

Temporal structure of learning to regulate ventral tegmental area using real-time fMRI neurofeedback

Shabnam Hakimi^{†1}

Jeffrey J. MacInnes^{†*2}

Kathryn C. Dickerson^{†‡3}

R. Alison Adcock^{†‡4}

[†] Center for Cognitive Neuroscience, Duke University, Durham, NC USA

[‡] Psychiatry and Behavioral Sciences, Duke University, Durham, NC USA

^{*} Institute for Learning & Brain Sciences, University of Washington, Seattle, WA USA

Abstract

The ventral tegmental area (VTA) and its dopaminergic projections are central to volitional behavior. Previous research from our group demonstrated that individuals can use real-time neurofeedback training to learn to reliably self-activate the VTA using self-generated motivational imagery (MacInnes, Dickerson, Chen, & Adcock, 2016). The mechanism of learning, however, is not yet known. Here, we investigated how the temporal structure of neurofeedback training impacts successful transfer of VTA self-activation. We analyzed veridical VTA neurofeedback during self-activation trials in one of three temporal contexts (individual trial, scanning run, or full training session) to test the extent to which slope of VTA response over time during each context explains change in self-activation ability. A comparison of the model evidence suggested that, relative to trial and run, the full session best explained the magnitude of transfer from pre- to post-training ($p < 0.01$). These preliminary data suggest that the overall training context may be a better predictor of learning from VTA neurofeedback than individual training episodes—regardless of their success.

Keywords: fMRI; neurofeedback; VTA; learning; motivation

Introduction

Biofeedback is a powerful tool for providing individuals insight into physiological states, creating a salient context for learning. Recent advances in neurofeedback techniques, such as real-time functional magnetic resonance imaging (rtfMRI) neurofeedback, offer new opportunities for shaping individual behavior using non-invasive interventions, although learning mechanisms are poorly understood (Sulzer et al., 2013). Through 'cognitive neurostimulation' (CN), or the use of self-generated thoughts and imagery to non-invasively engage neuromodulatory systems, individuals are able to use rtfMRI neurofeedback to learn to volitionally sustain activation of dopaminergic circuits (MacInnes et al., 2016). With this technique, individuals gain access to the neurobiological substrates of motivation and volitional behavior (Salamone & Correa, 2012; Jahanshahi, 1998), using this information to

learn to increase activation of the VTA as well as mesolimbic functional connectivity *in the absence of external rewards* (MacInnes et al., 2016). This new insight into the biological basis of motivational experience may be particularly impactful for translational efforts for enhancing individual health and well-being, as volitional, self-generated motivational states are closely linked to therapeutic success (Ryan & Deci, 2008) and may be required (Gneezy, Meier, & Rey-Biel, 2011; Bandura, 1991) for sustainable behavior change (Schwarzer, 2001). Understanding *how* individuals learn to use VTA neurofeedback, therefore, is critical for establishing the efficacy of CN and developing it further as an intervention. In the present study, we investigated this question by testing the extent to which the neurofeedback context itself—rather than any specific strategy—predicted learning.

Methods

19 right-handed individuals (9 female, mean age = 24) with no history of psychological or neurologic illness or MRI contraindications participated in this experiment as part of a larger study, which was approved by the institutional review board of Duke University. Participants and full experimental parameters are detailed in (MacInnes et al., 2016).

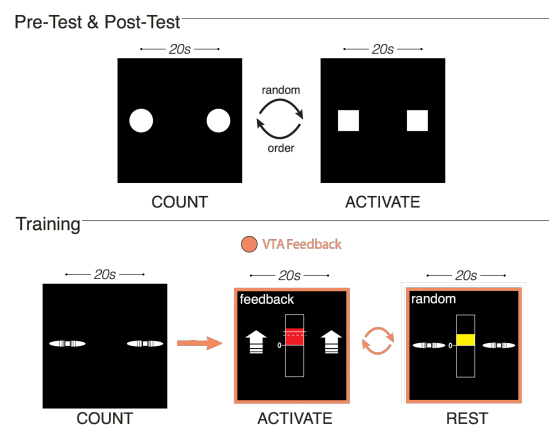


Figure 1: Cognitive neurostimulation task.

Neuroimaging of the CN Task The CN task consisted of two types of runs: test (Pre-Test and Post-Test) and training (as in Figure 1); only training runs included rtfMRI neurofeedback. Briefly, test runs were identical and included two

¹shabnam.hakimi@duke.edu

²jeffmaci@uw.edu

³kathryn.dickerson@duke.edu

⁴alison.adcock@duke.edu

trial types: activate and count. In activate trials, participants were instructed to get themselves into a heightened mental state of motivation using self-generated positive phrases and personally-relevant imagery or memories. Importantly, individuals were told to use only one strategy per trial and were provided with minimal guidance about how to sample motivational imagery/strategy space. In count trials, participants counted backward from 300 in increments of four. Each trial type was repeated five times and trial order was randomized. The training runs included three trial types: activate, count, and rest. For activate and count, instructions were the same as in the test runs, although activate trials now included a dynamic graphical thermometer display that presented real-time neurofeedback from the VTA. During rest, participants were instructed to rest and not think of anything in particular and saw a graphical thermometer that displayed random feedback. Again, each 20-s trial type was repeated five times in each of three runs (15 total trials per type). In both test and training, all trials had a duration of 20 s and were separated by a jittered ITI.

