered a blue halo and in the gain-loss condition one blue token was shown that traveled to the token bar where one already attained token was removed. When ≥5 tokens were collected in the token bar, fluid reward was delivered, and the token bar reset to zero. **b**, In successive learning blocks, different visual features were associated with reward and blocks alternated randomly the gain-only (G) and gain-loss (GL) conditions. **c**, Attentional load varied by increasing the number of object features that varied from block to block from 1, 2 and 3. **d**, In each sonication or sham session the experiment was paused after 6 learning blocks. There are four experimental conditions; high energy TUS in ACC (H-ACC; red), or anterior striatum (H-aSTR; green), or sham ACC (S-ACC; dimmed red), or sham anterior striatum (S-aSTR; dimmed green). 30-ms bursts of transcranial ultrasound stimulation were delivered every 100-ms over a duration of 80-second (40-second each hemisphere). **e**, Proportion of correct choices over trials since block begin for different attentional loads (left panel; 1-3D, light to dark grey) and learning contexts (gain-only, blue; gain-loss, red). **f**, The average fixation duration of objects during explorative sampling prior to choosing an object in the same format as in (**e**). The lines show the mean and the shaded error bars are standard error of the mean.

Fig. 2. Transcranial ultrasound stimulation (TUS) of ACC and anterior striatum. a. The maximum negative peak pressure of TUS in an example session of one monkey shows the focus is within ACC (left) and anterior striatum (right). **b.** Learning is slowed down (trials-to-criterion increased) after TUS in ACC in the gain-loss context (left, LME's, p=0.007) but not in the gain-only contexts (right, n.s.). Data represent means and the standard error of the mean. **c.** Learning is slowed with higher attentional load with a significant interaction in the gain-loss context (left, LME's, p=0.007) but not in the gain-only context (right, n.s., for the full multiple comparison corrected statistical results see **Supplementary Table 1,2**). **d.** Marginally normalized trials-to-criterion is significantly higher with ACC-TUS in the gain-loss (GL) learning context at higher (2D and 3D) attentional load (random permutation p<0.05). Each cell is color coded with the mean value ± standard error of the mean with a low to high value gradient from left to right. The white rectangle shows that learning in that TUS condition (*x-axis*) is different from other TUS conditions. White asterisks indicate that an attentional load condition in a learning context in a TUS condition is significantly different from other TUS conditions.

Fig. 3. Transcranial ultrasound stimulation effect on explorative behavior and post-error adjustment. a,b. Explorative sampling (the duration of fixating objects prior to choosing an object) is increased with TUS in ACC in gain-loss (**a**) but not gain-only (**b**) learning context**. c,d**. TUS in ACC reduced post-learning plateau accuracy in gain-loss contexts (**c**) but not in gain-only contexts (**d**) (LMEs, p=0.04) (see **Supplementary Table 2**). With FDR correction, this accuracy effect was significant at the session level, but not at the block level, indicating a low effect size (**Extended Data Fig. 8e** and **Supplementary Table 1**). **e.** The accuracy was overall reduced with TUS in ACC in the five trials after experiencing a token loss in the gain-loss context (random permutation, p<0.05). **f.** Gross token income (GTI) (*x-axis*) measures the signed average of tokens gained and lost in the preceding trials. TUS in ACC reduced the performance accuracy (*y-axis*) when monkeys had lost more tokens in the near past (negative GTI) (random permutation, p<0.05). Data indicate the mean and the standard error of the mean. Accuracy on the trial-level is normalized by the mean and standard deviation of the accuracy in the baseline (the first six blocks prior to the TUS) in the same TUS sessions.

References

1. Monosov, I. E. & Rushworth, M. F. S. Interactions between ventrolateral prefrontal and anterior cingulate cortex during learning and behavioural change. *Neuropsychopharmacol.* 1–15 (2021) doi:10.1038/s41386-021-01079-2.

2. Hunt, L. T. *et al.* Formalizing planning and information search in naturalistic decision-making. *Nat Neurosci* 1–14 (2021) doi:10.1038/s41593-021-00866-w.

3. Kaping, D., Vinck, M., Hutchison, R. M., Everling, S. & Womelsdorf, T. Specific Contributions of Ventromedial, Anterior Cingulate, and Lateral Prefrontal Cortex for Attentional Selection and Stimulus Valuation. *PLOS Biology* **9**, e1001224 (2011).

4. Oemisch, M. *et al.* Feature-specific prediction errors and surprise across macaque fronto-striatal circuits. *Nature Communications* **10**, 1–15 (2019).

5. White, J. K. *et al.* A neural network for information seeking. *Nat Commun* **10**, 5168 (2019).

6. Jezzini, A., Bromberg-Martin, E. S., Trambaiolli, L. R., Haber, S. N. & Monosov, I. E. A prefrontal network integrates preferences for advance information about uncertain rewards and punishments. *Neuron* **109**, 2339-2352.e5 (2021).

7. Banaie Boroujeni, K., Tiesinga, P. & Womelsdorf, T. Interneuron specific gamma synchronization indexes cue uncertainty and prediction errors in lateral prefrontal and anterior cingulate cortex. *eLife* **10**, e69111 (2021).

8. Boroujeni, K. B., Oemisch, M., Hassani, S. A. & Womelsdorf, T. Fast spiking interneuron activity in primate striatum tracks learning of attention cues. *PNAS* **117**, 18049–18058 (2020).