MR imaging data were collected using a GE MR750 3T scanner; physiological data (heart rate and respiration) were also collected during imaging. Functional imaging data were acquired using an echo-planar imaging sequence with partial-brain acquisition (18 oblique slices oriented to the AC-PC line) using the following parameters: TR = 1 s, TE = 28 ms, 90° flip angle, 3 x 3 x 3.8 mm voxel size. One 30-s resting state run was also acquired using these parameters; this functional scan was used with a high-resolution, whole-brain, T1-weighted structural scan (1 mm³ voxel size) to make participant-specific VTA masks using a probabilistic atlas (Murty et al., 2014; Ballard et al., 2011). rtfMRI data were reconstructed and analyzed in real-time using Pyneal (<https://github.com/jeffmacinnes/pyneal>), which output the weighted mean BOLD response from the subject-specific mask. This mean VTA response was presented to participants as neurofeedback during activate trials only, updating at a frequency of 1 Hz. Offline analyses of all other fMRI data were conducted using FSL v5.0.1 (<http://www.fmrib.ox.ac.uk/fsl>) using standard pipelines along with physiological noise correction. For each participant, the outcome of interest—training-mediated change in VTA self-activation ability from Pre-Test to Post-Test, or transfer—was computed using preprocessed, denoised, masked data for the contrast, [activate - count]. As noted, these analyses are described in detail elsewhere (MacInnes et al., 2016).

Modeling Approach

Here, we were interested in investigating the information provided by the neurofeedback signal itself over different temporal contexts and its ability to predict transfer. We defined three natural temporal contexts based on task design and experimental demands: individual trial, scanning run, and full training session. Distinct events were nested within run (3) and trial (15). These nested temporal contexts are illustrated in

Figure 2.



Figure 2: Temporal contexts in cognitive neurostimulation. Participants completed a Pre-Test, CN training session, and Post-Test. Training temporal contexts are shown as shaded boxes—CN training session in grey, the three scanner runs in purple, and the 15 individual trials in yellow—overlaid on a sample neurofeedback signal in blue. Contexts are numbered when appropriate; T='TRIAL'.

Critically, participants were provided with minimal instruction about how to use motivational imagery to self-activate and participant strategies were not reported during CN training. Thus, the analytical approach was blind to the strategies used and how each individual did—or did not—explore his or her unique motivational strategy space. Instead, we focus on how the temporal structure of CN impacted learning to self-activate the VTA, independent of the idiosyncratic content of each individual's learning experience.

To examine the effect of temporal structure, we extracted the activate neurofeedback signal at the single-participant level (i.e., the feedback each participant received), resampled it to 2 Hz, and rescaled it to [-1,1] using MATLAB 2016b (The MathWorks, Inc.). The analysis then proceeded in two stages. First, we computed the slope of the neurofeedback time series for each context using MATLAB `polyfit`, using the common assumption in the neurofeedback literature that learning is linear in time (e.g., (Lawrence et al., 2014)). This produced a set of parameters reflecting learning during each context (trial: 15, run: 3, session: 1). Next, we used linear regression to examine how well the parameters for each context predicted CN transfer, fitting one model per temporal context on the group-level with MATLAB `fitlm`. Finally, we estimated model evidence for each of the three models using two measures that penalized overfitting: the small-sample size equivalent of Akaike Information Criterion (AICc; (Akaike, 1974)) and the Bayesian Information Criterion (BIC; (Schwarz, 1978)).

Results

Initial model comparison suggests that a linear trend in neurofeedback signal at the session level is most predictive of transfer of CN training (adjusted $R^2 = 0.335$, $F_{2,17} = 8.56$, $p < 0.01$), as described in Table 1.

While AICc and BIC both supported training session as the best model, the differential penalties imposed by the two parameters provided conflicting evidence about the relative abilities of individual trial and scanner run to predict transfer.

Table 1: Model evidence by temporal context.

Temporal Context	AICc	BIC
Individual trial	243.0	-13.88
Scanner run	10.82	11.74
Training session	-2.584	-1.445

Discussion

We investigated whether the temporal structure of the rtfMRI neurofeedback environment predicts transfer of training to self-activate the VTA. We found initial evidence that successful learning from VTA neurofeedback is associated with the training context as a whole, although there may be a role for shorter, trial-level changes. These data are consistent not only with the function of dopamine in adaptive memory (Shohamy & Adcock, 2010), but also with the idea that the biofeedback episode provides a novel, motivationally salient context (MacDuffie & Strauman, 2017) that may support learning through an increase in metacognitive awareness (MacDuffie et al., 2018).