9. Arcizet, F. & Krauzlis, R. J. Covert spatial selection in primate basal ganglia. *PLOS Biology* **16**, e2005930 (2018).

10. Shenhav, A., Cohen, J. D. & Botvinick, M. M. Dorsal anterior cingulate cortex and the value of control. *Nat Neurosci* **19**, 1286–1291 (2016).

11. Shenhav, A., Straccia, M. A., Cohen, J. D. & Botvinick, M. M. Anterior cingulate engagement in a foraging context reflects choice difficulty, not foraging value. *Nature Neuroscience* **17**, 1249–1254 (2014).

12. Shenhav, A., Botvinick, M. M. & Cohen, J. D. The Expected Value of Control: An Integrative Theory of Anterior Cingulate Cortex Function. *Neuron* **79**, 217–240 (2013).

13. Monosov, I. E. Anterior cingulate is a source of valence-specific information about value and uncertainty. *Nature Communications* **8**, 134 (2017).

14. Yechiam, E. & Hochman, G. Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin* **139**, 497–518 (2013).

15. Boroujeni, K. B., Watson, M. & Womelsdorf, T. Gains and Losses Differentially Regulate Attentional Efficacy at Low and High Attentional Load. *bioRxiv* 2020.09.01.278168 (2020) doi:10.1101/2020.09.01.278168.

16. Fouragnan, E. F. *et al.* The macaque anterior cingulate cortex translates counterfactual choice value into actual behavioral change. *Nature Neuroscience* **22**, 797–808 (2019).

17. Verhagen, L. *et al.* Offline impact of transcranial focused ultrasound on cortical activation in primates. *eLife* **8**, e40541 (2019).

18. Bongioanni, A. *et al.* Activation and disruption of a neural mechanism for novel choice in monkeys. *Nature* **591**, 270–274 (2021).

19. Clennell, B. *et al.* Transient ultrasound stimulation has lasting effects on neuronal excitability. *Brain Stimulation* **14**, 217–225 (2021).

20. Amiez, C., Joseph, J. P. & Procyk, E. Reward Encoding in the Monkey Anterior Cingulate Cortex. *Cerebral Cortex* **16**, 1040–1055 (2006).

21. Shima, K. & Tanji, J. Role for Cingulate Motor Area Cells in Voluntary Movement Selection Based on Reward. *Science* **282**, 1335–1338 (1998).

22. Kennerley, S. W., Walton, M. E., Behrens, T. E. J., Buckley, M. J. & Rushworth, M. F. S. Optimal decision making and the anterior cingulate cortex. *Nature Neuroscience* **9**, 940–947 (2006).

23. Berwian, I. M. *et al.* Computational Mechanisms of Effort and Reward Decisions in Patients With Depression and Their Association With Relapse After Antidepressant Discontinuation. *JAMA Psychiatry* **77**, 513 (2020).

24. Gold, J. M. *et al.* Negative Symptoms of Schizophrenia Are Associated with Abnormal Effort-Cost Computations. *Biological Psychiatry* **74**, 130–136 (2013).

Extended Data Figures and Figure Legends

Extended Data Fig. 1. Attentional load effect on learning and fixational information sampling. a,b. Both monkeys reached the learning criterion of 80% or more correct trials (based on a 12 trials forward-looking window). Learning is fastest at low attentional load (*light grey*), and slowest at high attentional load (*dark grey*). In all panels, the left column shows the results for monkey W, the middle for monkey I, and the right for both monkeys combined. **c.** Post-learning accuracy is significantly reduced in higher attentional load (LME's, P<0.001). **d,e.** Explorative sampling is the duration in msec. of fixational sampling of objects prior to making a choice. Explorative sampling increased at the beginning of a block, reached a maximum during learning, and but remained elevated only at the highest attentional load (LMEs, p<0.001). **f.** Choice reaction time is the time from stimulus onset to the onset of the final fixation (that chooses the object). It increased with attentional load (Kruskal-Wallis test, p<0.001). **g.** Choice sampling is measured as the duration of sampling the chosen object prior to the final choice fixation. Choice sampling is increased in blocks with higher attentional load (Kruskal-Wallis test, p<0.001). **h.** Asset sampling duration quantifies how long subjects fixate the token bar prior to choosing an object. Asset sampling is independent of attentional load, and more extensive in monkey I. Data show means and standard error of the mean.

Extended Data Fig. 2. Effects of Gain-Loss and Gain-Only learning contexts on learning and gaze sampling behaviors. Same format as **Extended Fig. 1. a,b.** Learning is faster in the gain-only context (blue), than in the gain-loss context. In all panels, the left column shows the results for monkey W, the middle for the monkey I, and the right for both monkeys. **c.** Post-learning accuracy is significantly reduced in the gain-loss context (Wilcoxon test, P<0.001). **d,e.** Explorative sampling reaches a higher maximum during learning in the gain-loss context (red) compared to the gain-only context (blue) (Wilcoxon test, P<0.001). **f.** Choice reaction time is slower in the gain-loss context (Wilcoxon test, p<0.001). **g.** Sampling of the object that is subsequently chosen is longer in the gain-loss context (Wilcoxon test, p<0.001). **h.** Asset sampling duration is longer in the gain-loss context (Wilcoxon test, p<0.001). Data show mean and standard error of the mean.

Extended Data Fig. 3 Transcranial ultrasound stimulation localization, energy, and sonication focus specifications. In each experimental session we positioned the transducer sonication beam to focus on left and right hemisphere ACC and striatum and sonicated the area for 40 seconds in each hemisphere. **a.** Both hemispheres (ACC in red, and anterior striatum in green) in both monkeys were targeted precisely within an averaged sub-millimeter distance of the center of focal beam of the transducer to the anatomical target region. **b.** Real-time monitoring of output voltage to the transducer confirmed reliable feedforward voltage range for both monkeys (*see* Methods). **c.** Computer simulations show a reliable range of maximum peak pressure at the focal center of the transducer. **d,e.** Root mean squared (RMS) deviation values of negative peak pressure at the focus of the sonication attenuated at -3dB (**d**), and -6dB (**e**).

Extended Data Fig. 4. Transcranial ultrasound stimulation effect on learning for individual monkeys (left and middle column) and their average (right column). a,b. Learning curves for four different experimental conditions: high energy TUS in ACC (H-ACC; red), or anterior striatum (H-aSTR; green), or sham ACC (S-ACC; darkened red), or sham anterior striatum (S-aSTR; darkened green). Learning curves are shallower after ACC-TUS in the gain-loss context (**a**) but not in the gain-only context (**b**). **c,d.** The reduced learning speed (increased trials-to-criterion) with TUS-ACC in the gain-loss learning context (**c**) but not in the gain-only learning context (**d**). Detailed statistics are provided in **Supplementary Table 1**. Data show mean and standard error of the mean.

Extended Data Fig. 5. TUS interaction of attentional load and motivational context. a,b. TUS in ACC reduces learning speed at higher attentional load in the gain-loss learning context (**a**), but not the gain-only learning context (**b**) (for the full multiple comparison corrected statistical results see **Supplementary Table 2**) **c.** Marginally normalized trials-to-criterion are significantly higher with ACC-TUS in the gain-loss learning context, and the effect in ACC-TUS was only significant at higher attentional load conditions 2D and 3D (random permutation $p<0.05$). Error bars are standard error of the mean. The white rectangle in (**c**) shows learning in a TUS condition is different from other TUS conditions. Each cell is color coded with the mean value ± standard error of the mean with a low to high value gradient from left to right. The white asterisk shows attentional loads in a learning context in a TUS condition is significantly different from other TUS conditions. In all panels, the left column shows the results for monkey W, the middle for the monkey I, and the right for both monkeys combined.

Extended Data Fig. 6. Transcranial ultrasound stimulation effects on explorative behavior. a,b. Explorative sampling curves showing the average fixation durations on objects before making a choice. With the beginning of a new learning block explorative sampling increases to a maximum during learning and then slowly reduces to a baseline level. TUS in ACC in the gain-loss learning context (**a**), but not in the gain-only context causes elevated exploratory sampling (**b**). **c,d.** Explorative sampling is increased with TUS in ACC relative to baseline and other TUS conditions in the gain-loss context (**c**) but not the gain-only context (**d**). **e,f.** Post-learning accuracy is not significantly different from the baseline and other TUS conditions in the block-level analysis, neither in gain-loss context (**e**), nor the gain-only context (**f**) (detailed statistics in **Supplementary Table 1**).

Extended Data Fig. 7. Time course of the effect of transcranial ultrasound stimulation on learning and explorative behavior. a. Monkeys show a gradual deterioration of learning with TUS in ACC over time in the gain-loss conditions. **b**. Explorative sampling increases over time relative to the time of TUS in ACC. The x-axis shows the time in minutes relative to TUS stimulation within a session. The time-step resolution is 20 minutes. Both monkeys show a similar time course with the slowing of learning from 20-40 minutes after the stimulation to the end of the session (randomized permutation test, $p < 0.05$). Shaded error bars are standard errors of the mean.

Extended Data Fig. 8. Session level effect of transcranial ultrasound stimulation on learning and explorative sampling. Average results for each session and across sessions for monkey W and I, and both combined (left, middle, and right column, respectively). **a,b.** TUS in ACC caused slower learning in both monkeys in the gain-loss context (**a**), but not in the gain-only context (**b**). **c,d.** TUS in ACC causes longer explorative sampling than Sham aSTR in the gain-loss context (**c**). There was, however, no difference to baseline or in the gain-only context. (**d**). **e,f.** TUS in ACC reduced post-learning accuracy relative to baseline and other TUS conditions in the gain-loss context (**e**), but not the gain-only context (**f**). For detailed statistics see **Supplementary Table 1**. Data show session means for the pre-stimulation baseline blocks (*grey*) and across blocks after TUS/Sham (*colored*). Significant pairwise comparisons are indicated with black horizontal lines between each pair of conditions and the black asterisk shows significant session-level difference (post TUS versus baseline) of each behavioral measure.

Extended Data Fig. 9. Trial-level effects of transcranial ultrasound stimulation on post-outcome performance adjustment. a-d. Normalized performance accuracy in the five trials after receiving either a loss of one token (**a**), no loss after an incorrect response (**b**), a gain of 2 tokens (**c**), and a gain of 3 tokens (**d**). TUS in ACC (**red**) caused overall reduced performance accuracy in each monkey after losing 1 token or gaining 2 tokens which shows the effect is specific to the gain-loss context. **e.** Both monkeys show reduced performance (y-axis) after TUS in ACC when the Gross token income (GTI) was negative, i.e. then they on average had lost 1-3 tokens in the preceding five trials. Accuracy on the trial-level analysis is normalized by the mean and standard deviation of the accuracy of similar events in the baseline during the same TUS session. All statistics in these panels used randomization (permutation) tests with FDR correction of p-values for dependent samples with an alpha-level of 0.05. The black asterisk shows points significantly different from pre-TUS baseline and relative to other TUS conditions.