Our analytical approach was limited by a relatively small sample size, resulting in insufficient statistical power to test the effect of temporal context within one mixed model. Further, data reflected significant variability both within- and across-participants that may be driven by the nature and number of motivational strategies used. Ongoing work will interrogate 1) the effect of the linearity assumption by using data driven approaches (e.g., Gaussian process models) and 2) the specificity of these findings to veridical VTA neurofeedback learning by also examining two other feedback conditions, random noise and patterned false feedback. Additional work in larger samples is needed to better resolve the temporal dynamics of VTA neurofeedback learning, especially how multiple learning signals may emerge and interact over different time scales.

Acknowledgments

This work was generously supported by NIH grants to SH (F32 MH114608), KD (F32 MH100764), and RAA (R01 MH094743); additional support to RAA was provided by the Alfred P. Sloan Foundation, the Esther A. & Joseph Klingenstein Fund, and the Dana Foundation.

References

Akaike, H. (1974, Dec). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*(6), 716-723. doi: 10.1109/TAC.1974.1100705

Ballard, I. C., Murty, V. P., Carter, R. M., MacInnes, J. J., Huettel, S. A., & Adcock, R. A. (2011, Jul). Dorsolateral prefrontal cortex drives mesolimbic dopaminergic regions to initiate motivated behavior. *Journal of Neuroscience*, *31*(28), 10340–10346. doi: 10.1523/jneurosci.0895-11.2011

Bandura, A. (1991). Social cognitive theory of self-regulation. *Organizational Behavior and Human Decision Processes*,

50(2), 248 - 287. (Theories of Cognitive Self-Regulation) doi: [https://doi.org/10.1016/0749-5978\(91\)90022-L](https://doi.org/10.1016/0749-5978(91)90022-L)

Gneezy, U., Meier, S., & Rey-Biel, P. (2011, December). When and why incentives (don't) work to modify behavior. *Journal of Economic Perspectives*, *25*(4), 191-210. doi: 10.1257/jep.25.4.191

Jahanshahi, M. (1998, Sep). Willed action and its impairments. *Cognitive Neuropsychology*, *15*(6-8), 483–533. doi: 10.1080/026432998381005

Lawrence, E. J., Su, L., Barker, G. J., Medford, N., Dalton, J., Williams, S. C., ... others (2014). Self-regulation of the anterior insula: reinforcement learning using real-time fmri neurofeedback. *Neuroimage*, *88*, 113–124.

MacDuffie, K. E., MacInnes, J., Dickerson, K. C., Eddington, K. M., Strauman, T. J., & Adcock, R. A. (2018). Single session real-time fmri neurofeedback has a lasting impact on cognitive behavioral therapy strategies. *bioRxiv*. doi: 10.1101/258095

MacDuffie, K. E., & Strauman, T. J. (2017, Mar). Understanding our own biology: The relevance of auto-biological attributions for mental health. *Clinical Psychology: Science and Practice*, *24*(1), 50–68. doi: 10.1111/cpsp.12188

MacInnes, J. J., Dickerson, K. C., Chen, N., & Adcock, R. A. (2016). Cognitive neurostimulation: Learning to volitionally sustain ventral tegmental area activation. *Neuron*, *89*(6), 1331 - 1342. doi: <https://doi.org/10.1016/j.neuron.2016.02.002>

Murty, V. P., Shermohammed, M., Smith, D. V., Carter, R. M., Huettel, S. A., & Adcock, R. A. (2014). Resting state networks distinguish human ventral tegmental area from substantia nigra. *NeuroImage*, *100*, 580 - 589. doi: <https://doi.org/10.1016/j.neuroimage.2014.06.047>

Ryan, R. M., & Deci, E. L. (2008). A self-determination theory approach to psychotherapy: The motivational basis for effective change. *Canadian Psychology/Psychologie canadienne*, *49*(3), 186.

Salamone, J. D., & Correa, M. (2012). The mysterious motivational functions of mesolimbic dopamine. *Neuron*, *76*(3), 470 - 485. doi: <https://doi.org/10.1016/j.neuron.2012.10.021>

Schwarz, G. (1978, Mar). Estimating the dimension of a model. *The Annals of Statistics*, *6*(2), 461–464. doi: 10.1214/aos/1176344136

Schwarzer, R. (2001). Social-cognitive factors in changing health-related behaviors. *Current directions in psychological science*, *10*(2), 47–51.

Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends in cognitive sciences*, *14*(10), 464–472.

Sulzer, J., Haller, S., Scharnowski, F., Weiskopf, N., Birbaumer, N., Blefari, M., ... Sitaram, R. (2013). Real-time fmri neurofeedback: Progress and challenges. *NeuroImage*, *76*, 386 - 399. doi: <https://doi.org/10.1016/j.neuroimage.2013.03.033>