Supplementary Table 1. Statistical results considering data from individual blocks (block level, left) and average data from individual sessions (session level: right) for trials-to-criterion (**a**,**b**), explorative sampling durations (**c**,**d**), and the post-learning plateau accuracy (**e**,**f**). Each table shows results of 3 different tests for the 2 motivational learning contexts: gain-only (G, in blue font), and gain-loss (GL, in red font). The p-values on the diagonal show Wilcoxon tests for each TUS condition compared to its baseline (before the stimulation). The non-diagonal cells show p-values for pairwise comparisons of each pair of the TUS conditions. The overall Kruskal-Wallis test results is shown on the right outside of each table. All p-values are FDR-corrected with an alpha level of 0.05, as explained in the methods.

Supplementary Table 2. Statistical results split for low (left) and high (right) attentional load conditions for trials-to-criterion (**a,b**), explorative sampling durations (**c,d**), and the post-learning plateau accuracy (**e,f**). Each table shows results of 3 different tests for the 2 motivational learning contexts: gain-only (G, in blue font), and gain-loss (GL, in red font). The p-values on the diagonal show Wilcoxon tests for each TUS condition compared to its baseline (before the stimulation). The non-diagonal cells show p-values for pairwise comparisons of each pair of the TUS conditions. The overall Kruskal-Wallis test results is shown on the right outside of each table. All p-values are FDR-corrected with an alpha level of 0.05, as explained in the methods.

Methods

Experimental procedures. All procedures were in accordance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals, the Society for Neuroscience Guidelines and Policies, and approved by the Vanderbilt University Institutional Animal Care and Use Committee. Two male macaque monkeys (monkey I 13.6 kg, and monkey W 14.6 kg, 8-9 years of age) contributed to the experiments. They sat in a sound-proof booth in primate chairs with their head position fixed, facing a 21'' LCD screen at a distance of 63 cm from their eyes to the screen center. Behavior, visual display, stimulus timing, and reward delivery, were controlled by the *Unified Suite for Experiments* (USE), which integrates an IO-controller board with a unity3D video-engine based control for displaying visual stimuli, controlling behavioral responses, and triggering reward delivery[25]. Prior to the ultrasound experiment, the animals were trained on the feature learning task in a kiosk training station[26]. Monkeys first learned choosing objects to receive immediate fluid reward, before a token system was introduced that provided animals with green circles per correct choice that symbolized tokens later cashed out for fluid reward. Tokens were presented above the chosen object and traveled to a token bar where they accumulated with successive correct performance. When five tokens were collected, the token bar blinked red/white, fluid was delivered through a sipper tube and the token bar reset to five empty token placeholders (**Fig. 1a**). The animals effortlessly adopted the token reward system as documented in[15]. Here, we used in separate blocks of 35-50 trials a condition with 'gains-only' (3 tokens for correct choices, no penalties) and with 'gains-and-losses' (2 tokens for correct choices and 1 token lost, i.e. removed from the token bar, for incorrect choices). The introduction of gains-and-losses effectively changed behavior. Animals learned slower, showed reduced plateau accuracy, enhanced exploratory sampling and more checking of the token bar (**Extended Data Fig. 2b-h**).

Task paradigm. The task required monkeys to learn feature-reward rules in blocks of 35-60 trials through trial-and-error by choosing one of three objects. Objects were composed of multiple features, but only one feature was associated with reward (**Fig. 1a-c**). A trial started with the appearance of a black circle with a diameter of 1 cm (0.5° radius wide) on a uniform grey background on the screen. Monkeys fixated the black circle for 150 ms to start a trial. Within 500 ms after registration of the central gaze fixation, three objects appeared on the screen randomly at 3 out of 4 possible locations with an equal distance from the screen center (10.5 cm, 5° eccentricity). Each stimulus had a diameter of 3 cm (~1.5° radius wide). To choose an object, monkeys had to maintain fixation onto the object for at least 700 ms. Monkeys then had 5 sec. to choose one of three objects or the trial was aborted. Choosing the correct object was followed by a yellow halo around the stimulus as visual feedback (500 ms), an auditory tone, and either 2 or 3 tokens (green circles) added to the token bar (**Fig. 1a**). Choosing an object without the rewarded target feature was followed by a blue halo around the selected objects, a low-pitched auditory feedback, and in the loss conditions, the presentation of a grey 'loss' token that traveled to the token bar where one already attained token was removed. Timing of the feedback was identical for all types of feedback. In each session, monkeys were presented with up to 36 separate learning blocks, each with a unique feature-reward rule. Across all 48 experimental sessions monkeys completed on average 23.6 (±4 SE) learning blocks per session (monkey W: 20.5±4, monkey I: 26.5±4). For each experimental session a unique set of objects was defined by randomly selecting three dimensions and three feature values per dimension (e.g. 3 body shapes: oblong, pyramidal and ellipsoid; 3 arm types: upward pointy, straight blunt, downward flared; 3 patterns: checkerboard, horizontal striped, vertical sinusoidal; Watson et al.[27] have documented the

complete library of features). Of this feature set three different task conditions were defined: One task condition contained objects that varied in only one feature while all other features were identical, i.e. the object body shapes were oblong, pyramidal and ellipsoid, but all objects had blunt straight arms and uniform grey color. A second task condition defined objects that varied in two feature dimensions ('2D' condition), and a third task condition defined objects that varied in three feature dimensions ('3D' condition). Learning is systematically more demanding with increasing number of feature dimensions that could contain the rewarded feature (for computational analysis of the task, see[28]). We refer to the variations of the object feature dimensionality as attentional load because it corresponds to the size of the feature space subjects have to search to find the rewarded feature (**Figure 1e**, **Extended Data Fig. 1**). 1D, 2D, and 3D conditions varied randomly across blocks.

Experimental design. Each session randomly varied across blocks two motivational conditions (Gain/Loss and Gains-only) and three attentional load conditions (1D, 2D, 3D). Randomization ensured that all (six) combinations of conditions were present in the first six blocks prior to sonication and that all combinations of conditions were equally often shown in the 24 blocks after the first six blocks. After monkeys completed the first 6 learning blocks (on average 12 minutes after starting the experiment), the task was paused to apply transcranial ultrasound stimulation. After bi-lateral placement of the transducer and sonication (or sham sonication) of the brain areas the task resumed for 18-24 more blocks (**Fig. 1d**). Experimental sessions lasted 90-120 minutes.

Transcranial Ultrasound Stimulation (TUS). For transcranial stimulation, we used a single element transducer with a curvature of 63.2 mm and an active diameter of 64 mm (H115MR, Sonic

Concepts, Bothell, WA). The transducer was attached to a cone with a custom-made trackable arm. Before each session, we filled the transducer cone with warm water and sealed the cone with a latex membrane. A conductive ultrasound gel was used for coupling between the transducer cone and the shaved monkey's head. A digital function generator (Keysight 33500B series, Santa Rosa, CA) was used to generate a periodic burst of 30 ms stimulation with a resonate frequency of 250 kHz and an interval of 100-ms for a total duration of 40-second and 1.2 MPa pressure (similar to [17,29]) (**Fig. 1d**). A digital function generator was connected to a 150-Watt amplifier with a gain of 55 dB in the continuous range to deliver the input voltage to the transducer (E&I Ltd., Rochester, NY). We measured the transducer output in water using a calibrated ceramic needle hydrophone (HNC 0400, Onda Corp., Sunnyvale CA) and created a linear relationship between the input voltage and peak pressure. To avoid hydrophone damage, only pressures below a mechanical index (MI) of 1.0 were measured and amplitudes above this were extrapolated. We have previously measured transducer output at MI>1.0 with a calibrated optical hydrophone (Precision Acoustics, Dorchester, UK) to validate the linearity of this relationship at higher MI, but this calibrated device was not available during these studies. During stimulation, the bi-directionally coupled (ZABDC50-150HP+, Mini Circuits Brooklyn, NY) feedforward and feedback voltage were monitored and logged using a Picoscope 5000 series (A-API; Pico Technology, Tyler, TX) and a custom written python script.

Four different sonication conditions were pseudo-randomly assigned to the experimental days per week for a 12-week experimental protocol per monkey. These four conditions consisted of high energy TUS in anterior striatum (H-aSTR), high energy TUS in anterior cingulate cortex (H-ACC), sham anterior striatum (S-aSTR), and sham anterior cingulate cortex (S-ACC) (**Fig. 1d**). We

sequentially targeted an area in both hemispheres (each hemisphere for a 40 sec duration) with real-time monitoring of the distance of the transducer to the targeted area (**Extended Data Fig. 3a**) and monitoring of the feedforward power (**Extended Data Fig. 3b**). Sham conditions were identical to TUS conditions, only no power was transmitted to the transducer.

Neuro-navigation setup. An MR image was acquired before implantation of the head-post and a CT image was obtained containing the headpost with a fiducial array arm attached to it during the scan. The fiducial array in the CT scan was identical to its position during the neuro-navigation procedure in an experimental session. We co-registered MR and CT images into the same space (CT images were the fixed volume to prevent distortion of fiducial array space) using FSL[30]. 3D-Slicer software (http://www.slicer.org), with Image-guided therapy (IGT)[31] extension modules were used for tracking and positioning the transducer[2]. First, the co-registered brain images were imported to 3D-Slicer. Fiducial array points were marked on the images ordered from right to left. These points allowed reconstructing the relative distances of the sonication target locations in the brain from the fiducial array. At the beginning of each session, the fiducial array was mounted to monkey's headpost and fixed to the same location it had in the CT images. A stylus tool was then used to collect the physical fiducial array points in the same order as it was marked in the image space. After collecting the physical points, we transformed the physical space into image space. The focus to tracker transformation (which was used for tracking the focus of the transducer relative to the tracker attached to its enclosing cone) was also transformed with the same transformation to the image space[32,33]. We then navigated the focal point of the transducer (a circle with 3 mm radius) centered at the maximum focus estimated as sphere model (1-mm radius) placed at the focal maximum that was used for the location of targets in each area and hemisphere (these

sphere models were held the same across all experiments for each monkey). Under the IGT module, we tracked the real-time distance of the center of the transducer focus from the target sphere model.

Acoustic modeling of sonication beam. To validate sonication of the ACC and aSTR we performed 3D numerical simulations of the acoustic propagation through the monkey skull using the open-source MATLAB toolbox k-Wave (Http://www.k-wave.org). A single-element transducer with a radius of curvature of 63.2 mm and an active diameter of 64 mm (H115MR, Sonic Concepts, Bothell, WA) and resonant frequency of 250 kHz was modeled as a spherical section. To position the transducer for simulations, we used 3D-Slicer combined with optical tracking to recreate the optical tracking geometry. To localize the transducer relative to the transducer tracker we manually aligned a transducer mask's marked geometric focus created in k-Wave with the center of the focal location from reconstructed temperature images collected using a 2D gradient echo thermometry pulse. MR thermometry was performed on a 3.0T Pihlips Achieva 2D gradient echo thermometry pulse sequence with multi-shot EPI readout factor of 3. Imaging parameters were 150 x 150 mm^2 field of view (FOV), 112 x 112 matrix, 1.3 x 1.3 mm^2 voxel size, 5 slices, 4 mm slice thickness, TE/TR 12/500 msec. A continuous wave sonication for 30 sec. at input voltage of approximately 1 MPa was used to generate several degrees of heating. This calibration was performed once to generate a transform with an accurate offset and rotation relative to the optical tracker. A CT image of the respective monkey was resampled using the Resample Image (BRAINS) module in Slicer so that the resolution and spacing matched the calibrated transducer mask or integration with simulations. The resampled data was exported to MATLAB to run acoustic simulations, using the CT to estimate the speed of the sound and density of the

skull. The final spatial maps were imported back into 3D Slicer and overlaid for comparison with the estimated focus from the optical tracking and MR thermometry results. To account for heterogeneities of the monkey skull, a linear relationship between Hounsfield units (HU) from CT scans was used to calculate the speed of sound and density[34]. Brain tissue was assumed to have the same speed of sound and density as water. Values used included a padded grid size [Nx,Ny,Nz] of [384,288,280] with isometric voxels of 0.25 mm, attenuation = 8 dB/cm, power law absorption exponent of 1.1, $c_{min} = 1480 \ m/s$, $c_{max} = 3100 \ m/s$, $\rho_{water} = 1000 \ kg/m^3$, and $\rho_{bone} = 2100 \ kg/m^3$. Maximum pressure was recorded for every voxel in the simulation grid. The highest pressure was found at the skull, as it absorbs more sound than brain tissue. For each dataset, the spatial map cropped the maximum pressure of the skull so that an accurate root-mean-square (RMS) of the sound pressure could be calculated at the focus. The focal spot size of the sonicated area was calculated using $P > P_{max}/2$ for half-maximum pressure (-3 dB; **Extended Data Fig. 3d**) and $P^2 > P^2_{max}/2$ for half maximum intensity (-6 dB; **Extended Data Fig. 3e,** error bars showing standard error of the mean across sessions). The half-maximum pressure and intensity equations were used to create a mask of the focal spot volume and calculate the RMS pressure for each (**Fig. 2a**).

Analysis of behavioral learning and fixational sampling. The improvement of accuracy over successive trials relative to the beginning of a learning block reflects the learning curve, which we computed with a forward-looking 12-trial averaging window in each block (**Fig. 1e**). We defined learning speed in a block as the number of trials needed to reach criterion accuracy of ≥80% correct trials over the next 12 successive trials (**Fig. 2b,c**; **Extended Data Fig. 1c**, **Extended Data Fig. 2c**). When monkeys did not reach the learning criterion at the end of a block, we estimated the trial

number the monkeys would have learned the block with a linear regression fit to the performance accuracy of monkeys in the last 12 trials in the block. We computed the plateau performance accuracy, or *post-learning accuracy*, as the proportion of correct choices across all trials after the learning criterion was reached in a block (**Fig. 3c,d**). In blocks in which the learning criterion was not reached we computed the accuracy over the last 12 trials in the block. Analysis of exploratory eye movements used an adaptive velocity-based thresholding algorithm[35,36] to detect saccadic eye movements and fixations (*i*) onto objects during explorative sampling prior to choosing an object (*explorative sampling*), (*ii*) onto the chosen object prior to choosing it (*choice fixations*), and (*iii*) onto the token bar (*asset sampling*) (see e.g. **Extended Data Fig. 1e,g,h** and **Extended Data Fig. 2e,g,h**).

Trial-level statistical analysis. We tested TUS effects on behavior at the trial level using linear mixed effect models (LME)[37] with 4 main factors: *attentional load (Att_{Load})* with three levels (1D, 2D, and 3D distractor feature dimensions, ratio scale with values 1,2 and 3), *trial in block (TIB)*, *previous trial outcome* ($Prev_{Outc)}$) which is the number of tokens gained or lost in the previous trial, motivational token condition, which we call the *motivational Gain/Loss context ($MCtx_{Gain/Loss}$)* with two levels (1, for the loss condition, and 2 for the gain condition, nominal variable), TUS condition (TUS_{Cnd}) with four levels (Sham-ACC, High-ACC, Sham-STR, High-STR), and *time relative to stim (T2Stim)* with 2 levels (before versus after stimulation). We used 3 other factors as random effects, a factor *target features* (*Feat*) with 4 levels (color, pattern, arm, and shape), weekday of the experiment (*Day*) with 4 levels (Tuesday, Wednesday, Thursday, and Friday), and the factor *monkeys* with 2 levels (W and I). We used these factors to predict 3 metrics (*Metric*):

accuracy (*Accuracy*), reaction time (*RT*), explorative sampling ($Sample_{Explr}$). The LME is formalized as in eq. 1.

$$Metric = Att_{Load} + TIB + Prev_{Outc} + MCtx_{Gain/Loss} + TUS_{Cnd} + T2Stim + (1|Day) +$$

$$(1|Feat) + (1|Monkeys) + b + \varepsilon \qquad \qquad \text{(eq. 1).}$$

Analysis of behavioral adjustment to the recent token history (Gross Token Income). To quantify how TUS affected choice accuracy depending on the motivational status of the subjects, we calculated the gross sum of the earned token over four trials which we call *gross token income* (GTI; **Fig. 3f**, **Extended Data Fig. 9e**). For each value of the GTI (spanning from – 3 to +6 tokens), we calculated the accuracy on the subsequent trial and normalized it by subtracting mean and dividing by the standard deviation of the same GTI values during the baseline. We then asked whether in any of the TUS conditions accuracy of choice for a given GTI changed relative to baseline or other TUS conditions. We first used Wilcoxon signed-rank test to test for each TUS condition and each GTI, whether accuracy changes relative to baseline were different from zero. We then used FDR control of multiple comparison for dependent samples[38], with an alpha level of 0.05 to adjust p-values across different GTIs and the TUS conditions. In a second analysis, we applied randomization statistics to test whether for a given GTI accuracy changed relative to other TUS conditions. We randomly permuted trial labels of the TUS conditions and selected 1000 subsamples with a sample size of equal to the true TUS and GTI conditions. We then formed a randomized sampling distribution of the means and extracted the probability of randomly finding values more extreme than the true baseline-corrected accuracy value for each TUS and GTI conditions, under the null hypothesis that for a given GTI none of the TUS conditions show accuracy changes relative to other TUS conditions. After finding the p-values for all GTI

conditions, we corrected them for multiple comparisons using FDR control at an alpha level of 0.05. GTI values for a TUS condition were considered significant if it passed both tests, significantly different from the baseline and from other TUS conditions (**Fig. 3f**, **Extended Data Fig. 9e**).

Analysis of previous-trial-outcome effects on accuracy adjustment. We measured how accuracy changed 1-5 trials after experiencing gaining or loosing tokens and tested how TUS modulated this previous trial outcome effect (**Fig. 3e**, **Extended Data Fig. 9a-d**). Accuracy was baseline normalized for each TUS condition. Wilcoxon rank tests were applied to test the effects of the previous trial outcome after versus before TUS with FDR correction for dependent samples[38] with an alpha level of 0.05. In a next step, we tested post-outcome accuracy in one TUS condition to other TUS conditions. Thus, for each outcome condition, we randomly permuted the TUS condition label of the trials, and randomly selected 1000 subsamples with a size equal to the true outcome and the TUS condition (both in here and in the GTI analysis, only the TUS labels were shuffled and the rest of the labels for trials e.g., the order relative to TUS was still the same). We then used the 1000 subsamples and formed a randomized sampling distribution of the means and calculated the probability of finding a value more extreme than the true mean of changed trials accuracy following an outcome relative to the baseline, under the null hypothesis that the accuracy changes relative to the baseline after a given trial outcome is not different across different TUS conditions. After finding the p-values for each TUS condition and n^{th} (n=1-5) following trials, we used FDR correction for dependent samples with an alpha level of 0.05 to adjust p-values across different post outcome trials and the TUS conditions.

Block-level analysis of behavioral metrics. We used linear mixed effect models to analyze across blocks how learning speed (indexed as the 'learning trial' (LT) at which criterion performance was reached), post-learning accuracy (*Accuracy*), and metrics for explorative and choice sampling were affected by TUS. In addition to the fixed and random effects factors used for the LME in eq. 1 we also included the factor block *switching condition* ($Switch_{Cnd}$) with 2 levels (intra- and extra-dimensional switch). The LME had the form of eq. 2:

$Metric = Att_{Load} + MCtx_{Gain/Loss} + TUS_{Cnd} + T2Stim + (1|Day) + (1|Monkeys) + (1|Feat) + b + \varepsilon$

(eq. 2).

We extended the model to test for interactions of TUS_{Cnd}, $MCtx_{Gain/Loss}$, and $T2Stim$:

$Metric = Att_{Load} + MCtx_{Gain/Loss} + MCtx_{Gain/Loss} \times TUS_{Cnd} \times T2Stim + (1|Day) + (1|Monkey) + (1|Feat) + b + \varepsilon$ (eq. 3);

for interactions of $T2Stim$, Att_{Load} and TUS_{Cnd}:

$Metric = Att_{Load} + MCtx_{Gain/Loss} + TUS_{Cnd} \times T2Stim \times Att_{Load} \times TUS_{Cnd} + (1|Day) + (1|Monkey) + (1|Feat) + b + \varepsilon$ (eq. 5);

and for interactions of $T2Stim$, Att_{Load}, $MCtx_{Gain/Loss}$ and TUS_{Cnd}:

$Metric = Att_{Load} + MCtx_{Gain/Loss} + TUS_{Cnd} \times T2Stim \times Att_{Load} \times TUS_{Cnd} \times Frame_{Cnd} + (1|Day) + (1|Monkey) + (1|Feat) + b + \varepsilon$ (eq. 6).

Where appropriate we used a non-parametric approach by fitting generalized linear mixed effects models (GLME's) as in equations 2-6. While the formulas were the same as in LME's, we used a link *Identity* function, and a *Poisson* distribution for the response variable.

Pairwise comparisons of TUS effects. As a validation of the overall mixed effect modeling results, we report a Kruskal Wallis test for a main effect of TUS conditions on behavioral metrics and follow up with Wilcoxon tests for all pairwise comparisons between individual TUS conditions at each motivational context (gains-only and gains-loss). P-values were FDR corrected to control for multiple comparisons of dependent samples[38]. Complete statistical results are provided in **Supplementary Table 1**. We repeated the same procedure for pairwise comparisons of TUS effects at low attentional load (1D condition) and high attentional load (combining 2D and 3D conditions) (**Supplementary Table 2**).

Randomization statistics for normalized behavioral metrics. We calculated the normalized value of a behavioral metric across blocks in each TUS condition, motivational contexts, and attentional loads by subtracting their marginal means and dividing by the standard deviation. As shown in **Fig. 2d** and **Extended Data Fig. 5c** this procedure provides a table where each cell shows the mean of the marginally normalized metric value ± the standard error of normalized metric value (for each cell, the color shows from left to right the mean minus the standard error to the mean plus the standard error of that cell). We then performed randomization tests at two levels. First, at the level of the motivational contexts, we asked whether any TUS condition significantly differed from others by permuting the TUS condition labels and randomly selecting 1000 subsamples with a size equal to the size of the true TUS condition. This tested the null hypothesis

that the marginally normalized metric values in a given TUS condition in a given motivational context do not differ from other TUS conditions in the same motivational context. We calculated the p-values and adjusted them post-hoc using FDR correction for dependent samples with an alpha level of 0.05. The white rectangles in **Fig. 2d** and **Extended Data Fig. 5c** show the FDR corrected, significant TUS. In a second analysis, we applied the same rationale and procedure but this time the randomization test was done on each motivational context and attentional load condition separately to test the null hypothesis that TUS conditions had no effect on the behavioral metrics. Those TUS conditions that showed a significant difference from others in a given attentional load and motivational context are marked with a white asterisk in **Fig. 2d** and **Extended Data Fig. 5c**.

Session-Level analysis of TUS effects on behavioral metrics. In addition to analyzing behavior in individual blocks and performing block-level statistics (see above), we also analyzed the data across sessions and provide these session-level results in **Supplementary Table 1** and **Extended Data Fig. 8**. For each session, we calculated the mean of each behavioral metric for gain-only and gain-loss learning contexts separately before and after the TUS. We then used Wilcoxon tests for each TUS condition under the null hypothesis that the behavioral metric is not different before versus after TUS. To compare behavioral metrics across TUS conditions we first applied a Kruskal-Wallis test under the null hypotheses that the behavioral metrics did not differ between TUS conditions. Then we used pairwise Wilcoxon comparisons to compare each pair of TUS conditions. All p-values were corrected post-hoc using FDR for dependent samples with an alpha level of 0.05

References

25. Watson, M. R., Voloh, B., Thomas, C., Hasan, A. & Womelsdorf, T. USE: An integrative suite for temporally-precise psychophysical experiments in virtual environments for human, nonhuman, and artificially intelligent agents. *Journal of Neuroscience Methods* **326**, 108374 (2019).

26. Womelsdorf, T. *et al.* A Kiosk Station for the Assessment of Multiple Cognitive Domains and Enrichment of Monkeys. *bioRxiv* 2021.03.06.434198 (2021) doi:10.1101/2021.03.06.434198.

27. Watson, M. R., Voloh, B., Naghizadeh, M. & Womelsdorf, T. Quaddles: A multidimensional 3-D object set with parametrically controlled and customizable features. *Behav Res* **51**, 2522–2532 (2019).

28. Womelsdorf, T., Watson, M. R. & Tiesinga, P. Learning at variable attentional load requires cooperation between working memory, meta-learning and attention-augmented reinforcement learning. *bioRxiv* 2020.09.27.315432 (2020) doi:10.1101/2020.09.27.315432.

29. Khalighinejad, N. *et al.* A Basal Forebrain-Cingulate Circuit in Macaques Decides It Is Time to Act. *Neuron* **105**, 370-384.e8 (2020).

30. Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W. & Smith, S. M. FSL. *NeuroImage* **62**, 782–790 (2012).

31. Ungi, T., Lasso, A. & Fichtinger, G. Open-source platforms for navigated image-guided interventions. *Medical Image Analysis* **33**, 181–186 (2016).

32. Chaplin, V. *et al.* On the accuracy of optically tracked transducers for image-guided transcranial ultrasound. *Int J CARS* **14**, 1317–1327 (2019).

33. Kim, H., Chiu, A., Park, S. & Yoo, S.-S. Image-guided navigation of single-element focused ultrasound transducer. *International Journal of Imaging Systems and Technology* **22**, 177–184 (2012).

34. Aubry, J.-F., Tanter, M., Pernot, M., Thomas, J.-L. & Fink, M. Experimental demonstration of noninvasive transskull adaptive focusing based on prior computed tomography scans. *The Journal of the Acoustical Society of America* **113**, 84–93 (2003).

35. Nyström, M. & Holmqvist, K. An adaptive algorithm for fixation, saccade, and glissade detection in eyetracking data. *Behavior Research Methods* **42**, 188–204 (2010).

36. Voloh, B., Watson, M. R., Konig, S. & Womelsdorf, T. MAD saccade: statistically robust saccade threshold estimation via the median absolute deviation. *JEMR* **12**, (2019).

37. Pinheiro, J. C. & Bates, D. M. Unconstrained parametrizations for variance-covariance matrices. *Stat Comput* **6**, 289–296 (1996).

38. Benjamini, Y. & Yekutieli, D. False Discovery Rate–Adjusted Multiple Confidence Intervals for Selected Parameters. *Journal of the American Statistical Association* **100**, 71–81 (2005).